

What Really Matters in Designing School Choice Mechanisms*

Parag A. Pathak, MIT and NBER[†]

Draft date: January 2016

Abstract

In the last decade, numerous student assignment systems have been redesigned using input from economists in the large American cities and elsewhere. This article reviews some of these case studies and uses practical experiences to take stock on what has really mattered in school choice mechanism design so far. While some algorithm design details are important, many are less practically important than initially thought. What really matters are basic issues that market operators in other contexts would likely be concerned about: straightforward incentives, transparency, avoiding inefficiency through coordination and well-functioning aftermarkets, and influencing inputs to the design, such as applicant decision-making and the quality of schools.

KEYWORDS: school choice, market design, matching

*Prepared for the 2015 World Congress of the Econometric Society in Montreal. I'm particularly grateful to Federico Echenique, my discussant, and Fuhito Kojima for helpful feedback. I would also like to thank Atila Abdulkadiroğlu, Nikhil Agarwal, Alonso Bucarey, Umut Dur, Arda Gitmez, Peter Hull, Yusuke Narita, Cara Nickolaus, and Tayfun Sönmez for comments and Miikka Rokkanen for his assistance with Tables 4 and 5. Many of the ideas in this paper developed from long-standing collaborations with Atila Abdulkadiroğlu, Alvin Roth, and Tayfun Sönmez, for which I'm grateful. I've also benefitted from interactions with Neil Dorosin, Gabriela Fighetti, and many other school assignment practitioners. Financial support was provided by the National Science Foundation. I'm on the scientific advisory board of the Institute for Innovation in Public School Choice, a non-profit 501(c) that provides assistance to school districts on the implementation of school choice systems.

[†]Department of Economics, Massachusetts Institute of Technology, Cambridge MA 02142 and NBER, e-mail: ppathak@mit.edu

Contents

1	Introduction	3
2	Clearing the Market	4
2.1	Efficiency vs. Fairness	4
2.2	Student Optimality in DA	8
2.3	Constrained Rank Order Lists	10
2.4	Ranking Behavior in New Mechanisms	12
2.4.1	Before and After DA in Boston	13
2.4.2	Chicago	15
3	Considerations from the Field	17
3.1	Straightforward Incentives	17
3.2	Transparency	20
3.2.1	Aversion to TTC	21
3.2.2	Demise of Walk Zones in Boston	22
3.2.3	Seattle’s Setback	23
3.3	Coordinating Offers	25
3.4	Aftermarkets	27
4	Making the Market Work	28
4.1	Families as Consumers	28
4.2	Improving School Quality	29
5	Conclusion	32

1 Introduction

In recent years, there has been a great deal of research activity and excitement among economists who study the design of systems used to assign students to schools. Motivated by Turkish college admissions, Balinski and Sönmez (1999) first defined the student placement problem, and Abdulkadiroğlu and Sönmez (2003) defined the closely related school choice problem, motivated by K-12 public school admissions in the United States. Both articles showed how insights from matching theory could be used to re-engineer and potentially improve existing centralized school assignment systems. Abdulkadiroğlu and Sönmez (2003) proposed two alternative mechanisms, which are adaptations of widely-studied mechanisms in the literature on matching and assignment markets, following seminal contributions by Gale and Shapley (1962) and Shapley and Scarf (1974). Since that article was published, I have been involved in a number of efforts to re-design school choice systems including those in New York City (2003), Boston (2005), New Orleans (2012), Denver (2012), Washington DC (2013), and Newark (2014).¹ New systems have also been developed in England, Amsterdam, a number of Asian cities, and elsewhere.

The purpose of this article is to review some facts from the field about these re-design efforts and to take stock on what I think has been important in practice so far. This article is not a survey of research on school choice market design (for surveys see, e.g., Pathak (2011) and Abdulkadiroğlu and Sönmez (2013).) My inspiration comes from Klemperer (2002), who presents his views on what matters for practical auction design based on his experience designing auctions and advising bidders. Klemperer concludes that “in short, good auction design is mostly good elementary economics,” whereas “most of the extensive auction literature is of second-order importance for practical auction design.”

My argument in this paper proceeds along similar lines. I argue that what really matters for school choice market design are basic insights about straightforward incentives, transparency, avoiding inefficiency through coordination of offers and well-functioning aftermarkets, and influencing inputs to the design, including applicant decision-making and the quality of schools. Some of the issues examined in the extensive theoretical literature on school choice matching market design are less important for *practical* design. However, my conclusion is not as pessimistic as Klemperer (2002), and I will discuss a handful of issues examined in the theoretical literature on matching mechanisms that have proven to be first-order. It’s worth emphasizing that it is only with the benefit of several design case studies that we’re beginning to understand which issues are quantitatively important.

This paper is organized as follows. Section 2 introduces and reviews the literature on algorithms to match demand and supply, with a focus on some theoretical issues that have emerged from the field.

¹Abdulkadiroğlu, Pathak, Roth, and Sönmez (2005) and Abdulkadiroğlu, Pathak, and Roth (2005) report details on Boston and New York City, respectively.

Section 3 reviews the role of strategy-proofness, transparency, offer coordination, aftermarkets, and participation across a variety of field settings. Section 4 discusses work on inputs to a school choice market design. The last section concludes.

2 Clearing the Market

The initial literature on school choice mechanisms compared the properties of assignment algorithms given the demand and supply in the market, as expressed via student preferences, school priorities, and school capacities. This literature theoretically shows that classic economic tradeoffs between fairness, efficiency, and incentives are unavoidable in the allocation of school seats. With the benefit of datasets from districts using strategy-proof mechanisms, it is now possible to revisit the quantitative magnitudes of these tradeoffs. I begin by reviewing the theoretical background.

2.1 Efficiency vs. Fairness

Abdulkadiroğlu and Sönmez (2003) proposed two strategy-proof student assignment mechanisms. The first is based on the celebrated deferred acceptance (DA) algorithm of Gale and Shapley (1962) and is defined as follows:

Step 1) Each student proposes to her first choice. Each school tentatively assigns seats to its proposers one at a time, following their priority order. Any remaining proposers are rejected.

In general, at

Step k) Each student who was rejected in the previous step proposes to her next best 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 the school's 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.

The second mechanism builds on the Gale's top trading cycles (TTC) algorithm described in Shapley and Scarf (1974). Several authors have extended this algorithm and substantially broadened its use for potential applications including Abdulkadiroğlu and Sönmez (1999) and Papai (2000). I will refer to the version that Abdulkadiroğlu and Sönmez (2003) define for school choice as TTC.² First,

²Recent literature examines alternate versions of TTC for school assignment (e.g., Abdulkadiroğlu, Che, Pathak, Roth, and Tercieux (2015), Dur (2014), and Morrill (2014)).

assign a counter for each school that keeps track of how many seats are still available at the school. Initially set the counters equal to the capacities of the schools. The mechanism works as follows:

Step 1) Each student points to her 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, at

Step k) Each remaining student points to her 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.

Balinski and Sönmez (1999) and Abdulkadiroğlu and Sönmez (2003) observed the close relationship between fairness concepts in the student placement problem and stability in an associated college admissions problem. The following well-known example illustrates how these two mechanisms resolve the trade-off between efficiency and fairness:

Example: Consider an economy 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 are given by the following orderings:

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

The matching produced by DA 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 since there are no blocking pairs, it is stable. The matching produced by the TTC 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 for students and both student i_1 and i_2 obtain their top choice. However, student i_3 prefers school s_1 and has higher priority than student i_2 , who is assigned there. Therefore, student i_3 has *justified envy* at s_1 , or in the language of two-sided matching models, student i_3 and school s_1 form a blocking pair.

Abdulkadiroğlu and Sönmez (2003) framed the debate between these two alternatives in terms of the interpretation of priorities: when elimination of justified envy is a more important goal than efficiency, the DA Pareto should be used because it dominates any other fair outcomes; when efficiency is paramount, they argued for TTC.

In Boston, this article led to one of the first school choice market design efforts in the field after it was featured prominently in a local newspaper (Cook 2003). Boston has a long-standing but controversial school choice system following court-ordered busing in the 1970s. From 2003 to 2005, a Student Assignment Task force, led by civic leaders, evaluated possible improvements to the choice system. While they mostly debated the merits of alternative school zone configurations, which determine the schools a student can rank, they ended up recommending no change to Boston’s three-zone system. However, they did recommend changing the student assignment algorithm. Their recommendation favored TTC over DA. The description of the tradeoffs involved between the algorithms illustrates how these two alternatives were perceived (Landsmark, Dajer, and Gonsalves 2014):

[T]he Gale-Shapley algorithm [...] cuts down on the amount of choice afforded to families. The Top Trading Cycles algorithm also takes into account priorities while leaving some room for choice. [...] [C]hoice was very important to many families who attended community forums...

Based on this report and further consultation with the community and academic experts, the official recommendation was for DA. TTC was faulted for allowing students to trade as stated in the final school committee report (BPS 2005):

[TTC’s] trading shifts the emphasis onto the priority and away from the goals BPS is trying to achieve by granting these priorities in the first place.

Since the debate between DA and TTC sparked by Abdulkadiroğlu and Sönmez (2003) came to such prominence in Boston’s deliberations, it is natural to ask about the extent of differences between

the two systems using data from Boston. A major advantage of both mechanisms is it is a weakly-dominant strategy for applicants to report their true preferences. The strategy-proofness property motivates using preferences from applicants under DA to compute counterfactual assignments, holding the submitted preferences fixed.

Table 1 shows relatively little difference between DA and TTC for applicants in Boston’s choice plan. I use data from four school years when BPS employed DA and examine outcomes for roughly 6,000 elementary, middle and high school applicants per year. The fraction who are assigned their top choice is 65.4% under TTC, compared to 64.8% under DA. The fractions are similar comparing the number who obtain one of their top k choices for $k = \{2, \dots, 10\}$. The similarity of the aggregate rank distribution suggests that in practice, the choice between the two mechanisms does not involve large differences in overall efficiency. The table also tabulates the fraction of students who have justified envy under TTC. A small percentage of total applicants, 6.8%, prefer a school seat over what they are assigned and have higher priority for it.

Table 1 reports on a similar comparison of DA and TTC using data from the Recovery School District (RSD) in New Orleans. The RSD was formed in 2003 and came to oversee the majority of schools in New Orleans following Hurricane Katrina. In 2012, Louisiana state and district officials launched a pioneering universal enrollment process, known as OneApp, which coordinated admissions across traditional and charter school sectors.³ In the first year of OneApp, the assignment process was based on TTC, the first time to my knowledge that TTC was used in a real-life application (for more details, see Abdulkadiroğlu, Che, Pathak, Roth, and Tercieux (2015)). The RSD was initially a perfect test-bed for TTC since the priority structure involves a mix of sibling and neighborhood priorities, much like Boston, and did not involve any schools that screened applicants via test scores or some other criteria.

After one year with TTC, the Recovery School District switched the assignment mechanism to one based on DA. Officials adopted a new algorithm largely based on three reasons: 1) the desire to attract participation from schools in the neighboring Orleans Parish School Board district, many of which involve screened admissions criteria; 2) the need to include private scholarship schools in the system, which by law cannot allow for situations of justified envy;⁴ and 3) the simplicity of explaining results from DA relative to TTC. To investigate whether the switch from TTC to DA in RSD resulted in a

³A charter school is a publicly funded school that operates with more autonomy than a traditional district school. Since the 2014-2015 school year, the RSD consists entirely of charter schools. Abdulkadiroğlu, Angrist, Hull, and Pathak (2014) conduct an evaluation of RSD takeover charter schools.

⁴Louisiana’s Department of Education’s Act 2 expanded Louisiana’s Student Scholarships for Excellence Program statewide, a program that provided state-funded scholarships for low-income students who attend a C, D, or F school to enroll in state-approved non-public schools. The state’s legal counsel advised that blocking pairs potentially produced by TTC might be in violation of Act 2. Abdulkadiroğlu, Pathak, and Walters (2015) evaluate the effect of the program on student achievement.

large change in the aggregate distribution of ranks, Table 1 reports the ranks from TTC and DA for grades Pre-K through 9. As with Boston, slightly more students are assigned their top choice under TTC than DA, but slightly fewer are assigned their second choice. However, the overall distribution of top choices is similar across both mechanisms. 10.0% of students have justified envy under the TTC allocation, which is larger than in Boston.

Given that TTC and DA may generate substantially different allocations in theory, one question motivated by the results in Table 1 is what features of preferences and school priorities in the Boston and New Orleans school districts result in such similar aggregate outcomes. Ergin (2002) identified an acyclicity condition under which the outcome of DA with strict school priorities is Pareto efficient, and Kesten (2006) identified conditions for when two mechanisms are equivalent. Ehlers and Erdil (2010) generalized this result when school priorities have indifferences. Given that these conditions are not likely to be satisfied in practice (and are not in Boston and New Orleans), some have interpreted these results as negative ones (e.g., Abdulkadiroğlu and Sönmez (2013)). It seems more likely that the relatively small tension between efficiency and the elimination of justified envy is driven by aspects of preferences and priorities in those cities that are not captured by the Ergin (2002) and Ehlers and Erdil (2010) conditions. One conjecture involves the role of neighborhood priority, which exists in both cities. All else equal, families prefer schools closer to home, so many students rank schools for which they obtain a high priority. The correlation between preferences and priorities induced by proximity may in turn result in less scope for Pareto-improving trades across priority groups that involve situations of justified envy. This pattern may then result in a small degree of inefficiency in DA, though such an intuition remains to be formalized.⁵

2.2 Student Optimality in DA

A second widely studied issue with DA for school assignment involves turning coarse school priorities into strict ones. Unlike matching applications for labor markets, coarse priorities, such as neighborhood or sibling priority, are widespread. In such an environment, the mechanism must adjudicate claims for school seats between same-priority applicants. In footnote 14, Abdulkadiroğlu and Sönmez (2003)

⁵Table 2 of Abdulkadiroğlu, Pathak, and Roth (2009) compares a student-optimal stable matching to an efficient matching, computed by applying TTC to the economy where the student-optimal stable matching is treated as the students' endowments. Just over 5% of applicants obtain a more preferred assignment under TTC, whereas 44% of applicants are involved in a blocking pair, suggesting a larger efficiency and equity trade-off in New York than in Boston and the RSD. Since high schoolers appear to be more willing to travel for schools than students in earlier grades, this pattern is consistent with those in Boston and New Orleans, which have many more applicants in earlier grades. It's also worth emphasizing that New York City's assignment system involves a mix of screened and unscreened schools, which places it outside of the original Abdulkadiroğlu and Sönmez (2003) school choice model because screened schools submit rankings over pupils as in a two-sided matching model for the labor market context. Therefore, TTC-based mechanisms were never initially considered in New York's design.

suggest that using a single tie-breaking, in which each applicant receives a random number, may be a better idea under DA, and there are now several papers exploring this issue in much greater depth.

To illustrate the issue, consider the earlier example, but now suppose that schools s_1 and s_2 give equal priority to all applicants. That is, the priorities are:

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

where the notation indicates that schools s_1 and s_2 are indifferent between pupils, but pupils are ordered in a strict way at s_3 . If DA 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 each would be better off trading their placements with one another. If the priorities at s_1 had been strict, then one might justify prohibiting this trade because student i_3 forms a blocking pair with school s_1 after the trade. However, this is not a blocking pair in the traditional sense because s_1 is indifferent between applicants. This example shows that DA does not always produce a student-optimal stable matching, and therefore allows for efficiency loss even among stable allocations.

The issue of how best to break ties is a theoretical issue that arose from practical experiences with school districts. Pathak (2011) describes an intuition from policymakers in New York City, where the issue was a central part of designing the rules of the new mechanism. Questions related to tie-breaking have re-appeared in nearly every city I have interacted with using DA. For instance, in Washington DC, where DA was used 2013 and 2014 as part of their universal enrollment system, MySchoolsDC, there has been widespread skepticism over the use of a single lottery number for each applicant. For instance, a parent on a popular online forum summarized the concern (urbandcmoms.com 2014):

Man would it SUCK to get a crappy number for all 12 choices. Wouldn't that essentially knock you out of the whole lottery because, at least for the most popular schools, there is guaranteed to be more people with better numbers who ranked it high enough to be competing with you?

Abdulkadiroğlu, Pathak, and Roth (2009) show that any student-optimal stable matching can be produced by a single lottery draw, so that school-specific lotteries only generate matchings that are not student-optimal relative to a single lottery draw. Their results represent an ex post perspective, and as far as I know, there is no stronger ex ante argument for single versus multiple tie-breaking based on the distribution of matchings. Erdil and Ergin (2008) show that there is no strategy-proof mechanism that produces a student-optimal matching when there is coarse priority ordering at schools. They also

construct a mechanism that computes a student-optimal stable matching in polynomial time. Kesten (2010) defines an enhanced version of deferred acceptance, where students consent to waive certain priorities. Kesten and Ünver (2015) examine how definitions of fairness with deferred acceptance change from an ex ante view and propose two new mechanisms that produce fair allocations.⁶

How much do these issues matter in practice? Table 2, from Abdulkadiroğlu, Pathak, and Roth (2009), reports on the aggregate distribution of choices assigned under DA with single tie breaking (DA-STB) compared to DA with school-specific tie-breaking (DA-MTB). For both Boston’s elementary schools and New York City’s high schools, 3% more applicants obtain their top choice from single tie-breaking than multiple tie-breaking, but slightly more students obtain their 2nd or lower choice under DA-MTB. This phenomenon is also seen in a similar exercise in data from Amsterdam (de Haan, Gautier, Oosterbeek, and van der Klaauw 2015).

In column 3, I report the difference between DA-STB and a student-optimal matching computed via the stable improvement cycles algorithm for Boston. There is virtually no difference between the two assignments in the case of Boston; that is, deferred acceptance with single tie-breaking produces an allocation that is nearly student-optimal. Column 6 shows that for New York City, there is a slight improvement compared to DA-STB, where 1,488 students (or 1.9% of applicants) obtain a more preferred assignment. It’s worth emphasizing that any potential gain from a student-optimal stable matching requires using a mechanism that is no longer strategy-proof.

Quantitatively, multiple tie-breaking appears to result in a larger difference in allocations relative to computing a student-optimal stable matching using the Erdil and Ergin (2008) procedure, but the magnitudes of both exercises seem modest. Though understanding the effects of tweaks to DA is a worthwhile pursuit, there may be little reason to employ them in the field ahead of deferred acceptance with single tie-breaking.

2.3 Constrained Rank Order Lists

The initial theory of school choice mechanisms was based on the assumption that applicants can rank as many choices as they wish. In practice, applicants are usually not allowed to rank all choices, and mechanisms that are strategy-proof without constraints are no longer strategy-proof with this constraint. Pathak and Sönmez (2013) report on a number of systems based on DA, and in all cases aside from Boston, the mechanisms in the field do not allow participants to submit a complete rank order list.

The widespread presence of constraints inspired a theoretical and experimental literature examining

⁶A related literature examines the question of tie-breaking in the context of TTC-based mechanisms. See Abdulkadiroğlu and Sonmez (1998), Pathak and Sethuraman (2011), Carroll (2014), and Lee and Sethuraman (2014).

what happens to popular mechanisms when applicants cannot submit a complete rank order list. Haeringer and Klijn (2009) analyze the properties of the DA and TTC, showing that in DA, when a student may rank only k schools, (1) if a student prefers at most k schools, then she can do no better than submitting her true rank order list; (2) if a student prefers more than k schools, then she can do no better than employing a strategy that selects k schools among the set she prefers to being unassigned and ranking them according to her true preference ordering (Proposition 4.2). Pathak and Sönmez (2013) propose an ordering on manipulability and show that for $\ell > k$, DA when applicants can rank k schools is more manipulable than DA when applicants can rank ℓ schools. This result implies that Chicago’s reform to allow applicants to rank six choices rather than four has made the mechanism less manipulable. Calsamiglia, Haeringer, and Klijn (2010) report results from laboratory experiments showing that constraints do lead to preference manipulations, the most common of which involves participants including a “safety school” among their choices.

Since an applicant can only rank 12 schools in New York’s high school match, if the applicant ranks all 12 choices, then the applicant may potentially prefer more choices, and hence his incentives to truthfully rank schools is affected by this constraint on the number of choices. On the other hand, if an applicant ranks less than 12 choices, he could have ranked more choices but opted not to, so his preference submission problem was not constrained by the mechanism. This observation motivates tabulating the fraction of applicants who submit a rank order list equal to the maximum allowed in systems with constraints. Table 3 reports on New York City, Chicago, Denver and New Orleans, for years where each city used a mechanism that would be strategy-proof in the absence of a constraint.

The table shows that when there are few options, many more applicants rank the same number of schools as the constraint. In districts with many options, a relatively smaller fraction are at the constraint. For instance, in 2009, when students at Chicago’s exam high schools were allowed to rank four out of nine choices, more than 80% of applicants ranked all four schools. In 2010, the district allowed students to rank six choices, and roughly half of applicants listed all six choices. In contrast, in New York City, roughly one fifth of participants rank all 12 choices, but 12 choices represents a very small percentage of the choice set. Denver’s constraint of five choices is about 6% of all possible rankable schools, and about 25% of students rank five schools. Finally, the constraint in New Orleans seems least severe, with less than 5% of students ranking all eight possible choices.

Given how a constraint interferes with the incentive properties of mechanisms, it continues to be a puzzle why it is much more common for system operators to restrict the number of choices a participant can rank rather than having no constraint on what can be ranked. It seems unlikely that technical constraints prevent a system that allows participants to rank many choices. For instance, in the National Residency Matching Program, there is no limit on the number of choices an applicant can

rank in the Main residency match, though applicants may have to pay a fee to rank more programs than on the standard form. However, that market involves an important interview stage where a fair amount of sorting takes place, and a similar phase is not present in most school choice settings, in which applicants most often apply to a program without the need to interview or visit.

Market operators have made three different arguments to me about constraining rank order lists. First, they are concerned about the perceived complications involved in ranking many choices, but it's not clear that the costs of information acquisition motivate capping the number of choices for everyone. Second, they are reluctant to advertise that an applicant received their 20th choice because it makes salient that there is a shortage of desirable options. However, this hardly seems better than advertising that an applicant is unassigned. Finally, operators have emphasized that take-up rates would be lower when applicants can rank more choices. They are concerned that applicants would submit less "serious" choices only to not take them up after the main assignments are announced, leaving those seats vacant. One simple response is to over-book schools accounting for less than full yield. Abdulkadiroglu, Agarwal, and Pathak (2015) report that take-up of main round assignment in New York City is largely invariant to the length of an applicant's rank order list. For instance, the fraction of students who rank 5 choices (who are assigned on average their 2.5th choice) who enroll in the school assigned in their main round is 91.2%, while the fraction of students who rank 12 choices (who are assigned on average their 3.9th choice) is 94.3%. The take-up rates tend to be lower for students who have ranked few choices, though these students are much more likely to enroll in private schools. It's possible that these patterns would change if participants were able to rank more choices, but there is no direct evidence of that.

Overall, despite the fact that so many systems place ad hoc constraints on the number of schools an applicant can rank, it does not seem to be a pressing concern for operators even in cities where the constraint appears to bind more severely. This suggests that either there are aspects of the constraint that existing models don't capture or this design issue is less central in practice than initially thought.

2.4 Ranking Behavior in New Mechanisms

Abdulkadiroğlu and Sönmez (2003) emphasize that many popular assignment schemes do not encourage participants to truthfully report their preferences. During initial meetings in Boston, officials referenced results on preference manipulation in Boston, DA, and TTC in laboratory experiments in Chen and Sönmez (2006). The paper reports that there is a higher degree of preference manipulation under Boston than either alternative. For instance, in one of the treatments, less than 30% of applicants report their true preferences under the Boston mechanism, while over half of applicants report their true preferences under DA. They report that a common strategy involves ranking a district

school, for which they obtain priority, in a higher position than the true preference order.

Using field data from Boston from 1999-2005, Abdulkadiroğlu, Pathak, Roth, and Sönmez (2006) report that about 20% of applicants ranked two popular schools as top and second choices. They argued that this ranking behavior was a mistake because an applicant stands no chance of getting assigned their second choice if it is heavily oversubscribed because its seats will have been depleted by first choice applicants in the Boston mechanism. Not all applicants submit preferences naively, however. Without quantifying the fraction of strategic players, Abdulkadiroğlu, Pathak, Roth, and Sönmez (2006) collect anecdotes on strategic heuristics from parent groups who appear to understand incentive properties of the mechanism. Following this work, it's remained an open question about how many applicants understand and react to the incentives to manipulate preferences in non-strategy proof mechanisms. I next turn to field data from policy changes in Boston and Chicago to examine how families might be responding to mechanisms.

2.4.1 Before and After DA in Boston

It's clear that the extent of any behavioral response will depend crucially on outreach efforts and whether participants believe and react to the advice given by those who run the mechanism. School officials in Boston who studied the Chen and Sönmez (2006) experiment observed that in the laboratory experiment, descriptions of the alternative mechanisms did not indicate that truthful reporting is an optimal strategy. This experimental design decision was intended to isolate the actual incentive properties of the mechanism as experienced by laboratory subjects, not how participants responded to advice about mechanisms. In the experimental instructions, the authors went to great lengths to limit the variation in how the three mechanisms were described.

The absence of advice was notable to policymakers, given that BPS had been giving recommendations to participants in their school brochures. Prior to 2005, the Boston school brochure recommended (emphasis in original):

For a better chance of obtaining your “first” choice, consider choosing less popular schools.

Under a strategy-proof mechanism, it is no longer necessary to rank schools based on their popularity. Following the change in the mechanism, from 2006 through 2010, DA was described in school brochures as follows (emphasis in original):

Assignments are made by a computer that is programmed with a **mathematical formula**. The computer programs tries to assign students to their highest listed **choice** for which they have the highest **priority**.

The brochure also stated (BPS 2008):

List your school choices in your true order of preference. There is no need to “strategize.” If you list a popular school first, you won’t hurt your chances of getting your second choice school if you don’t get your first choice.

Aside from changes in how the mechanism was described in written communications, there were also outreach activities from BPS that trained family resource center staff members, who field questions on how school choice works in Boston. Of course, counselors in these centers are not necessarily algorithm experts, so it seems likely some of the intended effects of the Boston change would be gradual as information disseminates and families gained experience with the new system over time.

To focus our discussion on what happened after Boston’s algorithm change, I recount the deliberations prior to the vote to change the mechanism in official hearings. First, Boston officials conjectured that (BPS 2005):

selection habits will change, and fewer students are expected to receive their first choice school if more people are vying for seats in over-demanded schools.

Second, they speculated that:

the number of school choices made by families is likely to increase, as parents will be more inclined to list all schools they are interested in, not just those to which they are more likely to gain admittance.

Using data from years before and after the mechanism change, Figure 1 shows that applicants ranked more schools following the adoption of DA. The figure reports the fraction of applicants ranking four or more choices from 2000 to 2011 for grades 6 and 9.⁷ Applicants are therefore being more expressive by listing more schools, suggesting that under the Boston mechanism, applicants had less of an incentive to consider choices beyond their top one.

Figure 2 shows that there is a reduction in the number of students who obtain their stated top choice. For both grades 6 and 9, when the Boston mechanism was used, roughly 70% of students obtained their top choice, while in the following years, the number decreases to between 50-60%. Since the Boston mechanism prioritizes first choices, this outcome likely represents a best-case scenario for assigning as many as possible to their top choice, holding fixed applicant behavior.

⁷I do not report numbers for the elementary school entry point because alongside the new algorithm, there was an expansion of the number of kindergarten programs in the city. In particular, in early years grade K2 was the main entry point, but this shifted to K1, which caused a number of applicants at K2 to be effectively guaranteed at their choice.

While a more thorough investigation of the effects of Boston’s new mechanism is beyond the scope of this paper, the patterns in both figures illustrate that two of the conjectured effects of the new mechanism occurred in Boston. Of course, attributing all of these changes solely to the new mechanism does not take into account the possibility that other aspects of the market may also have changed. Moreover, there is a different set of applicants participating across years, so some of the variation in the submitted reports may reflect the preferences of these new applicants. I therefore turn to an examination of an unusual experiment in Chicago’s Public Schools.

2.4.2 Chicago

Pathak and Sönmez (2013) report on a stunning midstream change in Chicago’s assignment system for its elite exam schools in 2009. After eliciting preferences from over 13,000 8th graders, CPS officials announced a new system and allowed any 8th grader to re-rank preferences if they wished. Prior to this reform, Chicago’s school system placed children to exam schools using a version of the Boston mechanism, where schools were allowed to make uncoordinated offers and a racial quota was in place. Following the reform, the system was based on a version of DA, and the racial quota changed to an affirmative action system based on a socioeconomic index⁸ We therefore can observe the same participants submitting rankings under two different assignment mechanisms. Submissions occurred about two months apart, allowing for the possibility that preference changes reflect new information, so the magnitude of potential changes can be attributed to both the new mechanism and time elapsing between preference submissions. Nonetheless, unlike the comparisons I reviewed in Boston, the set of participants are fixed in Chicago.

Before reporting on applicant responses, it’s worth noting the type of advice given to participants about the new system. After their initial preference submission, each applicant received a letter stating:

...the original application deadline is being extended to allow applicants an opportunity to review and re-rank their Selection Enrollment High School choices, if they wish. It is recommended that applicants rank their school choices honestly, listing schools in the order of their preference, while also identifying schools where they have a reasonable chance of acceptance.

Table 4 shows that the vast majority of applicants do not change their rank order list at all after receiving this letter. 89% of applicants continue to apply to the same set of schools in the same order, whereas 93% of applicants continue to list the same four schools, though not necessarily in the same order. It’s possible that large switching costs are responsible for the fact that most applicants did

⁸For more details on the socioeconomic index, see Ellison and Pathak (2015) and Dur, Pathak, and Sönmez (2015).

not change their preferences, though I think that is unlikely because each participant received a new application form. Instead, I think that the 11% of applicants who changed their preferences are either reacting to new information obtained since making their initial preference submission or responding to the new mechanism's incentives.

Chicago's midstream change may be the best case so far for measuring how participants respond to the incentives of a mechanism in the field. However, our ability to draw inferences from this policy change comes with an important weakness: under both the old and new mechanism, an applicant can apply to at most four different schools. An applicant's incentives to change their rank ordering may be dominated by the fact that he can only rank four choices. In other words, it is possible that a preference manipulation under the old mechanism still remains a manipulation under the new mechanism. Pathak and Sönmez (2013) formally demonstrate that Chicago's new mechanism is less manipulable than the old one, and that the old mechanism was highly manipulable. That is, the old mechanism is the most manipulable among a large class of mechanisms. Therefore, it seems unlikely that an applicant's report under the old mechanism is still their the best strategy under the new mechanism.

In light of earlier theoretical work that distinguishes between sincere and sophisticated players, it is worthwhile to examine the attributes of the 11% of applicants who reported a different rank order list. In Table 5, I report on correlates of preference changers and applicant covariates. For instance, applicants who change at least one choice tend to have higher admissions test points and baseline scores. They tend to be more black and Hispanic, and they are less likely to be absent from school. There is almost no relationship between whether students are from high- or low-SES census tracts, where I define census tract characteristics using the same 7-factor model that is used by CPS in their assignment system.⁹

It therefore appears that the policy change in Chicago induced a modest response in how applicants ranked schools. The magnitude of the behavioral response is swamped, for instance, by the mechanical change in how applicants are processed by the mechanism now that an applicant with a higher test score can no longer be rejected from a school due to an application from someone with a lower score who ranked that school higher.

The findings from Chicago contribute to an ongoing debate in the literature on the extent of preference manipulations in data. Hastings, Kane, and Staiger (2009) show that applicant's preferences for school quality did not react to a change in neighborhood boundaries in Charlotte North Carolina, and they use this fact to motivate treating reported preferences as truthful even though Charlotte employs a variant of the Boston mechanism. This study, like the Chicago evidence, is based

⁹For more details on this affirmative action scheme, see Ellison and Pathak (2015) and Dur, Pathak, and Sönmez (2015).

on observing participant reactions to changes in the assignment mechanism. Other work develops methods to quantify preference manipulations using variation within a particular assignment mechanism. Agarwal and Somaini (2014) report discontinuities in preferences at the proximity priority boundaries that they attribute to preference manipulation. They then develop a method to recover preferences from potentially manipulated reports and find that more than half of participants find it optimal to report their true preferences. Calsamiglia, Fu, and Guell (2014) estimate a mixture model with sincere and sophisticated applicants and find that the overwhelming majority of participants are responding strategically in Barcelona’s adaption of the Boston mechanism. He (2012) considers the empirical implications of assuming that applicants do not submit strictly dominated rankings using data from Beijing, which employs a version of the Boston mechanism, though without any priorities. Hwang (2015) is another similar effort.

If we interpret the Chicago evidence as suggesting that only a small fraction of applicants respond to the incentives in manipulable mechanisms by manipulating their preferences, then it begs the question of why it was cited as a factor in Chicago’s decision to abandon the mechanism midstream. I next turn to describe some other experiences in the field, which show how participants perceive and reject highly manipulable mechanisms. These episodes suggest some new rationales for mechanisms with straightforward incentives.

3 Considerations from the Field

3.1 Straightforward Incentives

Abdulkadiroğlu and Sönmez (2003) describe existing manipulable assignment schemes as ones where “students and their parents are forced to play very complicated admissions games...[which] is confusing to students and their parents, but also results in inefficient allocation of school seats.” The aspect of Boston’s scheme that most strongly resonated with BPS officials in policy discussion was not the possibility of inefficiency, but rather that their manipulable mechanism might stand in the way of allowing for equal access to high-quality schools across the district. Differential knowledge about manipulation opportunities could exacerbate inequalities as sophisticated players benefit from understanding algorithmic details at the expense of those who were not aware. In public hearings, Boston school officials articulated this rationale as follows (BPS 2005):

a strategy-proof algorithm ‘levels the playing field’ by diminishing the harm done to parents who do not strategize or do not strategize well. [...] [T]he need to strategize provides an advantage to families who have the time, resources, and knowledge to conduct the necessary research.

Pathak and Sönmez (2008) present a formal investigation of this “equal access” argument for strategy-proof mechanisms in a model with sincere and sophisticated players in the preference revelation game induced by the Boston mechanism. Sincere players do not react to the incentives for preference manipulation, while sophisticated players best respond in the model. The paper shows that sophisticated players prefer the manipulable mechanism because they can exploit their knowledge about how the mechanism works. In particular, compared to the equilibrium outcome of the student-proposing deferred acceptance algorithm, sophisticated students weakly prefer their outcome under the Pareto-dominant Nash equilibrium of the Boston game. Even though sincere players lose priority to sophisticated players in equilibrium, it is sometimes possible that they prefer their outcome under the Boston mechanism. It’s interesting to note that when the consequences of manipulation are modeled in this way, they imply distributional effects across player sophistication, which are not present if everyone is assumed to play strategically as in Ergin and Sönmez (2006). Dur, Hammond, and Morrill (2015) identify sophistication by tracking applicants who log in to the online application portal and revise their preferences in response to feedback on existing demand. They then show how strategic sophistication allows applicants to get admissions into higher quality schools than sincere participants.

The idea that a manipulable mechanism frustrates participants and creates inequities for sincere participants is a theme that I have seen in cities other than Boston. Even if only a small fraction of participants manipulate their preferences, these players appear to bear substantial costs, with potentially uncertain benefits. Even in the context of the model of Pathak and Sönmez (2008), the benefit that a sophisticated player obtains from manipulation still requires coordination at the Pareto-dominant equilibrium. Moreover, the perception that the mechanism encourages “gaming” suggests that the mechanism and its operator cannot be trusted. I believe these two aspects – not the potential for allocative inefficiency resulting from manipulation - are central for rationalizing stunning dismissals of highly manipulable mechanisms.

I’ve already described how Chicago abandoned their mechanism after eliciting rankings from thousands of participants. Another indication of how society perceives highly manipulable mechanisms comes from England, where a manipulable mechanism was outlawed by an Act of Parliament (Pathak and Sönmez 2013). The 2007 English School Admissions Code, which regulates “National Offer Day” for hundreds of thousands of 10 and 11 year olds looking for school seats throughout England, outlawed a particular class of mechanisms (in Section 2.13):

In setting oversubscription criteria the admission authorities for all maintained schools must not:
give priority to children according to the order of other schools named as preferences by their parents, including 'first preference first' arrangements.

More specifically, a first preference first system is any “oversubscription criterion that gives priority to children according to the order of other schools named as a preference by their parents, or only considers applications stated as a first preference” (DfES 2007). The Boston mechanism is probably the most famous example of a first preference first system. Both the 2007 and 2010 Admissions Code outlaws use of first preference first systems at more than 150 Local Authorities England.

What is remarkable about this revision of the admissions code is that, to my knowledge, it occurred without any direct involvement of matching theorists or economists. Unlike in Boston, where economists played an active role lobbying against the mechanism, there is no evidence that academic research on school assignment either directly inspired rejections of the old mechanisms or led to the development of the new ones. Nonetheless, the arguments made by English authorities bear an uncanny resemblance to points discussed in the Boston debate. For instance, England’s Department for Education and Skills viewed first preference first arrangements as making the system “unnecessarily complex” (DfES 2007). Newspaper accounts alleged that first preference first arrangements “force many parents to play an ‘admissions game’ with their children’s future” (Smith 2007).

Even prior to this nationwide ban, numerous regions across England had already done away with first preference first arrangements. The Pan-London Admissions scheme adopted variants of DA known as “equal preference systems,” where all preferences were considered without reference to the rank order made by parents (Pennell, West, and Hind 2006). The report of the Pan-London Board and London Inter-Authority Admissions Group states that equal preference scheme was designed to “make the admissions system fairer” and “create a simpler system for parents” (Association of London Government 2005). The Newcastle Admissions Forum recommended that “the equal preference system was more parent-friendly as it would reduce anxiety among parents as they can set out their ranked preferences without having to calculate the chances of their getting a place” (Young 2003). As with Chicago, the poor incentive properties of the assignment mechanism appears to be central for its dismissal.

Pathak, Song, and Sonmez (2015) recount similar turmoil driven by a highly manipulable mechanism used for school assignment in Taiwan, which appear unrelated to concerns about allocative efficiency. In June 2014, more than five hundred parents and teachers marched in downtown Taipei, Taiwan, protesting the deduction point system used for senior high school placement largely because it encourages strategic ranking. Fiercely protesting parents held placards stating, “fill out the preference form for us,” while others complained that the process was akin to “gambling” (I-chia 2014). These protests were inspired by the 2014 Senior-High School Education Act, which changed both exam scoring and the assignment mechanism.

In Taiwan’s new assignment system, students obtain points or priority based on the order in which

preferences are ranked. For instance, in Jibei, the largest district with over 60,000 applicants, a student obtains thirty points on top of their maximum possible score of 90 points for their first choice. For their second choice, only 29 points are added to their test score. Compared to the first choice, the next ten choices experience a deduction of one point from the test score relative to one choice immediately higher. In other districts, the points are the same for a group of choices. In the limit where deduction points are large, the Taiwanese system is equal to a special version of the Boston mechanism, where the priorities are the same at each school. Therefore, protests in Taiwan represent yet another situation where a variant of the Boston mechanism generated public outcry and reproach.

The widespread condemnation of variations of the Boston mechanism from Chicago, England, and Taiwan involved participants themselves protesting and re-organizing market designs (and not matching theorists). In this respect, the school admissions reforms parallel changes in marketplace rules for the placement of medical residents in the early 1950s documented by Roth (1984). However, while reforms of the residency market were motivated in part by concerns about the inefficiency of unraveling, the school choice design efforts were motivated by concerns about excessive vulnerability to “gaming.” In that respect, these efforts challenge a traditional mechanism design paradigm that treats incentive compatibility only as a constraint and not as a direct design objective. That is, in the discussions about these new policies there is almost no attempt to establish a direct link between preference manipulation and the inefficiency of outcomes, suggesting that the rationales for new mechanisms are in part inspired by non-consequentialist objectives.

Roth (1991) recounts the history of regional markets for physicians in the UK and argues that stable marketplaces were more likely to succeed than centralized markets that produced unstable outcomes. He therefore attributes a positive interpretation to stability as a governing principle of market design in addition to its traditional normative rationale. Likewise, when interpreted through the lens of the public and policymaker’s revealed preferences, it is clear that strongly manipulable mechanisms are perceived as undesirable and inequitable, even absent overwhelming evidence that they generate widespread inefficiencies and that large numbers manipulate their preferences.

3.2 Transparency

Another property that shapes actual school choice market designs in the field involves transparency, or the ability to explain how the assignment process works. This aspect of admissions is particularly important when a new centralized scheme replaces a largely unregulated market, where both students and schools must trust the system operator, as was the case in New Orleans, Denver, and Washington DC.

Transparency is paramount even for modifications of algorithms within existing centralized sys-

tems. For instance, in Boston, the new mechanism was touted as one that (BPS 2005):

Adds transparency and clarity to the assignment process, by allowing for clear and straight forward advice to parents regarding how to rank schools.

Similarly, Abdulkadiroğlu, Pathak, and Roth (2009) describe how NYC school officials preferred a simple mechanism with straightforward incentive properties over one that might allow for strategic manipulation because of the ability to give advice. While hard to formalize, the idea that a mechanism is transparent and simple to understand plays as important a role in choosing mechanisms as the algorithmic properties of mechanisms. It has even played a role in decisions between strategy-proof mechanisms, as I describe next.

3.2.1 Aversion to TTC

Transparency and the ability to explain the process were two main reasons why Boston officials opted for DA over TTC, despite the fact that DA may not be efficient. Quoting BPS officials (BPS 2005):

... trading of priorities [under TTC] could lead families to believe they can still benefit from strategizing, as they may be encouraged to rank schools to which they have priority, even if they would not have put it on the form if the opportunity for trading did not exist. The behind the scenes mechanized trading makes the student assignment process less transparent.

Given that TTC is strategy-proof, it would be possible to advise participants that reporting preferences truthfully is the best course of action. However, BPS officials were concerned that when TTC was described, participants may perceive preference manipulations to be advantageous, even though it would be possible to provide a clear recommendation reporting the truth is best for applicants. The BPS position has some support in the laboratory experiments of Chen and Sönmez (2006), where the rates of preference manipulation under TTC are higher than under DA.

The difficulty in explaining TTC compared to DA is also a reason that Recovery School District in New Orleans switched mechanisms after one year with DA. Those overseeing the OneApp process frequently fielded questions from parents and schools on how TTC worked and why it produced a particular outcome, which they often found challenging to handle. In particular, under DA, when an applicant is not admitted to a school, they could explain that there was another applicant who applied to the school who had higher priority. Under TTC, such an explanation is not as simple, because lack of admittance is due to a higher priority applicant trading away his seat with someone who wanted the school (and may have in fact had lower priority than the rejected applicant). I believe that the difficulty of explaining TTC, together with the precedent set by New York and Boston's choice of DA,

are more likely explanations for why TTC is not used in more districts rather than the fact that it allows for justified envy, while DA does not.

3.2.2 Demise of Walk Zones in Boston

The reforms to the Boston choice system in 2013 are another setting where the desire for increased transparency led to dramatic changes in assignment mechanisms. Until 2013, Boston Public Schools used walk zone priority to assign students for half of each school’s seats. For the other half of the school seats, walk zone priority did not apply. BPS’s 50-50 seat split was official policy starting in 1999, following the end of racial and ethnic tie-breakers. At that point, there was a citywide debate, with a faction pushing for the elimination of the choice system and a return to neighborhood schools. Motivated by concerns that such a system would increase segregation across schools, the school committee chose to reduce the fraction of seats where walk-zone priority applies from 100% to 50% of seats within each school. When an applicant with a high lottery number also has walk zone priority, they can be assigned either to a walk zone slot or a non-walk slot.

Dur, Kominers, Pathak, and Sönmez (2014) theoretically and empirically examine the performance of DA in a matching model with slot-specific priorities motivated by Boston’s problem. The paper shows that because of how the 50-50 split was implemented, the outcome of the BPS mechanism is not at the midpoint of the allocation with no walk zone priority and the allocation when all seats use walk zone priority. Instead, the allocation produced by Boston’s implementation of DA is nearly the same as the allocation when none of the seats use a walk-zone priority. The reason is that BPS’s choice of precedence – the order in which seats are depleted by the mechanism – undermined the priority policy. Boston’s implementation ensured that any applicant with walk-zone priority who did not qualify for a walk-zone slot had little chance of being assigned a school seat at the other half. While there was active debate about the priority structure, the precedence, a much more arcane aspect of the assignment mechanism, was a detail left to the system operator. This detail ended up having significant implications for the final allocation.

Once the findings of Dur, Kominers, Pathak, and Sönmez (2014) were made public, there was widespread confusion about how walk-zone priority could matter little due to precedence. Many thought that walk-zone applicants had been advantaged by the current 50-50 rule, even though in practice they appear to have had little advantage. The discovery about how precedence became a central part of the fight between those favoring neighborhood assignment and those favoring increased choice, with additional details on the various factions, is outlined in Dur, Kominers, Pathak, and Sönmez (2014).

After drawn-out deliberations, policymakers decided that simply eliminating walk zone priority

may be the best solution, especially in light of the observation that it played a relatively small role in practice. Moreover, getting rid of walk zone priority altogether avoids the (false) impression that applicants from the walk zone are receiving a boost under the mechanism. In March 2013, the Superintendent recommended eliminating walk zones entirely, overturning an earlier recommendation (Johnson 2013):

Leaving the walk zone priority to continue as it currently operates is not a good option. We know from research that it does not make a significant difference the way it is applied today: although people may have thought that it did, the walk zone priority does not in fact actually help students attend schools closer to home. The External Advisory Committee suggested taking this important issue up in two years, but I believe we are ready to take this step now. We must ensure the Home-Based system works in an honest and transparent way from the very beginning.

The quote emphasizes that lack of transparency with the previous implementation of walk zone priority was a major reason for its demise. Alongside the elimination of walk zone priority, the city also adopted a new procedure, known as the Home-Based system, to determine where applicants can apply; see Pathak and Shi (2014) for more details. This policy change was inspired by a desire to reduce the district’s busing costs, while still ensuring equitable access across the district.

Boston’s recent experience also illustrates how the criteria used to allocate seats – taken as given by much of the literature – are just as important as how the mechanism processes applicants’ claims. The public process that led to the new system in 2013 followed many similar attempts that resulted in stalemates. One consequence of the increasing adoption of strategy-proof mechanisms is that the role of students’ property rights in determining the distribution of school access becomes more transparent. A small but growing body of research directly addresses how these property rights interact with the mechanism beyond considerations of justified envy, which is emphasized in the earlier literature (see, e.g., Calsamiglia and Miralles (2014), Echenique and Yenmez (2014), Shi (2014) and Dur, Pathak, and Sönmez (2015)).

3.2.3 Seattle’s Setback

While concerns for transparency motivated policymakers to consider simplifications of existing mechanisms, the lack of transparency allows system operators to achieve certain ends through indirect channels. One example of this phenomenon is the policy of “principal’s discretion” used for assignment at Chicago’s exam schools. This policy allows principals of exam schools to handpick applicants for five percent of incoming seats. The application process considers criteria aside from that used in

the centralized mechanism and has led to accusations of special assignments for politically connected applicants. A 2010 Audit reported that nearly all high school principals routinely received telephone calls and inquiries from politicians, CPS staff members, and others seeking preferential treatment for enrollment in schools, and there have been calls to eliminate this practice (Inspector-General 2010). Principals defend the policy as necessary for ensuring a diverse student body or to give students another chance for admissions. It is of course possible to accommodate some concerns for diversity within centralized choice plans (see, e.g., Erdil and Kumano (2012), Hafalir, Yenmez, and Yildirim (2012), Kojima (2012), Kominers and Sönmez (2014) and Echenique and Yenmez (2014)), but doing so in an opaque way such as principal’s discretion appears to undermine participants’ trust in the system.

Seattle Public School’s experience with school choice provides another example of the dangers of the lack of transparency. In 1999, Seattle switched from the Boston mechanism to DA, but called it the Barnhart-Waldman (BW) amendment in honor of two school board members who proposed the modification. Given that this change pre-dated academic work on school choice mechanism design, it is not surprising that it was not widely understood. For instance, in a court challenge to the Seattle choice plan by *Parents Involved in Community Schools*, a case eventually decided by the U.S. Supreme Court, confusion about the mechanism came up in the school board president’s deposition (US District Court of Appeals (2001)):

Q: Can you explain for me what the Barnhart/Waldman Amendment is and how it works?

A: If I could I’d be the first. The Barnhart/Waldman – this is my understanding. The Barnhart/Waldman Amendment affects the way that choices are processed. Before we adopted that amendment, all the first choices were processed in one batch and assignments made. If you did not get your first choice, it is my understanding that all the students who did not get the first choice fell to the bottom of the batch processing line, and then they would process the second choices, et cetera. Barnhart/Waldman says that after all the first choices are processed, in the next batch, if you don’t get your first choice, you don’t fall to the bottom of the list but you are then processed, your second choice, with all the other second choices together. The result is that instead of a high degree of certainty placed - or of value placed on first choice, people can list authentically their first, second and third choices and have a higher degree of getting their second and third choice if they do not get their first choice. Now, was that clear as mud?

In 2007, a parent obtained the computer code and verified that the BW amendment actually corresponds to the student-optimal stable mechanism (McGregor 2007). Interestingly, researchers did not learn about the change until Seattle returned to the Boston mechanism in 2009.¹⁰

¹⁰The first reference to Seattle is in Abdulkadiroğlu, Che, and Yasuda (2011), who describe the episode as the “clock turning back.”

Pathak and Sönmez (2013) present evidence that lack of understanding about the mechanism was exploited by some who wanted to reduce transportation costs in the current mechanism. I corresponded with some of the school committee members involved in this decision. While they mentioned a few hard-to-square reasons (such as computer implementation costs), one suggested that the BW amendment encouraged mobility among students since parents could freely express their choices, without having to take their geographic priority into account. By forcing families to adopt more conservative strategies such as ranking their neighborhood schools, the return to the Boston mechanism discouraged student movement and therefore reduced transportation costs. My impression from these interviews is that the absence of a transparent design was exploited by a faction that was uncomfortable with the idea of school choice in the first place. Returning to the Boston mechanism was a politically attractive alternative to decrease mobility across neighborhood zones while still maintaining the illusion of choice, even though that goal could have been achieved in a more transparent manner by simply changing priorities and choice menus within the deferred acceptance algorithm.

3.3 Coordinating Offers

Some of the largest allocative improvements due to new market designs are from relatively simple ideas. One important example involves eliminating multiple offers in favor of systems that produce a single offer for each applicant. Proponents of multiple-offer systems believe that parents value the added flexibility of investigating their options following assignment, but this comes at the cost of making sure extra offers percolate to those initially not offered. There's often not enough time for the market to clear in multiple offer systems, and managing yields through over-booking of school seats requires lots of experience and fine-tuning. In the labor market context, Roth and Xing (1997) highlighted that congestion can generate inefficiencies, and automation of the offer process can go a long way towards improving the market's performance.

Abdulkadiroglu, Agarwal, and Pathak (2015) measure the effects of coordinating admissions by examining the change in New York City's high school assignment system in 2003. Prior to 2003, roughly 80,000 aspiring high school students applied to five out of more than 600 school programs; they could receive multiple offers and be placed on wait lists. Students in turn were allowed to accept only one school and one wait list offer, and the cycle of offers and acceptances repeated two more times. Students initially expressed their preferences on a common application, but admissions offers were not coordinated across schools, allowing some students to obtain more than one offer. In Fall 2003, the system was replaced by a single-offer assignment system based on DA for the main round. Applicants were allowed to rank up to 12 programs for enrollment in 2004-05, and a supplementary round placed those unassigned in the main round.

Abdulkadiroglu, Agarwal, and Pathak (2015) document how few rounds of offer processing and a limited number of applications make it difficult for schools to make enough offers to assign to clear the market. In the uncoordinated mechanism, more than 25,000 applicants were assigned in the administrative round. Roughly one-fifth of applicants received more than one first round offer, generating about 17,000 extra offers in the first round. However, the second and third round only processed 10,000 applicants, and this led to the large number unassigned students after the main round. Students assigned administratively are placed at schools with attributes that differ from what they initially ranked. They are also more likely to enroll in schools other than where they are assigned and more likely to leave the public schools altogether.

Using preference estimates from econometric models of school demand, Abdulkadiroglu, Agarwal, and Pathak (2015) compare the aggregate utility associated with several alternative allocations. On one extreme is neighborhood assignment, where pupils are allotted to schools closest to their home, given capacity constraints. The other extreme is the utilitarian optimal assignment, computing by taking estimated cardinal utility parameters and finding the allocation that maximizes the sum. This interval provides a way to gauge the magnitude of different aspects of school choice market design. The paper identifies large allocative gains from a choice system, given substantial heterogeneity in student preferences. It also investigates the effects of relaxing mechanism design constraints and finds that they are much smaller than the gains from moving from the uncoordinated to the coordinated assignment scheme. This finding suggests that the gains from fine-tuning algorithmic details are swamped by simply moving from multiple to single offers.

It's important to note that the automation that comes from market design inspired solutions is not necessarily always beneficial. For instance, another growing trend in school assignment involves the adoption of a common application, where an applicant can apply to all schools in one place, and then schools make offers in an uncoordinated way. A large fraction of US colleges participate in a common application, and these systems are often intended to streamline the process of applying to colleges. The fact that offers are not coordinated suggests that it will be necessary to develop a way to ensure that offers percolate through the system to avoid congestion problems like those experienced in New York City. For colleges, there may be more experience with yield management and smaller costs for either over or under-booking (see Che and Koh (2015) for an analysis of yield management when there is uncertainty over student preferences).

Avery, Nickolaus, and Pathak (2015) develop a model of decentralized school assignment with and without a common application and compare it to a centralized scheme based on DA. They show how a common application eliminates an important screening component of a system where applications are costly, and that it does not necessarily result in better matches. They then compare their model

to data on magnet elementary and high schools, which adopted a common application in Fall 2010. They show that while the number of applications increased, a common application does not seem to lead to improved matches.

3.4 Aftermarkets

Another issue brought to light by experiences in the field involves how applicants who are unmatched or who wish to change their assignment are processed. Some cities operate wait lists, while others conduct second round matches. For instance, in Boston Public Schools, DA is actually run four separate times, with adjustments to priorities in each round that give applicants the highest priority at the school (if any) that they were assigned in the previous round. This approach allows applicants to resubmit preferences if they wish, but also does not switch the assignment of applicants who do not wish to move. Most applicants in rounds after the first are not in the main transition grades and therefore have not participated in the first round. But there are a handful who have changed their mind after the first round, and BPS's system tries to mimic the continuation of DA for these applicants. This well-functioning process for handling applicants after the main round contrasts with that in place in Denver's school choice system. Currently, Denver runs an ad hoc school-by-school application process, which is unrelated to what occurred in the match.

Most models of school assignment are static, even though circumstances change and students who were not assigned need to be placed at the conclusion of the assignment process. A major lesson of the reform of New York City's assignment mechanism is that it is undesirable to be placed in the administrative round. Abdulkadiroglu, Agarwal, and Pathak (2015) show welfare gains are larger for populations of students who are more likely to have been assigned administratively. This fact raises the question of what can be done to improve the administrative round or assignments following the initial match.

One approach is to develop a centralized procedure for applicants who are unassigned or who wish to change assignments. Before the adoption of DA in NYC, those unassigned were simply manually placed to schools that had extra capacity after the main round. Aware of the fact that there were going to be unassigned applicants, in the new mechanism, NYC officials developed a supplementary round where unassigned applicants could re-express preferences. But even after this round, students could be unassigned, and they were placed manually. Why not simply elicit preferences from these applicants again? Narita (2015) investigates properties of schemes that take evolving preferences into account. He shows that a two-stage mechanism potentially improves welfare over a static mechanism.¹¹

¹¹There have been similar proposals for a two-stage match in the National Residency Matching Program, the main clearinghouse for U.S. doctors.

Interestingly, he finds that the effects of his dynamic mechanism are much larger when combined with efforts to speed up learning about school choices. Given the mismatch associated with the administrative round, there are likely large returns to efforts to place students who are left over after the match or to reduce the fraction who need to be assigned post-match.

4 Making the Market Work

The two ingredients of any school assignment problem are the students and schools. Much of the recent school choice market design literature focuses on market clearing algorithms because there are not enough desirable schools available to accommodate the demand of all applicants. There is growing interest in understanding the relationship between the assignment process and what is being assigned. I start by reviewing research that focuses on the student side before turning to an examination of the school side.

4.1 Families as Consumers

Starting with Abdulkadiroğlu and Sönmez (2003), models of the school choice market design take preferences of families as given. All else equal, families tend to prefer schools that are closer to home. The relationship between the housing market and school assignment rules is widely documented (see, e.g., Black (1999), Bayer, Ferreira, and McMillan (2007)), and a central aspiration of school choice plans is to delink residential location from access to schools. It is therefore important to consider how residential choices might be influenced by school assignment when comparing market designs.

Avery and Pathak (2015) develop a model to compare neighborhood and school choice when households make a simultaneous decision about schools and residential choices.¹² Though their model of school choice abstracts away from many market design elements, it highlights a counterintuitive force behind school choice when also considering residential choices. Compared to neighborhood assignment, choice results in a compression of school qualities in a city, which is reflected in housing prices. This compressed distribution generates incentives for both the highest and lowest types to move out of cities with school choice, typically producing worse outcomes for low types than neighborhood assignment rules. Paradoxically, even when choice results in improvement in the worst performing schools, the lowest type residents may not benefit. By incorporating feedback between residential and school choices, the model suggests that analysis of school assignment that does not account for the

¹²This paper builds on an earlier literature that is not focused on market design aspects of school assignment, but does integrate housing and school choices, such as Epple and Romano (2003) and Nechyba (2003). These papers incorporate multi-dimensional student types, define school quality as functions of tax funding and peers, and allow for property taxes and the housing market to be determined endogenously, often relying on computational methods.

possibility of residential resorting may lead to an incomplete understanding about the consequences of school choice.

Aside from proximity to schools, recent work has started to investigate the extent to which student preferences reflect productive dimensions of school quality. Proponents of school choice as a reform strategy embrace the market paradigm: when families act like consumers when choosing schools, it can unleash valuable competitive forces as schools compete to attract parents. Rothstein (2006), on the other hand, argues that parental demand may reflect peers rather than value added, and such sorting may not generate productivity improvements. MacLeod and Urquiola (2012) show that with imperfect information schools may have incentives to invest in their reputation, which may not be enhance their productivity. The extent to which preferences reflect what a system operator might value for productive reasons (such as school value-added) remains an open question.

In the context of high-performing charter schools, Walters (2014) shows that the applicants who are most likely to benefit are actually those who are the least likely to apply. This evidence comes from a decentralized application process, so it's possible that a centralized system may make it easier for disadvantaged applicants to access higher-quality providers. Abdulkadiroğlu, Angrist, Hull, and Pathak (2014) compare applicants who arrive at high-performing charter schools passively to those who apply via lottery and find larger gains for the former, suggesting large potential gains from making it easier to exercise choice. However, datasets from centralized market designs consistently show that more disadvantaged families place more weight on proximity than various measures of school quality (see, e.g., Hastings, Kane, and Staiger (2009), Pathak and Shi (2014), and Abdulkadiroglu, Agarwal, and Pathak (2015)). Such a finding implies that choice reforms might exacerbate inequalities.

While many design initiatives aspire to make school selections easier, the ranking decision still remains complex for many families, even in strategy-proof mechanisms. Hastings and Weinstein (2008) report on a field experiment showing that small information queues can change how parents rank schools in Charlotte's variant of the Boston mechanism. Efforts to improve how participants interact with market designs, including decision aids and other informational interventions, hold great promise to complement research on market clearing algorithms.

4.2 Improving School Quality

On the school side, encouraging participation of highly-desired schools diminishes the importance of a market clearing algorithm. Participation by schools in a centralized system, however, is not always an easy issue to navigate in practice. Charter and other autonomous schools often view centralized assignment as a threat to their autonomy. For instance, in New Orleans, there are two school districts, the Recovery School District and the Orleans Parish School Board (OPSB). OPSB schools have been

reluctant to join the RSD’s OneApp process because school administrators like the ability to screen applicants, including groups of applicants (for arts and music and sports programs). There is an ongoing debate about participation of OPSB schools, and OPSB charter schools are currently mandated to participate in the system upon re-authorization (Dreilinger 2013). Ekmekci and Yenmez (2014) study participation by schools in a centralized clearinghouse. After showing that schools may have an incentive not to participate, they propose modifications to current mechanisms that encourage greater participation. Relatedly, Hatfield, Kojima, and Narita (2014) study which market clearing algorithms influence a school’s incentive for improving quality.

Market design schemes hold great potential to influence the portfolio of schools in a district indirectly through the data they generate. First, centralized systems systematically elicit preferences for schools, which may be used to guide enrollment and planning. For instance, if a school is perennially undersubscribed, then a district might consider intervening in the school, either by changing staff, educational programming, or by closing the school altogether. Often efforts to reconstitute schools require the district to decide which schools to target (see, e.g., Abdulkadiroğlu, Angrist, Hull, and Pathak (2014), who report on charter school takeovers of traditional district schools in New Orleans and in Boston). The ranking data available from strategy-proof mechanisms provides complementary information about perceived school desirability that may allow for more sophisticated systems of school accountability and portfolio planning. When school supply decisions do not reflect any information on school demand, it seems like a missed opportunity to harness competitive forces that could generate productivity improvements.

Second, the data generated by centralized systems allow for more effective ways to measure the performance of schools or school sectors. Research exploiting this aspect of school choice market designs has become increasingly popular. A centralized assignment system not only generates systematic admissions data, but it often lends itself to natural research designs. Abdulkadiroğlu, Angrist, Dynarski, Kane, and Pathak (2011) exploit random assignment embedded in DA to measure the effects of attendance on student achievement at one Boston’s pilot schools. Such schools have some of the independence of charter schools, but they are regulated by Boston Public schools and are covered by some of its collective bargaining provisions. The paper reports mixed results for pilots compared to results for Boston’s charter schools, which generate impressive achievement results on statewide assessments. Using the DA-induced variation in New York City’s centralized match, Abdulkadiroğlu, Hu, and Pathak (2013) estimate the effects of New York’s small schools, more than 200 of which were created in the 2000s. The paper shows that small schools created as part of the NYC’s *Children’s First Initiative* in 2002 produce significant achievement gains on standardized assessments and lead to more high schoolers attending college. Abdulkadiroğlu, Angrist, and Pathak (2014) use data from

a centralized exam school to construct regression discontinuity estimates of value added of New York and Boston’s elite exam schools. Even though these schools are heavily over-subscribed, the estimates show little evidence of value-added. These three studies are examples of how evidence on the performance of particular school sectors generated via a centralized assignment mechanism can yield new information that system operators can use for decisions about what schools to expand or contract.

Centralized school market designs also allow researchers to sidestep a major difficulty that has hindered evaluations in decentralized systems in the past: poor record-keeping on how students were admitted into schools. For instance, Abdulkadiroğlu, Angrist, Dynarski, Kane, and Pathak (2011) collected admissions records individually from Boston’s charter schools, since each runs their own admissions lottery with large variation in record-keeping.¹³ Had these schools been part of a unified enrollment system, it would be relatively simple to understand how a given applicant was assigned to a particular school. Understanding the admissions process in this precise way, then, allows for quasi-experimental approaches to measuring the performance of schools. Examples of papers exploiting data from centralized systems include Dobbie and Fryer (2014), Ajayi (2014), Lucas and Mbiti (2014), Pop-Eleches and Urquiola (2013), Jackson (2010), Bergman (2014), Hastings, Neilson, and Zimmerman (2013), Kirkeboen, Leuven, and Mogstad (2015), Angrist, Hull, Pathak, and Walters (2015) and Abdulkadiroglu, Pathak, and Walters (2015).

Abdulkadiroğlu, Angrist, Narita, and Pathak (2015) develop econometric techniques to use data from a DA match to measure school effectiveness in the best possible way. Since assignments depend on both non-random student preferences and school priorities, centralized assignment schemes generate complex stratified random assignments to schools. After developing easily-implemented empirical strategies that fully exploit the random assignment embedded in DA, the paper estimates large achievement gains from charter school attendance in Denver, which uses DA to assign both charter and traditional district schools. Compared to *ad hoc* methods that fail to exploit the full richness of the lotteries generated by centralized assignment with random tie-breaking, the new method results in substantial efficiency gains. This paper therefore shows how market design feeds into powerful research designs for credible impact evaluation.

In summary, school choice market design has had a positive spillover on work that seeks to understand the impact of schools on student outcomes. Although this work has not been explicitly linked to the actual assignment algorithm, it has the potential to guide the district’s offering of schools by exploiting the data emerging from well-designed assignment schemes. This feedback may ultimately lead to improved educational outcomes.

¹³Lack of systematic record-keeping was also a significant hurdle for charter school evaluations in Angrist, Dynarski, Kane, Pathak, and Walters (2012) Angrist, Dynarski, Kane, Pathak, and Walters (2012) and Angrist, Pathak, and Walters (2013).

5 Conclusion

Much of the school choice market design literature evaluates different ways to assign students given student preferences, school priorities, and school capacities. This literature has generated numerous insights on the tension between efficiency and fairness, the role of incentives, the implications of coarseness in school priorities, and constraints on rank order lists. Experiences from the field highlight new issues that have not been the focus of this earlier literature such as non-consequentialist rationales for straightforward incentives, the importance of transparency and simplicity in influencing designs, the value of single-offer systems over multiple-offer alternatives, the need to streamline aftermarkets, and the importance of participation on the school and student side. It is my hope that further interaction between theory and practice will sharpen focus on these and potentially other significant issues that have received relatively little attention from the theoretical literature.

While much of the market design literature has taken student preferences and schools as given, I believe there is great potential for work that examines the feedback between market clearing algorithms and these aspects of demand and supply. Without a deeper investigation of these broader aspects of school choice markets, our understanding of what really matters in designing school choice mechanisms will be incomplete.

Table 1. Comparing Deferred Acceptance and Top Trading Cycles

Choice Assigned	Boston Elementary, Middle and High School		New Orleans OneApp (All Grades)	
	TTC (1)	DA (2)	TTC (3)	DA (4)
1	65.4%	64.8%	73.2%	72.0%
2	17.9%	18.8%	11.7%	13.1%
3	7.8%	8.1%	3.3%	3.9%
4	3.3%	3.4%	1.0%	1.2%
5	1.3%	1.3%	0.4%	0.4%
6	0.4%	0.4%	0.1%	0.1%
7	0.2%	0.2%	0.0%	0.0%
8	0.1%	0.1%	0.0%	0.0%
9	0.0%	0.0%	0.0%	0.0%
10	0.0%	0.0%	0.0%	0.0%
Unassigned	3.5%	2.9%	10.3%	9.3%
Number of Students	5,927	5,927	7,789	7,789
Number of Students with Justified Envy	6.8%		10.0%	

Notes. Data covers four schools years, 2009-2010 to 2012-2013, from Boston Public Schools and one year, 2012-2013, from the New Orleans Recovery School District. Boston data are from grades K2, 6, and 9, while New Orleans data are from grades PK-9. DA is the student-proposing deferred acceptance algorithm. TTC is Gale's Top Trading Cycles algorithm as defined by Abdulkadiroglu and Sonmez (2003). The same lottery number is used for TTC and DA.

Table 2. Tie-breaking with Deferred Acceptance

Choice Assigned	Boston Elementary			NYC High Schools		
	DA-STB	DA-MTB	Student-Optimal	DA-STB	DA-MTB	Student-Optimal
	(1)	(2)	(3)	(4)	(5)	(6)
1	77%	74%	77%	41%	38%	42%
2	11%	12%	11%	18%	18%	18%
3	5%	6%	5%	12%	13%	12%
4	2%	3%	2%	8%	8%	8%
5	1%	1%	1%	5%	6%	5%
6	0%	0%	0%	3%	4%	3%
7	0%	0%	0%	2%	2%	2%
8	0%	0%	0%	1%	2%	1%
9	0%	0%	0%	1%	1%	1%
10	0%	0%	0%	1%	1%	1%
11	0%	0%	0%	0%	0%	0%
12	0%	0%	0%	0%	0%	0%
Unassigned	4%	4%	4%	7%	7%	7%
N	2927	2927	2927	78729	78728	78729

Notes. Extracted from Abdulkadiroglu, Pathak, Roth (2009). Boston and NYC data are from 2006-07. Boston data is for Elementary (K2), while NYC is for high school applicants. DA-STB is the student-proposing deferred acceptance algorithm with single tie breaking, DA-MTB is the student-proposing deferred acceptance algorithm with school-specific tie-breaking, and Student-Optimal Stable Matching is the result of the Erdil-Ergin (2008) Stable Improvement Cycles from the DA-STB allocation. Each column reports the average from 250 lottery draws.

Table 3. Constrained Rank Order Lists

Application Year	Constraint (1)	Number of Options (2)	Fraction at Constraint (3)
A. New York City High Schools			
2006	12	777	22.85%
2007	12	795	20.67%
2008	12	781	17.70%
B. Chicago Public Schools			
2009	4	9	84.60%
2010	6	10	56.80%
2011	6	10	49.50%
2012	6	10	46.60%
2013	6	10	46.02%
C. Denver Public Schools			
2012	5	78	26.80%
2013	5	81	26.60%
D. New Orleans RSD			
2013	8	41	4.97%

Notes. Denver applicants are from major transition grades (EC, K,6,9). RSD applicants are from transition grades (PK-9). The number of possible choices is equally-weighted across these applicant grades. Chicago Public Schools applicants are non-special education applicants in Chicago Public Schools in 8th grade.

Table 4. Chicago Preference Changers

Applicants who:	Change Choice (1)	Change School (2)
1	1.33%	5.29%
2	3.48%	1.62%
3	3.74%	0.38%
4	2.38%	0.04%
No change	89.07%	92.68%

Notes. N=13,121. This table tabulates the fraction of applicants to Chicago's selective enrollment high schools who change a listed choice in 2009.

Table 5. Attributes of Preferences Changers

	Change at least one choice (1)	Change at least one school (2)
Admission test points	55.157*** (4.683)	28.481*** (5.496)
Baseline Math	0.178*** (0.026)	0.082*** (0.030)
Baseline Reading	0.122*** (0.022)	0.028 (0.026)
White	-0.007 (0.008)	-0.031*** (0.008)
Black	0.047*** (0.014)	0.078*** (0.017)
Hispanic	-0.060*** (0.014)	-0.056*** (0.016)
Lowest SES Census Tracts	-0.016 (0.012)	-0.006 (0.014)
Second Quartile Census Tracts	-0.018 (0.012)	-0.008 (0.015)
Third Quartile Census Tracts	0.026** (0.013)	0.007 (0.015)
Highest SES Census Tracts	0.009 (0.012)	0.007 (0.014)
Days Absent	-1.108*** (0.161)	-0.669*** (0.200)

Notes. N=13121. Baseline Math is from grade 8 ISAT Math. Baseline Reading is grade 8 ISAT Reading. Baseline scores and days absent are not available for all applicants. Standard errors in parentheses, * p<0.1, ** p<0.05, *** p<0.01.

Figure 1. Length of Applicant Preference Lists in Boston




Figure 2. Fraction Receiving Top Choice in Boston



References

- ABDULKADIROĞLU, A., N. AGARWAL, AND P. PATHAK (2015): “The Welfare Effects of Coordinated Assignment: Evidence from the NYC HS Match,” NBER Working Paper 21046.
- ABDULKADIROĞLU, A., P. A. PATHAK, AND C. WALTERS (2015): “School Vouchers and Student Achievement: First-Year Evidence from the Louisiana Scholarship Program,” NBER Working Paper 21839.
- ABDULKADIROĞLU, A., AND T. SONMEZ (1998): “Random Serial Dictatorship and the Core from Random Endowments in House Allocation Problems,” *Econometrica*, 66(3), 689–701.
- ABDULKADIROĞLU, A., J. ANGRIST, AND P. PATHAK (2014): “The Elite Illusion: Achievement Effects at Boston and New York Exam Schools,” *Econometrica*, 82(1), 137–196.
- ABDULKADIROĞLU, A., J. D. ANGRIST, S. M. DYNARSKI, T. J. KANE, AND P. A. PATHAK (2011): “Accountability and Flexibility in Public Schools: Evidence from Boston’s Charters and Pilots,” *Quarterly Journal of Economics*, 126(2), 699–748.
- ABDULKADIROĞLU, A., J. D. ANGRIST, P. HULL, AND P. PATHAK (2014): “Charters without Lotteries: Testing Takeovers in New Orleans and Boston,” NBER Working Paper 20792.
- ABDULKADIROĞLU, A., J. D. ANGRIST, Y. NARITA, AND P. PATHAK (2015): “Research Design meets Market Design: Using Centralized Assignment for Impact Evaluation,” NBER Working Paper, 21705.
- ABDULKADIROĞLU, A., Y.-K. CHE, P. PATHAK, A. ROTH, AND O. TERCIEUX (2015): “Minimal Envy Efficient Assignment Mechanisms: The Design of New Orleans OneApp,” Working paper, Columbia University.
- ABDULKADIROĞLU, A., Y.-K. CHE, AND Y. YASUDA (2011): “Resolving Conflicting Preferences in School Choice: The “Boston Mechanism” Reconsidered,” *American Economic Review*, 101(1), 399–410.
- ABDULKADIROĞLU, A., W. HU, AND P. PATHAK (2013): “Small High Schools and Student Achievement: Lottery-Based Evidence from New York City,” NBER Working paper 19576.
- ABDULKADIROĞLU, A., P. A. PATHAK, AND A. E. ROTH (2005): “The New York City High School Match,” *American Economic Review, Papers and Proceedings*, 95, 364–367.

- (2009): “Strategy-proofness versus Efficiency in Matching with Indifferences: Redesigning the New York City High School Match,” *American Economic Review*, 99(5), 1954–1978.
- ABDULKADIROĞLU, A., P. A. PATHAK, A. E. ROTH, AND T. SÖNMEZ (2005): “The Boston Public School Match,” *American Economic Review, Papers and Proceedings*, 95, 368–371.
- (2006): “Changing the Boston School Choice Mechanism,” NBER Working paper, 11965.
- ABDULKADIROĞLU, A., AND T. SÖNMEZ (1999): “House Allocation with Existing Tenants,” *Journal of Economic Theory*, 88, 233–260.
- (2003): “School Choice: A Mechanism Design Approach,” *American Economic Review*, 93, 729–747.
- (2013): “Matching Markets: Theory and Practice,” D. Acemoglu, M. Arellano, and E. Dekel (eds), *Advances in Economics and Econometrics*, Econometric Society, vol. 1, 3–47.
- AGARWAL, N., AND P. SOMAINI (2014): “Demand Analysis Using Strategic Reports: An Application to a School Choice Mechanism,” NBER Working Paper 20775.
- AJAYI, K. (2014): “Does School Quality Improve Student Performance? New Evidence from Ghana,” Working paper, Boston University.
- ANGRIST, J., P. HULL, P. PATHAK, AND C. WALTERS (2015): “Leveraging Lotteries for Value Added: Testing and Estimation,” NBER Working Paper 21748.
- ANGRIST, J. D., S. M. DYNARSKI, T. J. KANE, P. A. PATHAK, AND C. R. WALTERS (2012): “Who Benefits from KIPP?,” *Journal of Policy Analysis and Management*, 31(4), 837–860.
- ANGRIST, J. D., P. A. PATHAK, AND C. R. WALTERS (2013): “Explaining Charter School Effectiveness,” *American Economic Journal: Applied Economics*, 5(4), 1–27.
- AVERY, C. N., C. NICKOLAUS, AND P. A. PATHAK (2015): “The Common Application and the Deferred Acceptance Algorithm,” Working paper, MIT.
- AVERY, C. N., AND P. A. PATHAK (2015): “The Distributional Consequences of Public School Choice,” NBER Working Paper 21525.
- BALINSKI, M., AND T. SÖNMEZ (1999): “A Tale of Two Mechanisms: Student Placement,” *Journal of Economic Theory*, 84, 73–94.

- BAYER, P., F. FERRIERA, AND R. MCMILLAN (2007): “A Unified Framework for Measuring Preferences for Schools and Neighborhoods,” *Journal of Political Economy*, pp. 588–638.
- BERGMAN, P. (2014): “Educational Attainment and School Desegregation: Evidence from Randomized Lotteries,” Working Paper.
- BLACK, S. (1999): “Do Better Schools Matter? Parental Valuation of Elementary Education,” *Quarterly Journal of Economics*, 114(2), 577–599.
- BPS (2005): “Recommendation to Implement a New BPS Assignment Algorithm,” Presentation to the Boston School Committee by Carleton Jones, May 11.
- (2008): “Introducing the Boston Public Schools: A Guide for Parents and Students,” Available from Boston Public Schools.
- CALSAMIGLIA, C., C. FU, AND M. GUELL (2014): “Structural Estimation of a Model of School Choices: the Boston Mechanism vs. Its Alternatives,” Working Paper, University of Wisconsin.
- CALSAMIGLIA, C., G. HAERINGER, AND F. KLJIN (2010): “Constrained School Choice: An Experimental Study,” *American Economic Review*, 100(4), 1860–1874.
- CALSAMIGLIA, C., AND A. MIRALLES (2014): “Catchment Areas and Access to Better Schools,” Working Paper, Universitat Autonomy de Barcelona.
- CARROLL, G. (2014): “A General Equivalence Theorem for Allocation of Indivisible Objects,” *Journal of Mathematical Economics*, 51, 163–177.
- CHE, Y.-K., AND Y. KOH (2015): “Decentralized College Admissions,” forthcoming, *Journal of Political Economy*.
- CHEN, Y., AND T. SÖNMEZ (2006): “School Choice: an Experimental Study,” *Journal of Economic Theory*, 127, 202–231.
- COOK, G. (2003): “School assignment flaws detailed,” *Boston Globe*, pp. Metro Desk, September 12.
- DE HAAN, M., P. A. GAUTIER, H. OOSTERBEEK, AND B. VAN DER KLAAUW (2015): “The Performance of School Assignment Mechanisms in Practice,” Working Paper.
- DFES (2007): “School Admissions Code, Department for Education and Skills,” Available at: <http://www.dcsf.gov.uk/sacode/>, Last accessed: January 31, 2011.

- DOBBIE, W., AND R. G. FRYER (2014): “Exam High Schools and Academic Achievement: Evidence from New York City,” *American Economic Journal: Applied Economics*, 6(3), 58–75.
- DREILINGER, D. (2013): “Orleans Parish School Board pulls back on OneApp, lets schools choose students,” Times-Picayune, June 18.
- DUR, U. (2014): “A Characterization of the Top Trading Cycles Mechanism in the School Choice Problem,” Working paper, NC State.
- DUR, U., R. HAMMOND, AND T. MORRILL (2015): “Identifying the Harm of Manipulable School Choice Mechanisms,” Working paper, NC State.
- DUR, U., S. KOMINERS, P. PATHAK, AND T. SÖNMEZ (2014): “The Demise of Walk Zones in Boston: Priorities vs. Precedence in School Choice,” NBER Working Paper 18981.
- DUR, U., P. PATHAK, AND T. SÖNMEZ (2015): “Equity vs. Merit in Affirmative Action: Explicit and Statistical Preferential Treatment at Chicago’s Exam Schools,” Unpublished working paper, MIT.
- ECHENIQUE, F., AND B. YENMEZ (2014): “How to Control Controlled School Choice?,” forthcoming, *American Economic Review*.
- EHLERS, L., AND A. ERDIL (2010): “Efficient Assignment Respecting Priorities,” *Journal of Economic Theory*, 145(3), 1269–1282.
- EKMEKCI, M., AND M. B. YENMEZ (2014): “Integrating Schools for Centralized Admissions,” Working paper, Carnegie-Mellon.
- ELLISON, G., AND P. PATHAK (2015): “The Inefficiency of Race-Neutral Affirmative Action,” Working paper, MIT.
- EPPLE, D., AND R. ROMANO (2003): “Neighborhood Schools, Choice, and the Distribution of Educational Benefits,” in *The Economics of School Choice*, ed. by C. Hoxby. University of Chicago Press.
- ERDIL, A., AND H. ERGIN (2008): “What’s the Matter with Tie-Breaking? Improving Efficiency in School Choice,” *American Economic Review*, 98, 669–689.
- ERDIL, A., AND T. KUMANO (2012): “Prioritizing Diversity in School Choice,” Unpublished working paper, Washington University.

- ERGIN, H. (2002): “Efficient Resource Allocation on the Basis of Priorities,” *Econometrica*, 70, 2489–2498.
- ERGIN, H., AND T. SÖNMEZ (2006): “Games of School Choice under the Boston Mechanism,” *Journal of Public Economics*, 90, 215–237.
- GALE, D., AND L. S. SHAPLEY (1962): “College Admissions and the Stability of Marriage,” *American Mathematical Monthly*, 69, 9–15.
- HAERINGER, G., AND F. KLIJN (2009): “Constrained School Choice,” *Journal of Economic Theory*, 144, 1921–1947.
- HAFALIR, I., B. YENMEZ, AND M. A. YILDIRIM (2012): “Effective Affirmative Action in School Choice,” *Theoretical Economics*, forthcoming.
- HASTINGS, J., T. J. KANE, AND D. O. STAIGER (2009): “Heterogenous Preferences and the Efficacy of Public School Choice,” Working paper, Yale University.
- HASTINGS, J., C. NEILSON, AND S. ZIMMERMAN (2013): “Are Some Degrees Worth More than Others? Evidence from College Admission Cutoffs in Chile,” NBER Working paper 19241.
- HASTINGS, J., AND J. M. WEINSTEIN (2008): “Information, School Choice and Academic Achievement: Evidence from Two Experiments,” *Quarterly Journal of Economics*, 123(4), 1373–1414.
- HATFIELD, J. W., F. KOJIMA, AND Y. NARITA (2014): “Promoting School Competition Through School Choice: A Market Design Approach,” Working paper.
- HE, Y. (2012): “Gaming the Boston School Choice Mechanism in Beijing,” Unpublished working paper, Toulouse School of Economics.
- HWANG, S. I. (2015): “A Robust Redesign of High School Match,” Working paper, University of Chicago.
- I-CHIA, L. (2014): “Parents, teachers protest senior-high entrance process,” Taipei Times, Front Page, June 22.
- INSPECTOR-GENERAL (2010): “Office of the Inspector General: Chicago Board of Education, Annual Report, July 1, 2009 - June 30, 2010,” City of Chicago.
- JACKSON, K. (2010): “Do Students Benefit from Attending Better Schools? Evidence from Rule-based Student Assignments in Trinidad and Tobago,” *Economic Journal*, 120(549), 1399–1429.

- JOHNSON, C. R. (2013): “Speech to Boston School Committee, March 13,” Available at: <http://bostonschoolchoice.files.wordpress.com/2013/03/3-13-13-superintendent-sc-memo-on-assignment.pdf>, last accessed: April 14, 2013.
- KESTEN, O. (2006): “On two competing mechanisms for priority-based allocation problems,” *Journal of Economic Theory*, 127, 155–171.
- (2010): “School Choice with Consent,” *Quarterly Journal of Economics*, pp. 1297–1394.
- KESTEN, O., AND U. ÜNVER (2015): “A Theory of School Choice Lotteries,” *Theoretical Economics*, 10, 543–595.
- KIRKEBOEN, L., E. LEUVEN, AND M. MOGSTAD (2015): “Field of Study, Earnings, and Self-Selection,” Working Paper.
- KLEMPERER, P. (2002): “What Really Matters in Auction Design,” *Journal of Economic Perspectives*, 16(1), 169–189.
- KOJIMA, F. (2012): “School Choice: Impossibilities for Affirmative Action,” *Games and Economic Behavior*, 75(2), 685–693.
- KOMINERS, S. D., AND T. SÖNMEZ (2014): “Matching with Slot-Specific Priorities: Theory,” forthcoming, *Theoretical Economics*.
- LANDSMARK, T., H. DAJER, AND L. GONSALVES (2014): “Report and Recommendations of the Boston Public Schools Student Assignment Task Force,” Report to the Boston School Committee, September 22.
- LEE, T., AND J. SETHURAMAN (2014): “Equivalence results in the allocation of indivisible objects: A unified view,” Working paper, Columbia University.
- LUCAS, A., AND I. MBITI (2014): “Effects of School Quality on Student Achievement: Discontinuity Evidence from Kenya,” *American Economic Journal: Applied Economics*, 6(3), 234–63.
- MACLEOD, W. B., AND M. URQUIOLA (2012): “Anti-Lemons: School Reputation, Relative Diversity, and Educational Quality,” IZA Discussion Paper 6805.
- MCGREGOR, C. (2007): “The REAL Truth about Seattle Public Schools Kindergarten Assignments,” Available at <http://www.cybermato.com/projects/school-assignments/>, Last accessed: December 12, 2011.

- MORRILL, T. (2014): “Making School Assignments Just,” Working paper, NC State.
- NARITA, Y. (2015): “Match or Mismatch: Learning and Inertia in School Choice,” Working paper, MIT.
- NECHYBA, T. (2003): “Introducing School Choice into Multi-District Public School Systems,” in *The Economics of School Choice*, ed. by C. Hoxby. University of Chicago Press.
- OF APPEALS, T. U. S. C. (2001): “Parents Involved in Community Schools vs. Seattle School District,” Page 58 of the U.S. Court of Appeals for the Ninth Circuit, No. 01-35450, No. 1.
- PAPAI, S. (2000): “Strategyproof Assignment by Hierarchical Exchange,” *Econometrica*, 68, 1403–1433.
- PATHAK, P., F. SONG, AND T. SONMEZ (2015): “Turmoil in Taiwan: Preference-Driven Priority Mechanisms in School Choice,” Working paper, MIT.
- PATHAK, P. A. (2011): “The Mechanism Design Approach to Student Assignment,” *Annual Reviews*, 3, 513–536.
- PATHAK, P. A., AND J. SETHURAMAN (2011): “Lotteries in Student Assignment: An Equivalence Result,” *Theoretical Economics*, 6(1), 1–17.
- PATHAK, P. A., AND P. SHI (2014): “Demand Modeling, Forecasting, and Counterfactuals, Part I,” NBER Working Paper 19859.
- PATHAK, P. A., AND T. SÖNMEZ (2008): “Leveling the Playing Field: Sincere and Sophisticated Players in the Boston Mechanism,” *American Economic Review*, 98(4), 1636–1652.
- (2013): “School Admissions Reform in Chicago and England: Comparing Mechanisms by their Vulnerability to Manipulation,” *American Economic Review*, 103(1), 80–106.
- PENNELL, H., A. WEST, AND A. HIND (2006): “Secondary School Admissions in London,” Centre for Educational Research, Department of Social Policy, LSE, Clare Market Papers, No. 19.
- POP-ELECHES, C., AND M. URQUIOLA (2013): “Going to a Better School: Effects and Behavioral Responses,” *American Economic Review*, 103(4), 1289–1324.
- ROTH, A. E. (1984): “The Evolution of the Labor Market for Medical Interns and Residents: A Case Study in Game Theory,” *Journal of Political Economy*, 92, 991–1016.

- (1991): “A Natural Experiment in the Organization of Entry Level Labor Markets: Regional Markets for New Physicians and Surgeons in the U.K.,” *American Economic Review*, 81, 415–440.
- ROTH, A. E., AND X. XING (1997): “Turnaround Time and Bottlenecks in Market Clearing: Decentralized Matching in the Market for Clinical Psychologists,” *Journal of Political Economy*, 105, 284–329.
- ROTHSTEIN, J. (2006): “Good Principals or Good Peers: Parental Valuation of School Characteristics, Tiebout Equilibrium, and the Incentive Effects of Competition Among Jurisdictions,” *American Economic Review*, 96(4), 1333–1350.
- SHAPLEY, L., AND H. SCARF (1974): “On Cores and Indivisibility,” *Journal of Mathematical Economics*, 1, 23–28.
- SHI, P. (2014): “Guiding School-Choice Reform through Novel Applications of Operations Research,” *Interfaces*, 45(2), 117–132.
- SMITH, A. (2007): “Schools admissions code to end covert selection,” *Guardian*, Education section, January 9.
- URBANDCMOMS.COM (2014): “More Info on Common Lottery Algorithm,” URL: <http://www.dcurbanmom.com/jforum/posts/list/354001.page#4517696>, January 8.
- WALTERS, C. R. (2014): “The Demand for Effective Charter Schools,” UC Berkeley, Working paper.
- YOUNG, P. (2003): “First choice for schools may go,” *Evening Chronicle*, Newcastle, November 26.