

Reserve Design: Unintended Consequences and the Demise of Boston's Walk Zones

Umut Dur

North Carolina State University

Scott Duke Kominers

Harvard University and National Bureau of Economic Research

Parag A. Pathak

Massachusetts Institute of Technology and National Bureau of Economic Research

Tayfun Sönmez

Boston College

We show that in the presence of admissions reserves, the effect of the precedence order (i.e., the order in which different types of seats are filled) is comparable to the effect of adjusting reserve sizes. Either low-

We thank Kamal Chavda, Carol Johnson, Carleton Jones, Tim Nicollette, and Jack Yessayan for their expertise and for granting permission to undertake this study. Edward L. Glaeser, Arda Gitmez, Fuhito Kojima, Yusuke Narita, Peng Shi, and several seminar audiences provided helpful comments. Kominers is grateful for the support of National Science Foundation grants CCF-1216095 and SES-1459912, the Harvard Milton Fund, an American Mathematical Society–Simons Foundation Travel Grant, the Ng Fund and the Mathematics in Economics Research Fund of the Harvard Center of Mathematical Sciences and Applications, and the Human Capital and Economic Opportunity Working Group sponsored by the Institute for New Economic Thinking. Much of this work was conducted while Kominers was a research scholar at the Becker Friedman Institute at the University of Chicago. Pathak and Sönmez's work was supported by NSF grant SES-1426566. Sönmez also acknowledges the research support of Goldman Sachs Gives via Dalinc Ariburnu—Goldman Sachs Faculty Research Fund.

Electronically published November 6, 2018
[*Journal of Political Economy*, 2018, vol. 126, no. 6]

© 2018 by The University of Chicago. All rights reserved. 0022-3808/2018/12606-0006\$10.00

ering the precedence of reserve seats at a school or increasing the school's reserve size weakly increases reserve-group assignment at that school. Using data from Boston Public Schools, we show that reserve and precedence adjustments have similar quantitative effects. Transparency about these issues—in particular, how precedence unintentionally undermined intended policy—led to the elimination of walk zone reserves in Boston's public school match.

I. Introduction

Admissions policies in school systems are often shaped by historical circumstances and modern-day compromises between competing interest groups. At many publicly funded Indian engineering colleges, for example, seats are reserved for applicants from disadvantaged caste and gender groups (see Bagde, Epple, and Taylor 2016). In the Indian system, an applicant from a disadvantaged group who qualifies for a school without invoking caste/gender priority is assigned one of the school's regular seats instead of a reserve seat; the reserve seats are held for students who otherwise would not be able to gain admission. The public school administration in Boston also devised a reserve scheme but based it on neighborhood boundaries rather than student types. The Boston policy came after 1970s-era court-ordered desegregation divided the city into geographically segregated communities. At each school in Boston, half of the seats were made open to all applicants, while the other half prioritized applicants from local neighborhoods. Unlike the Indian system, the Boston system filled reserve seats ahead of open seats.

Indian engineering admissions are decentralized in some states, while Boston's school choice program is centralized. Under both systems, however, there are two types of seats at each school—reserve seats and open seats—and it is common for some applicants to qualify for both seat types. When a student can be admitted to a school via multiple routes, an admissions policy must specify the relative precedence of different admissions tracks; in the cases of Indian engineering colleges and Boston public schools, this means that policy must account for the order in which reserve and open seats are processed. In this paper, we formally show that precedence plays a central role that is qualitatively and quantitatively similar to the impact of reserve sizes in achieving distributional objectives. We then relate these results to a recent policy discussion in Boston, showing how an oversight leading to the wrong precedence policy completely undermined the city's stated objectives in a subtle way.

Boston's 50-50 reserve–open seat split emerged from a citywide discussion after racial and ethnic criteria for school placement ended in 1999. Many stakeholders advocated abandoning school choice and returning to neighborhood schooling, but the school committee chose instead to

maintain school choice while making neighborhood, “walk zone” priority apply to 50 percent of each school’s seats (app. C excerpts the official policy). In popular accounts, the 50-50 seat split was described as “striking an uneasy compromise between neighborhood school advocates and those who want choice,” while the superintendent hoped the “plan would satisfy both factions, those who want to send children to schools close by and those who want choice” (Daley 1999, B1).

The fragile compromise between the pro-neighborhood schooling faction and the pro-school choice faction has resurfaced in numerous debates about Boston’s school admissions policies.¹ Boston Mayor Thomas Menino’s 2012 State of the City Address forcefully argued in favor of assigning students to schools closer to home.² Proposals from BPS and other community members became the center of a year-long, citywide school choice discussion that featured over 70 public meetings and input from more than 3,000 parents.³ Boston’s decision to revisit its reserve policy was partly motivated by a persistent empirical puzzle: while 50 percent of seats at each school were reserved for students living in the neighborhood/walk zone, the fraction of neighborhood students assigned to popular schools consistently hovered around 50 percent. With half the seats reserved for neighborhood students and the other half open to everyone, one would expect more than 50 percent neighborhood assignment, as Boston’s official policy suggests (see app. C).

In this paper, we show that Boston’s assignment puzzle was an unintended consequence of the chosen implementation of the walk zone reserve: because the precedence order filled reserve seats before open seats, the 50-50 compromise was completely subverted, resulting in an allocation almost indistinguishable from a counterfactual setting without any reserve seats at all.

Our first formal result shows that reserves and precedence are policy tools with similar qualitative effects for any given school. For any precedence order, replacing an open slot with a reserve slot weakly increases

¹ A December 2003 community engagement process in Boston considered six different proposals for alternative neighborhood zone definitions. However, the only recommendation adopted by the school committee was to switch the assignment algorithm (Landsmark and Dajer 2004; Abdulkadiroğlu et al. 2005). In 2009, Boston Public Schools (BPS) renewed the discussion with a proposal for a five-zone plan, which eventually was not approved (Vaznis 2009).

² Constituents had long believed that students were traveling too far to attend schools and sought to alter the plan to assign students to schools closer to home (Landsmark 2009).

³ For more on this debate, see the materials available at <http://bostonschoolchoice.org> and accounts by Goldstein (2012), Handy (2012), and Seelye (2012). In fall 2012, BPS proposed five different plans that all restricted participant choice by reducing the number of schools that students could rank; the idea behind these plans was to reduce the fraction of nonneighborhood applicants competing for seats at each school. (The initial plans suggested dividing the city into 6, 9, 11, or 23 zones, or doing away with school choice entirely and reverting to assignment based purely on neighborhood.)

the assignment of reserve-eligible applicants. Similarly, for any given reserve size, swapping the precedence order position of a reserve slot with that of a lower-precedence open slot weakly increases the assignment of reserve-eligible applicants. Next, we investigate how our within-school results extend to centralized assignment systems that use the deferred acceptance algorithm. We find that for a given school, increasing the number of reserve slots (relative to open slots) or raising the precedence of open slots (relative to reserve slots) increases admission of reserve-eligible applicants under deferred acceptance. This result is, to our knowledge, the first-ever comparative static result for multiagent priority improvements in matching models. Because of interactions across schools in the deferred acceptance algorithm, the comparative statics do not necessarily extend to aggregate increases in assignment of reserve-eligible applicants across all schools. However, even though pathological cross-school interactions are possible, they do not appear to be relevant in practice: Our comparative statics extend to the whole market in a two-school model, and we can also bound the worst case when reserves privilege the same group throughout the school system. Moreover, our theoretical analysis closely matches the empirical patterns observed in Boston: we show that Boston's implementation of the 50-50 reserve-open compromise was in practice closer to a 10-90 system once implemented.

This paper contributes to a broader agenda, examined in a number of recent papers, that introduces diversity concerns into the literature on school choice mechanism design (see, e.g., Erdil and Kumano 2012; Kojima 2012; Budish et al. 2013; Hafalir, Yenmez, and Yildirim 2013; Kominers and Sönmez 2013, 2016; Echenique and Yenmez 2015). When an applicant ranks a school with seats that employ different admissions criteria, it is as if she is indifferent between that school's seats. Therefore, our work parallels investigations of indifferences in school choice problems (see, e.g., Erdil and Ergin 2008; Abdulkadiroğlu, Pathak, and Roth 2009; Pathak and Sethuraman 2011). Yet results on school-side indifferences do not extend to indifferences in student preferences. Finally, our goal here is to establish comparative static results based on Boston's policy developments. In subsequent work, Dur, Pathak, and Sönmez (2016) characterized optimal admissions policies motivated by Chicago's place-based affirmative action system.

Our paper proceeds as follows. Section II describes the puzzle Boston faced in more detail. Section III formally studies admissions policies in which applicants can be admitted via multiple routes. Section IV examines how schools' admissions policies interact with a centralized admissions system based on deferred acceptance. Section V reports on data from Boston, and Section VI presents conclusions. All proofs are presented in appendix B. (All appendixes are available online.)

II. Motivation

A. Boston's "50-50 Puzzle"

Despite widespread perception and policy intent that, since 1999, the BPS school choice system had prioritized walk zone applicants, those applicants appear to have had little advantage in practice. Even though 50 percent of seats at each school were reserved for walk zone students, the assignment outcomes were close to those that would have arisen under a system without any walk zone reserve. To see this, we compute the fraction of students assigned to walk zone schools in Boston for the extreme case with no walk zone priority—the 0% Walk system.⁴ Table 1 shows that despite the 50 percent walk zone reserve, assignment outcomes under BPS's system are nearly identical to those under 0% Walk; they differ for only 3 percent of grade K1 students.

One might suspect that similarity between the BPS outcome and 0% Walk is driven by strong preferences among applicants for neighborhood schools, as such preferences would bring the two policies' outcomes close together. However, this is not the case. We compare the BPS outcome to the 100% Walk counterfactual in which all seats give priority to walk zone applicants. Under 100% Walk, 19 percent of grade K1 students obtain an assignment different from what they receive in the BPS outcome.⁵ Thus, the remarkable proximity of the BPS outcome and the 0% Walk ideal of school choice proponents neither suggests nor reflects negligible stakes in school choice.⁶ Rather, it presents a puzzle: Why does Boston's assignment mechanism result so closely resemble that of a system without any neighborhood priority, even though half of each school's seats prioritize neighborhood students? Or, more qualitatively, why did Boston's 50 percent reserve have so little impact in practice? Why did the policy not result in, say, an outcome halfway between 0% Walk and 100% Walk? To obtain intuition, we turn to a simple, single-school example that illustrates Boston's 50-50 seat split as implemented.

⁴ To compute counterfactual assignments, we use internal preference data from BPS and the same lottery numbers BPS used to break ties in its assignment system. It is worth noting that strategy-proofness (i.e., truthfulness) of the assignment mechanism used in Boston justifies recomputing the assignment without modeling how applicants might submit preferences under counterfactual mechanisms (see Abdulkadiroğlu et al. 2006, 2009; Pathak and Sönmez 2008; Agarwal, Abdulkadiroğlu, and Pathak 2017).

⁵ We have repeated our calculations for 500 different random lottery draws. Under 0% Walk, the average differences are 3 percent, 4 percent, and 1 percent for grades K1, K2, and 6, respectively. Under 100% Walk, the corresponding average differences are 20 percent, 18 percent, and 10 percent.

⁶ The patterns we observe are similar for grades above K1, with the smaller differences between 0% Walk and 100% Walk for higher grades driven by a larger share of continuing students who obtain guaranteed priority for higher grades.

TABLE 1
DIFFERENCES BETWEEN THE BOSTON 50-50 IMPLEMENTATION
AND ALTERNATIVE WALK ZONE RESERVE SIZES

	DIFFERENCE FROM BPS IMPLEMENTATION		
	No. Students (1)	0% Walk (2)	100% Walk (3)
A. Grade K1			
2009	1,770	46 3%	336 19%
2010	1,977	68 3%	392 20%
2011	2,071	50 2%	387 19%
2012	2,515	88 3%	504 20%
All	8,333	252 3%	1,619 19%
B. Grade K2			
2009	1,715	28 2%	343 20%
2010	1,902	62 3%	269 14%
2011	1,821	90 5%	293 16%
2012	2,301	101 4%	403 18%
All	7,739	281 4%	1,308 17%
C. Grade 6			
2009	2,348	54 2%	205 9%
2010	2,308	41 2%	171 7%
2011	2,073	4 0%	225 11%
2012	2,057	24 1%	247 12%
All	8,786	123 1%	848 10%

NOTE.—The table reports the fraction of applicants whose assignments differ between BPS's 50-50 implementation and two alternative mechanisms, one without any walk zone priority (0% Walk) and the other with walk zone priority at all seats (100% Walk).

B. A One-School Example

Consider a single school with 100 seats. Suppose there are 100 applicants with walk zone priority and 100 applicants without walk zone priority. A lottery used for tie breaking is such that, of the 100 applicants with

the highest lottery numbers, 50 are from the walk zone and 50 are not. Figure 1 illustrates the situation, with both walk zone applicants and non-walk zone applicants ordered by the random tiebreaker.

In panel A of figure 1, there is no walk zone priority at the school, so students are admitted solely on the basis of the random tiebreaker. Given the tiebreaker, the school admits an equal number of students from each group. That is, the school admits 50 students from the walk zone and 50 students from outside the walk zone.

In panel B of figure 1, half the seats grant walk zone priority and the other half are open. Under Boston's school choice system, students from both groups first apply to the walk zone half. For the walk zone half, students who have walk zone priority are admitted ahead of students who do not, and the admitted walk zone students are those with the most favorable random tiebreakers. Therefore, 50 students from the walk zone with the most favorable tiebreakers take up all of the seats in the walk zone half. Next, the remaining applicants from the walk zone—who have less favorable random tiebreakers—apply to the open half of the school together with all applicants from outside the walk zone. For the open seats, students are admitted only on the basis of the random tiebreaker. But at this point, the remaining walk zone applicants are disadvantaged because they have systematically less favorable tiebreakers; consequently, only non-walk zone applicants are assigned to the 50 seats in the open half. The final allocation results in half of the school's seats being assigned to walk zone applicants, with the remaining half assigned to applicants from outside the walk zone.

The preceding logic, illustrated in figure 1, shows how the 50-50 compromise can have the same outcome as a situation without any walk zone priority. However, our example is stylized in several ways: There are an equal number of applicants with walk zone priority and without it,⁷ and the tiebreaker has an equal number of students from each group among the top 100.⁸ Nevertheless, we capture the main intuition for the phenomenon documented in table 1.

Our example shows that the precedence order under which seats are processed significantly affects the outcome. Had all the applicants first applied to the open half, 75 walk zone applicants and 25 non-walk zone applicants would have been admitted—even holding fixed the 50-50 seat split. At the time of Menino's 2012 speech, precedence order's dramatic

⁷ It is easy to see that the same arguments work whenever there are more walk zone applicants than non-walk zone applicants. Moreover, if there are more non-walk zone applicants than walk zone applicants, the outcomes will differ only for a small set of applicants who are admitted at the end of the process.

⁸ The law of large numbers implies that this would be the expected outcome across repeated tiebreaker lottery realizations.

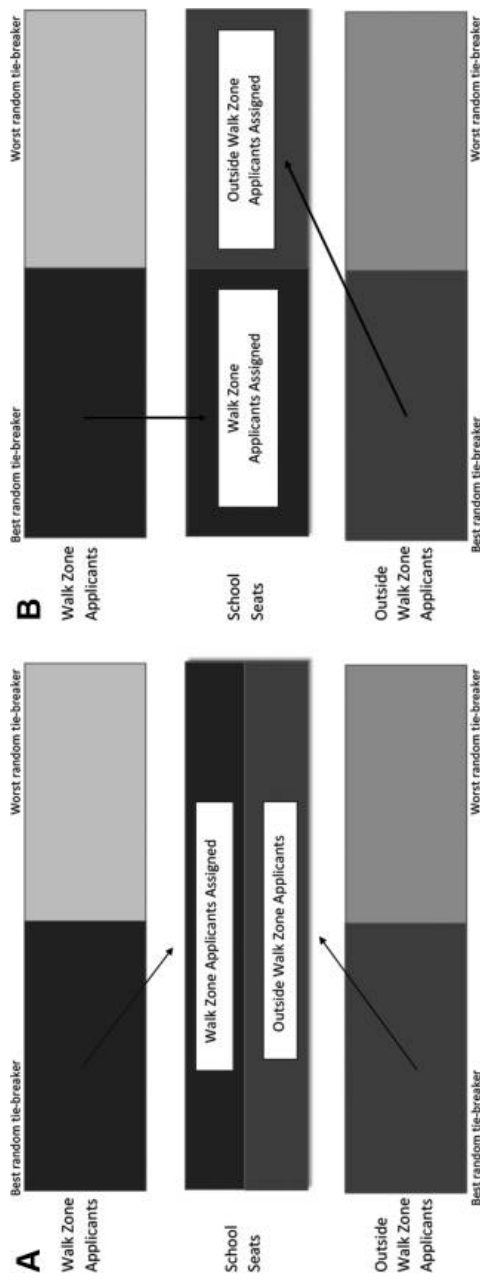


FIG. 1.—A single-school illustration comparing assignment without walk zone priority (left panel) to Boston's implementation of a 50-50 walk zone—open seat split (right panel). Color version available as an online enhancement.

role in disadvantaging walk zone students came as a surprise to many—including us—and motivated the formal analysis we now describe.

III. Admissions Policies with Reserves

To formalize the intuition presented in the preceding section, we develop a model of school admissions policies in which some seats at each school may be reserved for members of distinguished groups (e.g., disadvantaged castes or walk zone students). We prove comparative statics illustrating that both (1) increasing the number of reserve seats and (2) raising the precedence order positions of open seats will (weakly) increase the number of reserve-eligible students who are accepted.

A. Decentralized Model

There is a finite set I of students and a school a with a finite set of slots S^a . Each slot $s \in S^a$ has a linear priority order π^s over students in I . The linear priority order π^s captures the “property rights” of the students for slot s , in the sense that the higher a student is ranked under π^s , the stronger claim he or she has for slot s of school a . The total capacity of a is $q_a \equiv |S^a|$.

We are interested in situations in which slot priorities are heterogeneous. A consequence of such within-school heterogeneity is that we must determine how slots are assigned when a student is “qualified” for multiple slots that have different priority rankings. We suppose that the slots in S^a are ordered according to a (linear) *order of precedence* \triangleright^a . Given two slots $s, s' \in S^a$, the expression $s \triangleright^a s'$ means that slot s at school a is to be filled before slot s' whenever possible.

Given school a with set of slots S^a , profile of slot priorities $(\pi^s)_{s \in S^a}$, and order of precedence \triangleright^a with

$$s_a^1 \triangleright^a s_a^2 \triangleright^a \dots \triangleright^a s_a^{q_a},$$

the choice of school a from set of students J , denoted by $C^a(J)$, is obtained as follows: slots at school a are filled one at a time following the order of precedence \triangleright^a . The highest-priority student in J under $\pi^{s_a^1}$, say student j_1 , is chosen for slot s_a^1 of school a ; the highest-priority student in $J \setminus \{j_1\}$ under $\pi^{s_a^2}$ is chosen for slot s_a^2 , and so on.

We are particularly interested in slot priority structures in which some of the slots are reserved for applicants of a particular type (the “reserve-eligible”), while the remaining slots are open. Suppose there is a master priority order π^o that is uniform across all schools. This master priority is often determined by a random tiebreaker or by performance on an admissions exam (or in previous grades). For school a , there is a set $I_a \subseteq I$ of *reserve-eligible* students. Students who are not reserve-eligible are called *reserve-ineligible*. There are two types of slots:

1. Priorities at *open slots* correspond to the master priority order: $\pi^s = \pi^o$ for each open slot s .
2. Priorities at *reserve slots* grant all reserve-eligible students priority over all reserve-ineligible students, with the priority order within each group determined according to the master priority order π^o .

In Indian affirmative action systems, the reserve-eligible students are those from disadvantaged castes (Bagde et al. 2016). Aygün and Bo (2013) describe reserves for public universities in Brazil, where the reserve-eligible are racial minorities, applicants from low-income families, and applicants from public high schools. In BPS, the reserve-eligible groups are students who live in the school's walk zone, and thus, at times, we refer to BPS reserve slots as *walk zone seats* (and refer to BPS open slots as *open seats*).⁹

B. The Effects of Priority and Precedence Changes

We first examine the effects of increasing the reserve size given a precedence order. Suppose that slot s_* at school a is an open slot under priority profile π but is a reserve slot under priority profile $\tilde{\pi}$. Suppose that $\pi^s = \tilde{\pi}^s$ for all slots $s \neq s_*$. Let C^a and D^a , respectively, be the choice functions for a induced by the priorities π and $\tilde{\pi}$ under precedence order \succ^a . We obtain the following result.

PROPOSITION 1. Suppose that D^a is the choice function for school a obtained from C^a by changing an open slot to a reserve slot (fixing all other slots' priorities, as well as the precedence order). For any set of students $\bar{I} \subseteq I$,

- i. all students who are reserve-eligible at school a and are chosen from \bar{I} under choice function C^a are chosen under choice function D^a ;
- ii. all students who are reserve-ineligible at school a and are chosen from \bar{I} under choice function D^a are chosen under choice function C^a .

Proposition 1 states that when a school increases its reserve size, it admits weakly more reserve-eligible students and weakly fewer reserve-ineligible students. For Boston, this result suggests that increasing the walk zone percentage beyond 50 percent may increase neighborhood assignment.

What is much less apparent, however, is that swapping the precedence order of a reserve slot and a subsequent open slot has the same qualitative effect as increasing the reserve size. Suppose now that s_r is a reserve

⁹ BPS also uses sibling priority, but for our theoretical analysis we consider a simplified priority structure that depends only on walk zone status; using data from BPS, we show that this is a good approximation.

slot of school a that immediately precedes an open slot s_o under the precedence order \triangleright^a . Suppose, moreover, that precedence order $\tilde{\triangleright}^a$ is obtained from \triangleright^a by swapping the positions of s_r and s_o and leaving all other slot positions unchanged. Let C^a and D^a , respectively, be the choice functions for a induced by the precedence orders \triangleright^a and $\tilde{\triangleright}^a$ under slot priorities π . We obtain the following analogue to proposition 1.

PROPOSITION 2. Suppose that D^a is the choice function for school a obtained from C^a by swapping the precedence of a reserve slot and a subsequent open slot (fixing all slot priorities, as well as all other precedence order positions). For any set of students $\bar{I} \subseteq I$,

- i. all students who are reserve-eligible at school a and are chosen from \bar{I} under choice function C^a are chosen under choice function D^a ;
- ii. all students who are reserve-ineligible at school a and are chosen from \bar{I} under choice function D^a are chosen under choice function C^a .

Together, propositions 1 and 2 show how priority and precedence changes are substitute levers for influencing the assignment of reserve-eligible applicants. While the role of the number of reserve slots is quite apparent, the role of the order of precedence is much more subtle. Indeed, the choice of precedence order is often considered a minor technical detail, and, to our knowledge, precedence never explicitly entered school choice policy discussions until we raised the topic in Boston in parallel with the present work.

Qualitatively, the effect of decreasing a reserve slot's precedence order position is similar to the effect of replacing an open slot with a reserve slot. While this may initially appear counterintuitive, the reason is simple: decreasing the precedence of a reserve slot means that a reserve-eligible student with high enough master priority to be eligible for both open and reserve slots may now be assigned to an open slot. This in turn increases competition for open slots and decreases competition for reserve slots.

Our observation about how changing applicant processing orders influences access for reserve-eligible applicants also surfaced in debates on affirmative action policies in India. India's constitution stipulates that government-funded educational institutes and public-sector jobs, including seats in parliament, hold reservations for disadvantaged groups. In 1975, a debate about applicant processing made its way to the Supreme Court, where a judge ruled that the "benefits of the reservation shall be snatched away by the top creamy layer of the backward class, thus leaving the weakest among the weak and leaving the fortunate layers to consume the whole cake."¹⁰ In the context of our model, if reserve seats have higher

¹⁰ The court case is *State of Kerala vs. N. M. Thomas* (1974).

precedence than open seats, then priority goes to applicants who do not need it (the “creamy layer”), leaving the remaining reserve-eligible students without opportunities to obtain open seats because they are out-competed by the reserve-ineligible students.

So far, our results are for a single school with a given choice function; this analysis directly informs us about reserves implemented in decentralized admissions in India and elsewhere. Since many centralized systems can be seen as iterated applications of choice functions, our results also yield an approximation for those centralized systems. We next formally examine how our results extend to centralized systems that use the deferred acceptance algorithm.

IV. Centralized Admissions Systems with Reserves

Suppose now that there is a set of schools A . We use the notation a_0 to denote a “null school” representing the possibility of being unmatched; we assume this option is always available to all students. Let $S \equiv \bigcup_{a \in A} S^a$ denote the set of all slots (at non-null schools). Each school $a \in A$ has a reserve-eligible population $I_a \subseteq I$, slot priorities $(\pi^s)_{s \in S^a}$, precedence order \succ^a , and choice function C^a , as described in the preceding section. Meanwhile, each student i has a strict preference relation P^i over $A \cup \{a_0\}$ (with associated weak preference relation R^i). If a_0 is preferable to $a \in A$ under P^i , then we say that a is *unacceptable* to i .

A *matching* $\mu: I \rightarrow A \cup \{a_0\}$ is a function that assigns each student to a school (or the null school) so that no school is assigned to more students than it has slots. This model generalizes the school choice model of Abdulkadiroğlu and Sönmez (2003) in that it allows for heterogeneous priorities across a given school’s slots. Nevertheless, a mechanism based on the celebrated (student-proposing) deferred acceptance algorithm (Gale and Shapley 1962) easily extends to our model, given our earlier description of schools’ choice functions.

For a given profile of slot priorities $(\pi^s)_{s \in S}$ and an order of precedence \succ^a for each school $a \in A$, the *outcome of the (student-proposing) deferred acceptance mechanism* can be obtained as follows.

STEP 1: Each student i applies to his or her most-preferred school in $A \cup \{a_0\}$. Each school $a \in A$ with a set of step 1 applicants J_1^a tentatively holds the applicants in $C^a(J_1^a)$ and rejects the rest.¹¹

STEP ℓ : Each student rejected in step $\ell - 1$ applies to his or her most-preferred school in $A \cup \{a_0\}$ that has not yet rejected him or her. Each school $a \in A$ considers the set J_ℓ^a comprising the new applicants to a

¹¹ Here and in future steps, the null school a_0 always holds the full set of students who apply to it.

and the students held by a at the end of step $\ell - 1$, tentatively holds the applicants in $C^a(J_\ell^a)$, and rejects the rest.

The algorithm terminates after the first step in which no students are rejected and assigns students to the schools holding their applications at that time.

A. Comparative Statics for Deferred Acceptance

In the context of deferred acceptance, we look at a single school and consider the effects of replacing open slots with reserve slots and swapping the precedence order positions of reserve slots and lower-precedence open slots; we find that both changes weakly increase the number of reserve-eligible students assigned to the school.

PROPOSITION 3. Consider centralized assignment under (student-proposing) deferred acceptance, and fix a school $a \in A$.

- i. Replacing an open slot of school a with a reserve slot (fixing all other slots' priorities, as well as all precedence orders) weakly increases the number of reserve-eligible students assigned to school a .
- ii. Switching the precedence order position of a reserve slot of school a with the position of a subsequent open slot (fixing all slot priorities, as well as all other precedence order positions) weakly increases the number of reserve-eligible students assigned to school a .

Proposition 3 is analogous to the results for decentralized admission for a single school in propositions 1 and 2. However, proposition 3's proof is substantially more involved, as it is necessary to consider how changes at one school cascade through the system under the deferred acceptance. That is, when either the precedence order or reserve size changes, a different set of applicants may apply to school a , thereby initiating a sequence of applications to other schools at subsequent steps of the deferred acceptance algorithm. These subsequent applications need to be tracked carefully. Indeed, neither comparative static result in proposition 3 follows from earlier comparative static approaches used in simpler models (e.g., Balinski and Sönmez 1999) because our comparative static involves simultaneous priority improvements for multiple students. As a result, we have had to develop a new proof strategy, which may be of independent interest in settings involving multiagent priority improvements in matching models.

B. Aggregate Effects

Our analysis thus far has focused on assignment of reserve-eligible students to a particular school at which there is a change in reserve size or precedence. A natural question is whether increased reserve-eligible stu-

dent assignment at a particular school always translates into increased overall reserve-eligible assignment across all schools. As usual in going from partial to general equilibrium analysis, the aggregate comparative static is not a foregone conclusion. In particular, it is well known that in matching models, interactions across the market can lead to counter-intuitive overall predictions. The following example shows that our results for a single school do not always imply an aggregate increase in reserve-eligible assignment.

EXAMPLE 1. There are three schools, $A = \{k, l, m\}$. Schools k and m each have two slots and school l has three slots. There are seven students $I = \{i_1, i_2, i_3, i_4, i_5, i_6, i_7\}$. The reserve-eligible students are given by $I_k = \{i_1, i_7\}$, $I_l = \{i_2, i_3, i_4\}$, and $I_m = \{i_5, i_6\}$. The master priority π^o orders the students as

$$\pi^o : i_7 \succ i_2 \succ i_5 \succ i_3 \succ i_1 \succ i_6 \succ i_4.$$

The preference profile is as follows:¹²

P^k	P^m	P^l	P^l	P^k	P^m	P^l
k	m	l	l	k	l	k
l			m	l	m	

First consider the case in which school k 's first and school l 's second slots are reserve slots, and all other slots are open slots. The outcome of deferred acceptance for this case is¹³

$$\begin{pmatrix} i_1 & i_2 & i_3 & i_4 & i_5 & i_6 & i_7 \\ l & m & l & l & k & m & k \end{pmatrix}.$$

Observe that in addition to the two reserve slots assigned to reserve-eligible students i_4 and i_7 , two of the open slots (namely, those assigned to i_3 and i_6) are also assigned to reserve-eligible students. As such, four students are assigned to schools at which they are reserve-eligible.

Next, we replace the open slot at school k with a reserve slot, so that both slots at school k are reserve slots. We keep the slot sets and precedence orders of the other schools the same. The deferred acceptance outcome for the second case is

¹² In our notation here and in future examples, preference relations are read vertically, and we omit the null school (as well as unacceptable schools) from the end of each preference relation, for notational simplicity. Thus, for example, P^i as stated means that i_1 prefers k to l and finds only schools k and l acceptable (i.e., preferable to a_0).

¹³ Here and hereafter, we use this notation to indicate that i_1 is assigned to l , i_2 is assigned to m , and so forth.

$$\begin{pmatrix} i_1 & i_2 & i_3 & i_4 & i_5 & i_6 & i_7 \\ k & m & l & m & l & l & k \end{pmatrix}.$$

Observe that while all three reserve slots are assigned to reserve-eligible students (students i_1 , i_3 , and i_7), none of the open slots are assigned to reserve-eligible students. That is, the total number of reserve-eligible student assignments decreases when the open slot at school k is replaced with a reserve slot. \square

The preceding example illustrates that the direct “first-order” effect of a reserve change at a given school may be undone by the indirect effect on other schools. Moreover, it is easy to modify example 1 to show that when the precedence order position of a reserve slot is swapped with that of a subsequent open slot, the overall reserve-eligible student assignment need not increase (see example 2 in app. A). These negative findings highlight the complexity of distributional comparative statics in matching models with slot-specific priorities (see also Kominers and Sönmez 2013, 2016).

C. Aggregate Effects under Uniform Reserve Priority

One important feature of example 1 is that the set of reserve-eligible students differs by school. When reserves represent walk zone seats, as in Boston, we would expect reserves to differ by school since families are dispersed geographically (and thus live in different walk zones). However, in a case like India, in which the reserve is intended to remedy a nongeographical disadvantage, that is, membership in a particular caste, the set of reserve-eligible students is the same for each school.

If we have $I_a = I_{a'}$ for all pairs of schools $a, a' \in A$, then we say that we have *uniform reserve priority*.¹⁴ In case of uniform reserve priority, it is still possible that reserve-eligible assignment can decrease when an open slot is replaced with a reserve slot (and likewise when a reserve slot is swapped with a subsequent open slot; see example 3 in app. A). However, even in the worst-case scenario, only one fewer reserve-eligible student can be assigned under uniform reserve priority.

PROPOSITION 4. Consider centralized assignment under (student-proposing) deferred acceptance, and suppose that we have uniform reserve priority (i.e., for any two schools $a, a' \in A$, we have $I_a = I_{a'}$). Then

- i. replacing an open slot of a school with a reserve slot (fixing all other slots' priorities, as well as all precedence orders) cannot decrease the

¹⁴ Aygün and Turhan (2016) have described a centralized admissions procedure with uniform reserve priority in the Indian state of Maharashtra.

- total assignment of reserve-eligible students across all schools by more than 1;
- ii. switching the precedence order position of a reserve slot of a school with the position of a subsequent open slot (fixing all slot priorities, as well as all other precedence order positions) cannot decrease the total assignment of reserve-eligible students across all schools by more than 1.

In contrast to proposition 4, if we do not have uniform reserve priority, then it is easy to show that when an open slot is replaced with a reserve slot, the number of reserve-eligible students assigned can be reduced by more than 1 (see example 4 in app. A).

D. Aggregate Effects in the Two-School Case

When there are only two schools and there are enough slots for all the students, the effects of the reserve and precedence order changes described in Section IV.A can be sharpened: either change weakly increases total reserve-eligible assignment. For the next result, we assume that there are only two schools, that each student is reserve-eligible at exactly one school, that there are enough slots for all the students, and that all students rank both schools.

PROPOSITION 5. Consider centralized assignment under (student-proposing) deferred acceptance, and suppose that there are two (nonnull) schools, that each student is reserve-eligible at exactly one school, that there are enough slots for all the students, and that all students rank both schools. Then

- i. replacing an open slot of either school with a reserve slot (fixing all other slots' priorities, as well as all precedence orders) weakly increases the total assignment of students to schools at which they are reserve-eligible;
- ii. switching the precedence order position of a reserve slot of either school with the position of a subsequent open slot (fixing all slot priorities, as well as all other precedence order positions) weakly increases the total assignment of students to schools at which they are reserve-eligible;
- iii. the number of students assigned to their most-preferred schools is independent of both the number of reserve slots at each school and the precedence order profile.

The first two parts of proposition 5 show that when there are only two schools, the aggregate effects of the reserve and precedence changes examined for a particular school in proposition 3 extend across all schools.

The last part of proposition 5 shows that in the two-school case, changes in the reserve size or precedence orders are entirely distributional: both instruments leave unchanged the aggregate number of students obtaining their most-preferred schools.

Proposition 5 suggests a method to compute the policies with the minimum and maximum numbers of students assigned to schools at which they are reserve-eligible: In the two-school case, at least, the minimum (across all priority and precedence policies) is obtained when all slots are open slots; the maximum is obtained when all slots are reserve slots.

The analysis from the two-school case suggests that the outcomes illustrated in our negative examples require an elaborate sequence of applications and rejections involving more than two schools. We next turn to data from Boston and see that the results of proposition 5 better approximate Boston's situation than what is suggested by our negative examples, which depend critically on carefully constructed rejection chains.

V. Precedence and Reserves in Boston

Table 2 reports the number of walk zone students assigned to schools under different walk zone percentages using the same tiebreaker lottery numbers that BPS used.¹⁵ For each grade, more students are assigned to schools in their walk zones under 100% Walk than under 0% Walk. For grade K1, the range between the two reserve policies is 11.2 percent of all students, which corresponds to 938 students. The range is 9.3 percent for grade K2 and 5.4 percent for grade 6. Consistent with proposition 5, a higher reserve size corresponds to more walk zone assignments.

The *Walk-Open* precedence order has all walk zone seats precede open seats, while the *Open-Walk* precedence order has all open seats precede walk zone seats. Consistent with proposition 5, table 2 shows that with a 50-50 seat split, more walk zone students are assigned under Open-Walk precedence than under Walk-Open precedence. Moreover, the outcome of the Actual BPS policy is nearly identical to that of Walk-Open.¹⁶ Table 2 also shows that the outcome of Walk-Open with 50% Walk is very close to that of 0% Walk, whereas the outcome of Open-Walk with 50% Walk is substantially different from that of 100% Walk. For grade K1, the range between the two extremal precedence policies is 8.3 percent, or 691 students. This range corresponds to roughly three-quarters of the range between the two extremal reserve policies. For grade K2, the precedence range is also three-quarters of the reserve range, while for grade 6 it is about two-thirds. These empirical observations show that the maximal ef-

¹⁵ Appendix E provides details on the sample.

¹⁶ Appendix E elaborates on the differences between Actual BPS and Open-Walk.

TABLE 2
NUMBER OF STUDENTS ASSIGNED TO WALK ZONE SCHOOLS, USING ONE LOTTERY NUMBER

		PRIORITIES = 50% WALK: CHANGING PRECEDENCE			PRIORITIES =
	PRIORITIES = 0% WALK (1)	Walk-Open (2)	Actual BPS (3)	Open-Walk (4)	100% WALK (5)
A. Grade K1					
Walk zone	3,849	3,879	3,930	4,570	4,787
	46.2%	46.6%	47.2%	54.8%	57.4%
Outside walk zone	2,430	2,399	2,353	1,695	1,468
	29.2%	28.8%	28.2%	20.3%	17.6%
Unassigned	2,054	2,055	2,050	2,068	2,078
	24.6%	24.7%	24.6%	24.8%	24.9%
B. Grade K2					
Walk zone	3,651	3,685	3,753	4,214	4,374
	47.2%	47.6%	48.5%	54.5%	56.5%
Outside walk zone	2,799	2,764	2,694	2,214	2,036
	36.2%	35.7%	34.8%	28.6%	26.3%
Unassigned	1,289	1,290	1,292	1,311	1,329
	16.7%	16.7%	16.7%	16.9%	17.2%
C. Grade 6					
Walk zone	3,439	3,476	3,484	3,797	3,907
	39.1%	39.6%	39.7%	43.2%	44.5%
Outside walk zone	4,782	4,750	4,743	4,419	4,309
	54.4%	54.1%	54.0%	50.3%	49.0%
Unassigned	565	560	559	570	570
	6.4%	6.4%	6.4%	6.5%	6.5%

NOTE.—The table reports the fraction of applicants assigned to walk zone schools under several alternative assignment procedures, using data from 2009–12. The 0% Walk procedure implements the student-proposing deferred acceptance mechanism with no walk zone priority; 100% Walk implements the student-proposing deferred acceptance mechanism with all seats having walk zone priority. Columns 2–4 hold the 50-50 school seat split fixed. Walk-Open implements the precedence order in which all walk zone seats are ahead of all open seats. Actual BPS implements BPS’s exact system (see app. E). Open-Walk implements the precedence order in which all open seats are ahead of all walk zone seats.

fect of changes in precedence policy is nearly as large as that corresponding to maximal changes in reserve size.

What policy would implement BPS’s intended 50-50 compromise? To answer this question, it is worth returning to the example in figure 1. In panel B of figure 1, walk zone applicants who are rejected from the walk zone half and then apply for open seats have systematically worse tiebreaker lottery numbers and are outnumbered by non-walk zone students. This results in two biases: (1) the walk zone students who remain have the least favorable tiebreakers among walk zone applicants, leaving them unlikely to be assigned ahead of non-walk zone applicants; and (2) there are twice as many non-walk zone applicants as walk zone appli-

cants in the residual pool for open seats. We refer to the first phenomenon as *random number bias* and the second as *processing bias*.

To examine the random number bias, table 3 investigates the effects of using separate tiebreaker lotteries for the walk zone and open seats. Column 2 reports on the Walk-Open precedence order with two tiebreaker lottery numbers. Even with two lotteries, there is processing bias, as the pool of walk zone applicants is still depleted by the time the open seats are filled. Walk-Open with two lottery numbers assigns 48.4 percent of students to walk zone schools at grade K1 and is still close to the 46.6 percent assigned when Walk-Open is used with only one lottery number. That is,

TABLE 3
NUMBER OF STUDENTS ASSIGNED TO WALK ZONE SCHOOLS, USING TWO LOTTERY NUMBERS

		PRIORITIES = 50% WALK: CHANGING PRECEDENCE		
	PRIORITIES = 0% WALK (1)	Walk-Open: Two Lotteries (2)	Open-Walk: Two Lotteries (3)	PRIORITIES = 100% WALK (4)
A. Grade K1				
Walk zone	3,849 46.2%	4,034 48.4%	4,556 54.7%	4,787 57.4%
Outside walk zone	2,430 29.2%	2,217 26.6%	1,709 20.5%	1,468 17.6%
Unassigned	2,054 24.6%	2,082 25.0%	2,068 24.8%	2,078 24.9%
B. Grade K2				
Walk zone	3,651 47.2%	3,880 50.1%	4,210 54.4%	4,374 56.5%
Outside walk zone	2,799 36.2%	2,539 32.8%	2,220 28.7%	2,036 26.3%
Unassigned	1,289 16.7%	1,320 17.1%	1,309 16.9%	1,329 17.2%
C. Grade 6				
Walk zone	3,439 39.1%	3,516 40.0%	3,784 43.1%	3,907 44.5%
Outside walk zone	4,782 54.4%	4,655 53.0%	4,415 50.3%	4,309 49.0%
Unassigned	565 6.4%	615 7.0%	587 6.7%	570 6.5%

NOTE.—The table reports the fraction of applicants assigned to walk zone schools under several alternative assignment procedures, using data from 2009–12. The 0% Walk procedure implements the student-proposing deferred acceptance mechanism with no seats having walk zone priority; 100% Walk implements the student-proposing deferred acceptance mechanism with all seats having walk zone priority. Columns 2 and 3 hold the 50-50 school seat split fixed. Walk-Open implements the precedence order in which all walk zone seats are ahead of all open seats, but different lottery numbers are used for walk zone and open seats. Open-Walk implements the precedence order in which all open seats are ahead of all walk zone seats, but different lottery numbers are used for walk zone and open seats. The same lottery numbers are used for each simulation.

Walk-Open with two lottery numbers is much closer to 0% Walk than to 100% Walk; this suggests that random number bias is only part of the reason Boston's assignment outcome is not midway between the 0% Walk and 100% Walk extremes.

Although it eliminates the random number bias, the remedy of using two lottery numbers has an important drawback. It is well known that using multiple lottery numbers across schools with the deferred acceptance algorithm may generate efficiency losses (Abdulkadiroğlu and Sönmez 2003; Abdulkadiroğlu et al. 2009; Ashlagi, Nikzad, and Romm 2015). Even though the two lottery numbers are within schools (and not across schools), the same efficiency consequence arises here. The unassigned row in the tables provides indirect evidence for this fact: Comparing table 2 to table 3, for each precedence policy under the 50-50 split, there are at least as many unassigned students and sometimes more with two tiebreaker lotteries.

Open-Walk eliminates both types of bias because neither the lottery number distribution nor the applicant pool is affected by application processing at the open half. (In the example illustrated in fig. 1, Open-Walk would result in 75 students from the walk zone being assigned.) Thus, distributional objectives may need to be accommodated by adjusting the reserve size.

To return to the Boston policy discussion, we conclude our investigation by examining how far the actual Boston system was from a 50-50 compromise. Table 4 computes the reserve size adjustment, under Open-Walk, that corresponds to BPS's implementation of the 50-50 reserve. Depending on the grade, BPS's implementation corresponds to Open-Walk with roughly a 5–10 percent walk zone reserve share. For grade K1, the Actual BPS implementation gives 47.2 percent of students walk zone assignments; this is just above the Open-Walk treatment with a 5 percent walk zone reserve (46.9 percent) but below the Open-Walk treatment with a 10 percent walk zone reserve (47.6 percent). For grade K2, the Actual BPS implementation has 48.5 percent walk zone assignment, a figure close to the Open-Walk treatment with a 10 percent walk zone reserve. For grade 6, the Actual BPS implementation is bracketed by Open-Walk with 5 percent and 10 percent walk zone reserves. An unbiased version of the BPS implementation reveals it to be a substantial distance from a 50-50 compromise; indeed, the BPS implementation is closer to a 10-90 compromise.

VI. Conclusion

Admissions policies in which applicants can be granted more than one type of seat raise questions about how seats should be processed. We have shown how both reserves and precedence are policy tools that have qual-

TABLE 4
WHAT POLICY WAS BEING IMPLEMENTED IN BOSTON?

	Priorities = 0% Walk (1)	Priorities = 5% Walk: Open-Walk: One Lottery (2)	Actual BPS (3)	Priorities = 10% Walk: Open-Walk: One Lottery (4)
A. Grade K1				
Walk zone	3,849 46.2%	3,906 46.9%	3,930 47.2%	3,965 47.6%
Outside walk zone	2,430 29.2%	2,369 28.4%	2,353 28.2%	2,304 27.6%
Unassigned	2,054 24.6%	2,058 24.7%	2,050 24.6%	2,064 24.8%
B. Grade K2				
Walk zone	3,651 47.2%	3,692 47.7%	3,753 48.5%	3,743 48.4%
Outside walk zone	2,799 36.2%	2,757 35.6%	2,694 34.8%	2,702 34.9%
Unassigned	1,289 16.7%	1,290 16.7%	1,292 16.7%	1,294 16.7%
C. Grade 6				
Walk zone	3,439 39.1%	3,461 39.4%	3,484 39.7%	3,496 39.8%
Outside walk zone	4,782 54.4%	4,751 54.1%	4,743 54.0%	4,715 53.7%
Unassigned	565 6.4%	574 6.5%	559 6.4%	575 6.5%

NOTE.—The table reports the fraction of applicants assigned to walk zone schools under several alternative assignment procedures, using data from 2009–12. The 0% Walk procedure implements the student-proposing deferred acceptance mechanism with no walk zone priority. Open-Walk implements the precedence order in which all open seats are ahead of all walk zone seats. The same lottery numbers are used for each simulation.

itatively similar impacts on school admission outcomes. We have also examined how those results generalize to centralized assignment systems.

Our analysis resolved a puzzle underlying a policy debate in Boston. Many groups in Boston felt that the BPS school assignment system did not sufficiently value children attending schools close to their homes despite the stated policy reserving half of each school’s seats for walk zone applicants. The resolution of this puzzle hinges on the important and surprising role played by Boston’s chosen precedence order.

In addition to our comparative static results, our empirical analysis shows how the chosen precedence order effectively undermined the policy goal of the 50-50 seat split in Boston. Moreover, our empirical results establish that, in Boston, the precedence order (1) is an important lever for achieving distributional objectives and (2) has quantitative impacts of magnitudes similar to those of changes in reserve sizes.

The role of precedence order on admissions was not understood at the time of Boston's 50-50 compromise, and it is clear that Boston did not intend to choose a precedence order that undermined the walk zone reserve (External Advisory Committee 2013). When our work first made clear the unintended consequences of Boston's precedence choice, our findings were immediately of interest to all sides of the 2012–13 Boston school choice debate. Neighborhood schooling advocates were upset to learn that the precedence order had rendered the walk zone reserve ineffective. School choice proponents, by contrast, pushed to either maintain the Walk-Open precedence order or eliminate the walk zone reserve entirely. (For details on policy discussions and the impact of our research, see app. D.) Central to our own view was the need to encourage transparency: it is not sufficient to express the reserve policy without also specifying the precedence order.

Pursuant to our work, Boston Superintendent Carol Johnson (2013) proposed eliminating walk zone priority entirely, as it had not been working as intended, and because the new choice menu system (Shi 2013) baked in a form of geographic preference under which students could apply only to schools relatively close to their homes. The new BPS admissions policy took effect for placing elementary and middle school students in the 2013–14 school year.

References

- Abdulkadiroğlu, A., P. A. Pathak, and A. E. Roth. 2009. "Strategy-Proofness versus Efficiency in Matching with Indifferences: Redesigning the NYC High School Match." *A.E.R.* 99 (5): 1954–78.
- Abdulkadiroğlu, A., P. A. Pathak, A. E. Roth, and T. Sönmez. 2005. "The Boston Public School Match." *A.E.R. Papers and Proc.* 95 (2): 368–71.
- . 2006. "Changing the Boston School Choice Mechanism." Working Paper no. 11965, NBER, Cambridge, MA.
- Abdulkadiroğlu, A., and T. Sönmez. 2003. "School Choice: A Mechanism Design Approach." *A.E.R.* 93 (3): 729–47.
- Agarwal, N., A. Abdulkadiroğlu, and P. A. Pathak. 2017. "The Welfare Effects of Coordinated School Assignment: Evidence from the NYC High School Match." *A.E.R.* 107 (12): 3635–89.
- Ashlagi, I., A. Nikzad, and A. Romm. 2015. "Assigning More Students to Their Top Choices: A Tiebreaking Rule Comparison." Working paper, Sloan School Management, Massachusetts Inst. Tech.
- Aygün, O., and I. Bo. 2013. "College Admission with Multidimensional Reserves: The Brazilian Affirmative Action Case." Working paper, Boston Coll.
- Aygün, O., and B. Turhan. 2016. "Dynamic Reserves in Matching Markets: Theory and Applications." Working paper, Bogazici Univ.
- Bagde, S., D. Epple, and L. Taylor. 2016. "Does Affirmative Action Work? Caste, Gender, College Quality, and Academic Success in India." *A.E.R.* 106 (6): 1495–1521.
- Balinski, M., and T. Sönmez. 1999. "A Tale of Two Mechanisms: Student Placement." *J. Econ. Theory* 84 (1): 73–94.

- Budish, E., Y.-K. Che, F. Kojima, and P. Milgrom. 2013. "Designing Random Allocation Mechanisms: Theory and Applications." *A.E.R.* 103 (2): 585–623.
- Daley, B. 1999. "Plan Drops Race Role in Enrollment, Compromise Misses Point, Critics Say." *Boston Globe*, October 20.
- Dur, U., P. A. Pathak, and T. Sönmez. 2016. "Explicit vs. Statistical Preferential Treatment in Affirmative Action: Theory and Evidence from Chicago's Exam Schools." Working Paper no. 22109, NBER, Cambridge, MA.
- Echenique, F., and M. B. Yenmez. 2015. "How to Control Controlled School Choice." *A.E.R.* 105 (8): 2679–94.
- Erdil, A., and H. Ergin. 2008. "What's the Matter with Tie-Breaking? Improving Efficiency in School Choice." *A.E.R.* 98 (3): 669–89.
- Erdil, A., and T. Kumano. 2012. "Prioritizing Diversity in School Choice." Working paper, Washington Univ.
- External Advisory Committee. 2013. "Minutes from the Meeting of the Committee." January 14. <http://bostonschoolchoice.files.wordpress.com/2013/01/eac-minutes-1-14-2013.doc>.
- Gale, D., and L. S. Shapley. 1962. "College Admissions and the Stability of Marriage." *American Math. Monthly* 69 (1): 9–15.
- Goldstein, D. 2012. "Bostonians Committed to School Diversity Haven't Given Up on Busing." *Atlantic*, October 10.
- Hafalir, I., M. B. Yenmez, and M. A. Yildirim. 2013. "Effective Affirmative Action in School Choice." *Theoretical Econ.* 8 (2): 325–63.
- Handy, D. 2012. "Debate on Overhauling Boston Schools' Assignment System Continues." 90.9 WBUR, November 13.
- Johnson, C. R. 2013. "Speech to Boston School Committee." March 13. <http://bostonschoolchoice.files.wordpress.com/2013/03/3-13-13-superintendent-sc-memo-on-assignmentment.pdf>.
- Kojima, F. 2012. "School Choice: Impossibilities for Affirmative Action." *Games and Econ. Behavior* 75 (2): 685–93.
- Kominers, S. D., and T. Sönmez. 2013. "Designing for Diversity in Matching." Working paper, Boston Coll.
- . 2016. "Matching with Slot-Specific Priorities: Theory." *Theoretical Econ.* 11 (2): 683–710.
- Landsmark, T. 2009. "It's Time to End Busing in Boston." *Boston Globe*, op-ed, January 31.
- Landsmark, T., H. Dajer, L. Gonsalves, et al. 2004. "Report and Recommendations of the Boston Public Schools Student Assignment Task Force." <https://www.tayfunsonmez.net/wp-content/uploads/2013/10/BPS-TaskForceReport.pdf>.
- Menino, T. 2012. "State of the City 2012." January 17. <http://www.cityofboston.gov/>.
- Pathak, P. A., and J. Sethuraman. 2011. "Lotteries in Student Assignment: An Equivalence Result." *Theoretical Econ.* 6 (1): 1–18.
- Pathak, P. A., and T. Sönmez. 2008. "Leveling the Playing Field: Sincere and Sophisticated Players in the Boston Mechanism." *A.E.R.* 98 (4): 1636–52.
- Seelye, K. Q. 2012. "4 Decades after Clashes, Boston Again Debates School Busing." *New York Times*, October 4.
- Shi, P. 2013. "Closest Types: A Simple Non-Zone-Based Framework for School Choice." Working paper, Massachusetts Inst. Tech.
- Vaznis, J. 2009. "Boston Shelves 5-Zone Proposal." *Boston Globe*, Education Desk, August 27.