INTERNATIONAL ECONOMIC REVIEW
Vol. 64, No. 4, November 2023 DOI: 10.1111/iere.12641

CATCHMENT AREAS, STRATIFICATION, AND ACCESS TO BETTER SCHOOLS*

By Caterina Calsamiglia and Antonio Miralles

ICREA and IPEG, Spain; Università degli Studi di Messina, Universitat Autonoma de Barcelona, and Barcelona School of Economics, Spain

School Choice provides students with the opportunity to attend better schools than those in their neighborhood. This is crucial for students from disadvantaged areas where schools may be of lower quality. Our theoretical model and numerical simulations show that the widely used Deferred Acceptance (DA) algorithm has limitations in providing access to better schools (ABS). When schools have varying levels of quality and when there are priorities linked to neighborhood schools, the DA algorithm experiences significant limitations in providing ABS. Top Trading Cycles, when compared to DA, offers greater ABS, particularly for disadvantaged students.

1. INTRODUCTION

In the last 20 years, a significant number of countries in the OECD have increased school choice options. Historically, children were assigned to a school within their neighborhood, but school choice policies now allow parents/guardians to select a school regardless of where they live. The primary objective of these policies is to provide families, especially those from disadvantaged backgrounds, with *access to better schools* (ABS) than the ones available in their local neighborhood. As noted on the Friedman Foundation for Educational Choice's website, the idea behind school choice is to facilitate access to better education opportunities.

We define ABS as the expected percentage of individuals who are assigned a place in a school they prefer over their local neighborhood school.

This study demonstrates that in situations where there are clearly inadequate schools, if priority for neighborhood schools apply, and the market is sufficiently large, ABS may be negligible using the commonly used Deferred Acceptance (DA) algorithm. This means that almost all individuals may be assigned to their local neighborhood school, regardless of their preferences.

We compare the DA algorithm to the Top Trading Cycles (TTC) algorithm,[1] both introduced in the school choice literature by Abdulkadiroğlu and Sönmez (2003). Demand that ex-

*Manuscript received November 2020; revised February 2023.

We thank Atila Abdulkadiroğlu, Eduardo Azevedo, Estelle Cantillon, Guillaume Haeringer, Onur Kesten, Tayfun Sonmez, Utku Ünver, seminar participants at MIE meeting in Chicago, at the "Frontiers in Market Design" Conference in Monte Verità, Switzerland, and seminars in Duke, Carnegie Mellon, Madison-Wisconsin, Barcelona, Budapest, Helsinki, Istanbul, Bonn, Belfast, Montréal, Padova, Napoli, Brisbane and Sydney for their comments and suggestions. Both authors acknowledge financial support by the Severo Ochoa program (CEX2019-000915-S). Antonio Miralles thanks the Ortygia Business School and acknowledges financial support from the Spanish Government (ECO2017-83534P) and from the Catalan Government (SGR2017 0711). Caterina Calsamiglia acknowledges financial support from the ERC Starting Grant 638893. Antonio Miralles also thanks the funding from the Italian Program PRIN 2022 (2022HLPMKN.) Please address correspondence to: Antonio Miralles Asensio, Università degli Studi di Messina, Dipartimento di Economia, Piazza Pugliatti, 1 - 98122 Messina ME, Italy. E-mail: *amirallesasensio@gmail.com*

[Correction added on October 10, 2023 after first Online publication: Acknowledgment on page 1.]

[1] We also briefly analyze the Boston Mechanism (BM; Immediate Acceptance) in Section 7.

ceeds school capacity is rationed by ordering applicants according to priorities and random lotteries. We consider the case with coarse priorities defined by residence in the neighborhood of the school. We introduce stratification between schools: there is a *bad* school, a school that all families believe is the worst. We study the extent to which families can obtain better placement than their neighborhood school, the school they are given priority at by the authorities, and the default school when there is no school choice.

An important and noteworthy result in this article is that, when all high-rated schools are weakly overprioritized (i.e., they have weakly more students from the neighborhood than their capacity), and the market is large, all students will be assigned by DA to their neighborhood school regardless of their preferences. To explain why the DA algorithm fails in this situation, consider a scenario where all schools have equal capacity and an equal number of prioritized students. In this case, no student in the catchment area of a good school can end up at the bad school under DA (because access to the neighborhood school is guaranteed in every round). In other words, all students who live in the catchment area of the bad school are stuck there, regardless of their tie-breaking lottery number. Students with priority at different good schools may want to "exchange" their slots, but if they apply to their preferred school without priority, they will have to have a higher lottery number than any student from the bad neighborhood. In a large enough market, the student with the highest lottery number in the bad neighborhood will win and be given DA at the good school. But this assignment will never be final if there is still a student with priority at a good school who has not been assigned to a good school. This creates a continuous competition for good schools that blocks all school exchanges, even among students with priority at good schools.

Hence, the well-known potential problem of interrupters reducing efficiency of the DA (Kesten, 2010) has far-reaching consequences in most school choice environments. Stability in the DA mechanism leads to low ABS.[2] In sum, *a stable allocation that provides limited good placements for disadvantaged students results in low ABS in general*.

Conversely, if good schools are underprioritized, meaning their capacity exceeds the number of students in the corresponding neighborhood, ABS delivered by DA is increased. Nevertheless, we demonstrate that *TTC is a superior alternative to DA as it surpasses DA in terms of ABS*. In TTC, individuals with a preference for each other's schools can trade, thereby preventing the blocking of trade among individuals with priority at good schools. TTC guarantees Pareto-optimality of the final allocation as established by Gale and Shapley (1964), however, this does not mean that it Pareto-dominates DA. Our finding that TTC dominates DA in terms of ABS is not a foregone conclusion.

Interestingly, we find that *disadvantaged students* (i.e., those with priority at bad schools) *have a higher chance of accessing better schools through TTC* compared to DA. Not only is TTC more efficient in terms of ABS, it is also fairer.

This result may seem surprising at first. TTC operates by enabling trades, but no one wants to trade "priority rights" with disadvantaged students. As a result, disadvantaged students would have a weakly lower chance of getting into a good school compared to other students. However, TTC offers opportunities for disadvantaged students when assigning "leftover" seats during the algorithm. As long as trade cycles occur during the implementation of TTC, unassigned disadvantaged students keep better lottery numbers than unassigned advantaged students when distributing the remaining seats. By not blocking trades, disadvantaged students have better access eventually to the still unassigned seats at good schools.

It should be noted that our findings do not suggest that the TTC allocation is always a better option for all disadvantaged students compared to the DA allocation. Although TTC of-

[2] Refer to Roth (2008) for more information on the limitations imposed by stability. This conclusion is similar to the findings in Combe et al. (2022), who found that DA in teaching assignments to schools allows only a limited number of reassignments, assuming teachers cannot be forced to leave their current school. For DA to be individually rational, a guaranteed priority for the school the teacher is already assigned to before applying for a new school must be in place. However, this limits access to different schools, similar to ABS. In this sense, their result is a special case of ours, where the number of prioritized seats is equal to the number of available seats in schools.

fers better chances for disadvantaged students to attend a school better than their worst option, it may have a lower chance of assigning them to their preferred school.

Empirical evidence. The basis of our model is the differentiation of schools based on quality. In the United States, the notion of "failing schools" is widespread in both policy discussions and media and refers to schools that have shown poor performance for two consecutive years, making up 10% of all schools in the country.[3] The media highlights the difficulties faced by some families in leaving underperforming schools in their neighborhood, and the limitations of school choice in improving opportunities for disadvantaged families. However, there is limited evidence of these families having similar preferences for the worst schools. The main issue is that preferences are not easily observable and the methods used to gather them can be manipulated. As a result, evidence for this comes from estimates of individual preferences. A study by He (2017) estimated preferences over four colleges in Beijing using the BM and found that one school was ranked fourth by at least 58% of participants. The worst school in terms of academic performance was also found to be the least preferred school by students. Only 5% of students ranked it as their first choice.

A study by Calsamiglia et al. (2020) found that 44% of schools in Barcelona were filled in the first round and 40% were never filled. Similarly, Agarwal and Somaini (2016) showed that one preschool in Cambridge (USA), King Open Ola, was ranked only by five families whereas the next best school already had 51 applicants. Combe et al. (2022) analyzed teacher assignments to schools in France and found that some schools were unwanted by teachers. This issue was partially resolved by granting priority to move to better preferred schools to teachers previously assigned to these schools. This demonstrates that market stratification is not only present in student school choice, but also among teachers.

Our results are further exacerbated by overprioritization, which occurs when the number of applicants with priority for a school is equal to or greater than its capacity. This is common in markets such as the teacher market in France or the school choice market in cities where individuals have a guaranteed spot in their neighborhood school, as seen in the Charlotte-Mecklenburg Public School District (Hastings and Weinstein, 2008). In most school choice markets, there is a shift from a neighborhood-based assignment to a centralized assignment with priorities for the neighborhood school, which becomes the prioritized school. Administrations aim to ensure families have access to a nearby school, so weak overprioritization is likely, even in cases where it is not explicitly imposed through a default school assignment.

Our results indicate that under both DA and with neighborhood priorities for all seats, a significant portion of families will be assigned to their neighborhood school. This is also true with other mechanisms like the BM, which we discuss further in Section 7. Calsamiglia and Güell (2018) found that in Barcelona, priorities have a major impact on the list of schools parents apply for under BM. They discovered that a large number of parents apply for the neighborhood school, regardless of their preferences, due to a change in the definition of neighborhoods in Barcelona. On the other hand, Calsamiglia et al. (2020) also conduct counterfactual analysis to determine the allocation results if DA or TTC were implemented instead. *Around 40% of families in Barcelona prefer a school outside of their neighborhood, illustrating that ABS is a quantitatively relevant aim*. Table 19 in their paper reveals that for families whose favorite school is not in their neighborhood, both BM and DA assign them to their favorite school less frequently than TTC, with the proportions being 47.2%, 41.8%, and 58.9% for BM, DA, and TTC, respectively. Despite the fact that families can be truthful and express their desire to leave their neighborhood under DA, the mechanism most often assigns them to the neighborhood school. On the other hand, TTC clearly makes it easier for families to leave their neighborhood, more so than DA and BM. Thus, the excessive assignment to the

[3] The U.S. No Child Left Behind (NCLB) Public Choice Program mandates that local school districts permit students attending academically unacceptable schools (F-rated) to switch to higher performing schools in the same district, if space is available. This requirement is outlined in the Title I Public School Choice for low-performing schools, see: http://www.ncpie.org/nclbaction/publicchoice.html.

neighborhood school imposed by both DA and BM limits the ability of families' preferences to determine the allocation of students to schools.[4]

The evidence offered does not directly confirm our assumptions or conclusions, however it suggests that our analysis may offer insight into why some cities with stratified school systems and neighborhood priorities in school choice have limited school choice options.

Literature. The literature on school assignment mechanisms has focused on three main properties: strategy-proofness, stability, and efficiency. Strategy-proofness means that it is the best strategy for individuals to submit their true preferences. DA is favored for not only being strategy-proof but also for producing a stable allocation. Stability requires that no individual would prefer a different school and that the preferred school has no one with lower priority admitted. The results of this article apply to any stable mechanism, but DA's allocation is not always Pareto-efficient, except for certain priority structures (Ergin, 2002). Pareto-efficiency means that it is not possible to make one individual better off without making another worse off. The TTC mechanism is strategy-proof and efficient but not stable. No mechanism has all three properties (Kesten, 2010). However, the efficiency costs of DA, as demonstrated in experiments such as Chen and Sönmez (2006), are minimal, leading to its adoption in cities like New York and Boston, where it has replaced the BM.[5]

Both DA and BM, or a combination of the two (see Chen and Kesten, 2013) have been *the most debated alternatives by far*.[6] TTC has been used in New Orleans, and for some time in San Francisco, as far as the authors are aware. [7] This article adds an important reason to question the adequacy of DA given its limited ABS in very realistic education markets.

An important reference is the seminal paper (Kesten, 2010), which shows that for any vector of school capacities and any set of students, there are priority structures and individual preferences such that the Pareto-dominant stable allocation gives each student one of her two worst options.[8] Kesten addresses this issue in his paper by proposing the Efficiency-Adjusted Deferred Acceptance Mechanism (EADAM). This mechanism requires students to give up priority rights that have no effect on their own assignment, but may harm other students' assignments. Our article takes a similar approach, but with some key differences. We demonstrate that if there is a commonly ranked low school (bad school) and there are neighborhood priorities, individuals will be assigned to their neighborhood school, regardless of their preferences for other schools. In addition, we highlight that the problem pointed out by Kesten has a significant impact in the realistic education market depicted in our model.

Our article explores priority structures in school choice that lie between two extremes in the literature: the strict priority model (e.g., Ergin and Sönmez, 2006; Pathak and Sönmez, 2008) and the no-priorities model (Miralles, 2008; Abdulkadiroğlu et al., 2011). Although these papers discuss their models beyond the adopted extreme assumptions, their most notable proofs are based on these assumptions. Ergin and Erdil (2008) is an exception in that it analyzes weak priority structures, but focuses on improving the assignment after ties have been resolved in a specified manner. In contrast, our article accounts for the randomness of tiebreakers. Troyan (2012) is another exception, in that it compares the ex ante efficiency of BM and DA when coarse priorities are included.

[4] Recently, Terrier et al. (2021) found that banning a more general version of the BM in the United Kingdom, which led to an expansion of the DA, did not benefit families of disadvantaged backgrounds. It is worth noting that the mechanism was changed but most school districts preserved neighborhood priorities.

[5] Experiments evaluating the efficiency cost have been done in the lab, and the simulated environments used did not contain *bad schools*, as we model them here or are found in the data. This article suggests that under the presence of bad schools efficiency losses may be very large, since preferences may have a rather small effect on the final allocation.

[6] Important contributions to this debate include Abdulkadiroğlu and Sönmez (2003), Abdulkadiroğlu et al. (2006), Ergin and Sönmez (2006), Miralles (2008), Pathak and Sönmez (2008), and Abdulkadiroğlu et al. (2011).

[7] Actually, the mechanism used in San Francisco was a manipulable variant of TTC. It was substituted by DA in 2019.

[8] Following Kesten, there is recent literature advocating for a relaxation of the stability concept (Ehlers and Morrill, 2019; Troyan et al., 2020)

This article's findings are relevant for the field of economics of education, which studies the impact of school choice on school outcomes, as seen in studies like Lavy (2010) and Hastings et al. (2010). These studies assume that choice affects the allocation of students to schools based on their preferences, but they overlook the impact of priority structures on the allocation mechanism. The results of this article may shed light on potential misattributions of effects in the empirical literature.

In addition, the article's results support recent insights by Abdulkadiroğlu et al. (2020), who found that among strategy-proof and Pareto-optimal mechanisms, TTC minimizes justified envy, making DA's main advantage over unstable mechanisms redundant.

Outline of the article. The article clarifies the meaning of ABS in Section 2 and explores its alternatives. In Section 3, the mechanisms being compared are introduced. Section 4 provides two examples to show how the mechanisms lead to different ABS outcomes due to main driving forces. The results in the article are presented in a simplified model for easy understanding in Section 5. This model is inspired by Aumann (1964) and assumes the allocation of a continuum of individuals to a limited number of schools. Example 1 and further simulations indicate that the results from the continuum model also apply to realistic scenarios with limited school capacities. Section 6 extends the findings and offers observations for models with additional schools, supported by numerical simulations. Section 7 evaluates the BM and the stratification assumption. The article concludes in Section 8. Appendices A.1 and A.2 contain a discussion of the discrete model example and all proofs, respectively. An online appendix[9] contains simulation codes and detailed results.

2. ACCESS TO BETTER SCHOOLS

ABS refers to the percentage of individuals who are assigned to a school that is preferred over their neighborhood school. The standard for measuring ABS is neighborhood assignment, where before a centralized school choice program was introduced, children were assigned to their local school when possible. ABS shows the proportion of students who have benefited from the school choice program.

ABS is particularly relevant for individuals living in areas with so-called failing schools, as it provides access to quality education. Low ABS is a measure of the influence of neighborhood priorities. When ABS is close to zero, students' choices are restricted to their neighborhood schools. As ABS increases, the impact of neighborhood priorities on students' assignments decreases. Some families may attempt to move to a neighborhood with their preferred school, but this option is only feasible for wealthy families. An assignment process that results in high ABS levels equalizes educational opportunities.[10]

A valid criticism of the concept of ABS is that it disregards the possibility of choice making some individuals worse off, that is, assigning them to a school that is less preferred than their neighborhood school. To account for this, Access to Worse Schools (AWS) could be used to calculate a "net" ABS, $NABS = ABS - AWS$.

In a stratified model of school choice, as discussed in later sections, there may be a bad neighborhood with a school that everyone dislikes. To minimize AWS, one could argue that students from such a neighborhood should have no chance of attending other schools and should be confined to that poor school. However, if a student from this neighborhood gains access to a better school, a student from a better neighborhood and school must be assigned to the poor school, resulting in a net count of zero. This means that Net Access to Better Schools (NABS) does not reward mechanisms that allow children to escape from disadvantaged areas, due to a crowding-out effect. Despite this, ABS is deemed crucial for families from disadvantaged neighborhoods, as per the NCLB initiative.

[9] Available at https://sites.google.com/view/antoniomiralles.

[10] We thank a referee for the ideas in this paragraph.

The ideas of the NCLB initiative are reflected in some school districts' policies, where students from areas with "bad schools" are given priority over students from areas with "good schools" in the allocation of school places. For example, the San Francisco Unified School District used to prioritize families living in areas with "bad schools" (the lowest 20 percentile of average test scores) in all schools. [11]

To better take all these arguments into account, one could suggest an aggregated welfare indicator of the type

$$W = ABS - \delta AWS + \gamma ABS_w,$$

where δ is a penalty factor for each student who obtains an allocation worse than her priority-giving school, and γ is a redistribution premium for each student with priority at a worst school who obtains a better allocation. Therefore, for each student who obtains a worse placement than her priority-giving school, we require a compensation of δ students from good catchment areas obtaining a better placement, or $\frac{\delta}{1+\gamma}$ if the students improving their positions come from disadvantaged areas. *NABS* would be a special case of this formula with $\delta = 1$ and $\gamma = 0$.

Our simulations in Section 6 yield that, even if we suppress the redistribution premium ($\gamma = 0$) and we introduce a penalty factor $\delta = 2$, TTC W-dominates DA in all of the environments we consider. A factor $\delta = 4$ still allows TTC to W-dominate DA in 90% of the scenarios we consider in which TTC and DA do not yield identical assignments.

There are various ways to compare the satisfaction of families under different school choice mechanisms. One approach is to use Pareto-domination, which guides the comparison of outcomes. Another criterion is rank domination (Featherstone, 2014), which considers one mechanism's outcome as rank-dominating another if a higher percentage of students is allocated to better ranked schools at every position. Pareto-dominance implies rank dominance, and rank dominance implies a higher ABS. However, the reverse implications are not necessarily true. ABS is a measure that provides a way to compare outcomes, whereas alternative domination criteria may not always provide clear rankings.

3. THE MECHANISMS

The mechanisms we compare in this article are DA and TTC. In both mechanisms, students submit a ranked list of schools they prefer. Their strategy space is the set of all possible school rankings. School catchment area is the primary priority criteria when demand exceeds supply, and in case of ties, a unique lottery number per student is used to break the tie. The outcome of the lottery is unknown at the time students submit their lists.

Deferred Acceptance:

- In every round, each student applies for the highest school in her submitted list that has not rejected her yet.
- For every round$k$, $k \geq 1$: Each school tentatively assigns seats to the students that apply to it or that were preaccepted in the previous round following its priority order (breaking ties through a fair lottery).[12] When the school capacity is attained the school rejects any remaining students that apply to it in that round.
- The DA mechanism terminates when no student is rejected. The tentative matching becomes final.[13]

[11] http://www.sfusd.edu/en/assets/sfusd-staff/enroll/files/2012-13/annual_report_march_5_2012_FINAL.pdf, page 81.

[12] We assume that there is a single tiebreaker that serves to break ties when necessary at all schools. In the absence of priorities, a single tiebreaker guarantees ex post efficiency, whereas a separate tiebreaker per school cannot guarantee such a property (Abdulkadiroğlu et al., 2015).

[13] Abdulkadiroğlu et al. (2015) show that this algorithm converges to an assignment in big continuum economies, even though not necessarily in finite time.

TABLE 1
EXPECTED PERCENTAGE OF STUDENTS WHO GET ACCESS TO BETTER SCHOOL (ABS)

Mechanism↓ \ n→	1	2	5	10	20	∞
DA	22	13.3	6.3	3.4	1.77	0
TTC	66.7	66.7	66.7	66.7	66.7	66.7

Top-Trading Cycles:

- In each round, a cycle is initiated by selecting a school s that has available seats. The first student in its priority list, i, is selected, and the school that this student prefers the most, s', is then selected. This process continues until a cycle is formed.
- Students in the cycle are assigned to the schools they point at. We remove these students and slots.
- The process is repeated until all students have been assigned, with schools that are filled to capacity being removed from both students' lists and schools' priority lists.[14]

4. TWO EXAMPLES AND TWO INTUITIONS

4.1. *No ABS in Moderately Large Markets with Bad Schools and Neighborhood Priorities.* To illustrate our findings, we present an example with a finite set of individuals. Our model features three neighborhoods, each with n families and a school with a capacity of n. Neighborhood $i \in \{i_1, i_2, i_3\}$ gives priority to school $s \in \{1, 2, 3\}$, and in case of ties, a unique fair lottery is used.

In this example, all students rank school 3 as the worst, with student i_1 favoring school 2, and the rest preferring school 1.

Under the DA mechanism, students with priority at a desirable school are assured a place at that school. This means that all students in the i_3 category will eventually be assigned to the worst school. For instance, if $n = 1$, students i_1 and i_2 would like to swap their guaranteed slots. However, since they do not have priority at their preferred school, this exchange will only occur if both i_1 and i_2 get better lottery numbers than the i_3 student with the highest lottery number. The i_3 student blocks this trade with a probability of 2/3, which is a specific case of the well-known interrupter problem (Kesten, 2010).

Appendix A.1 shows how, for $n > 1$, the probability of blocking the xth exchange rapidly increases with $x = 1, .., n$, since not doing so requires both xth best lottery numbers in i_1 and i_2 to beat the best lottery number in i_3.

The results from determining the expected percentage of students who get into a better school than the one in their neighborhood are presented in Table 1. As n increases, this percentage decreases quickly to zero. Even when n is small, the percentage is very low, at 1.77% when there are 20 students per school.[15]

[14] TTC converges in the continuum (Leshno and Lo, 2021). An idea is to discretize both the mass of applicants and school capacity and to show that the discrete version converges as the size of the units goes to 0. To do this on the demand side, define a given type t by individuals with particular preferences and priorities (before ties are broken). Since the set of priorities and schools is finite, so is the set of orderings and types. Next, divide each type into *units* of size $1/n$. Let n be a natural number such that each type and each school capacity is larger than $1/n$, so that each type and school is composed of at least one unit u_n . However, each type and school capacity may not be divisible by an integer number of units. Note that the total mass of leftovers on the demand side is divisible by an integer number of units, since the total mass is of unit 1. Similarly for the supply side. Now define the preference ordering for the leftover units on demand side as a random preference ordering of the leftover types in that unit. Similarly distort capacities so that the remaining seats are all of one of the schools. We can now run TTC on units of individuals and schools. The assignment is distorted by mass of the leftovers. But one can show that the mass of leftovers on both sides goes to 0 as n goes to infinity.

[15] Table A6 in the Online Appendix contains similar results with four good schools and one worst school.

TABLE 2
MASS OF STUDENTS WHO GET ACCESS TO BETTER SCHOOL (ABS)

Assigned to↓ Priority at→	1	2	b	w
1	6,2	0,**4**	0,0	0,0
2	$\frac{4}{3}$,**4**	4,0	$\frac{4}{3}$,2	$\frac{4}{3}$,2
b	0,0	0,0	6,6	0,0
w	$\frac{2}{3}$,2	0,0	$\frac{2}{3}$,0	6+$\frac{2}{3}$,6

Under TTC, the blocking of an "exchange" between students of type i_1 and i_2 does not occur. Schools 1 and 2 form a cycle, with school 1 pointing at a student of type i_1 and school 2 pointing at a student of type i_2, who then points back to school 1. Regardless of the lottery numbers for students of type i_3, students of types i_1 and i_2 are assigned to schools better than their priority-giving schools.

Table 1 summarizes the expected number of students who access to a better preferred school (weighted by the size $3n$ of the market).

The expected percentage of students who secure a better school assignment than their neighborhood school decreases quickly to zero in the DA mechanism. Even with only 10 students per school, only 3.4% are expected to receive a better placement outside their catchment areas. With 20 students per school, this percentage drops to just 1.77%, showing that the poor results of DA seen later in the model are not simply due to the use of a continuum model in Section 5.

4.2. *More Access to "Leftovers".* Under TTC, a student's chances of getting assigned to their top-preferred school depend on how popular the school giving them priority is to others. This means that a disadvantaged student has weaker chances of getting assigned to a good school compared to others. DA is seen as more equal as it does not rely on trading ideas and gives disadvantaged students the same chances at good schools as any nonprioritized student. Counterintuitively, it is actually the opposite–disadvantaged students are more likely to be assigned to better schools under TTC than under DA.

Let us observe the following example: We have four schools, two good (1 and 2), one "bad" (b) and one "worst" (w). Consider the set of good schools to be $G = \{1, 2, b\}$ in the sense that everyone wants to avoid the worst school. However, inside G, nobody likes the bad school. Indeed, everyone ranks schools b and w third and fourth, respectively.

Capacities and prioritized students are $q_1 = q_b = 6, q_2 = q_w = 8, n_1 = n_b = n_w = 8, n_2 = 4$. These are treated as masses instead of nondivisible units. Students with priority at school 2 prefer school 1. Any other student prefers school 2.

Table 2 summarizes the mass of allocated students to each school, according to their priority group. In each cell, the first number corresponds to DA; the second number, to TTC.

In DA, it is observed that nonprioritized students will not have access to school 1. This is because all students would like to attend school 2, except for those with priority there. If all of these prioritized students take up four slots at school 2, the remaining four slots must be divided among the nonprioritized students, all of whom have equal chances at school 2. A mass 4/3 of the prioritized students at school 1 occupy slots at school 2. Of the remaining prioritized students at school 1, six are assigned to all the slots of their second-choice school 1. This means that school 1 will not be accessible to nonprioritized students. The rest of the assignment is straightforward: students with priority at school b fill all its slots, whereas the remaining students are placed at the worst school.

In contrast, TTC begins with a trade of slots between prioritized students from school 1 and those from school 2. Half of the students with priority at school 1 receive slots at school 2, and the minimum lottery number among unassigned students from this group becomes 1/2. The next step is to assign the four remaining slots at school 2. These are divided equally among students with priority at schools b and w, who enjoy an advantage over those prioritized by school 1. The two remaining slots at school 1 are assigned to prioritized students, followed by the six slots at school b and the remaining slots at school w.

This example highlights the differences between DA and TTC:

1. Under TTC, students with priority at underdemanded schools, including the worst school, benefit from the trade among other students. When it is time to assign leftover slots from underprioritized good schools, they have an advantage in the tiebreaker lottery.
2. Under TTC, students with priority at highly demanded schools have better chances of improving their assignments through trade, but also face a higher risk of being assigned to the worst school.

This example also shows who would prefer DA to TTC: if a student's top choice school is also the priority-giving school and is highly sought after by other students, the trades carried out in TTC would bring no benefits for the student and increase the likelihood of being placed in a worse school.

5. MODEL

We present a simple model in order to illustrate our insights.[16] We have a mass N of students $i \in I = [0, N]$, each of them to be allocated to one of three schools. I is endowed with the uniform Lebesgue measure λ. Two of the schools are "good" and one is "bad," in the sense that all students rank it as worst. Good schools are labeled 1 and 2, respectively, whereas the bad school is labeled w (as for "worst"). Schools have strictly positive capacities q_1, q_2, and q_w that add up to N. Students $i \in I$ have preferences $\succ_i$ over the schools. No student is indifferent between any two schools.

There is a measurable catchment area function $\pi : I \to \{1, 2, w\}$. Each student has a unique catchment area where she has priority over students outside the catchment area. There is a mass n_1, n_2, and n_w of students for the catchment areas of schools 1, 2, and w, respectively. We denote with $\Pi_s = \{i \in I : \pi(i) = s\}$ the set of students prioritized by school s. Students living in the catchment area have priority over students living outside.

Student i's preferences over schools $\succ_i$ could be summarized by the identity of the most-preferred school, since w is ranked last by everyone. Therefore, $\Pi_{ss'}$ denotes the set of students with priority at school s whose favorite school is s', and $n_{ss'}$ denotes its associated mass.

Other ties are resolved when needed using a *fair*[17] lottery outcome $l : I \to [0, 1]$ that assigns one number to each student. We apply the convention that a lower lottery number beats a higher lottery number.

For each school s, define $\rho_s = q_s/n_s$. We say school s is *overprioritized* if $\rho_s < 1$ (capacity is smaller than the number of individuals with priority in the school), and underprioritized in the opposite case. Notice that we cannot have the three schools being either all overprioritized or all underprioritized, since we have assumed that total capacity is equal to total mass of students. For two schools s and s', we say that s is more prioritized than s' if $\rho_s < \rho_{s'}$.

A *matching*[18] is a function $\mu = I \to \{1, 2, w\}$. For each matching μ, we compute the mass of students who obtain a slot in a school preferred to that of their catchment areas, as a measure of students' real choice. We call this measure *Access to Better Schools*, denoted *ABS*. We also

[16] In previous versions of this article, we use a more general model with an arbitrary number of schools. Results are qualitatively similar to the ones we find here. For DA, there is an upper bound to ABS that collapses to zero when good schools are weakly overprioritized.

[17] "Fair" meaning that for every interval $[l', l''] \subset [0, 1]$ and every group $\Pi_{ss'}$ we have $\lambda(\{i \in \Pi_{ss'} : l(i) \in [l', l'']\}) = n_{ss'}(l'' - l')$.

[18] In this article, we use the terms matching, assignment, and allocation indistinctively.

compute ABS for priority groups Π_s. More formally:

$$ABS(\mu) = \lambda\big(\{i \in I : \mu(i) \succ_i \pi(i)\}\big)$$
$$ABS_s(\mu) = \lambda\big(\{i \in \Pi_s : \mu(i) \succ_i \pi(i)\}\big).$$

We compare two matchings: the one induced by truth telling in DA (the DA matching), and that induced by truth telling in TTC (the TTC matching).

We find the following result:

PROPOSITION 1. *For every school $s \in \{1, 2, w\}$, we have $ABS_s(TTC) \geq ABS_s(DA)$.*

Note that the statement is also true for the set Π_w of disadvantaged students. Thus, TTC is superior to DA in terms of allocative efficiency, as measured by ABS. Furthermore, it provides more opportunities for disadvantaged students to enhance their positions. The proof for this can be found in Appendix A.2, where a range of scenarios are analyzed individually. The challenge in extending this result to models with an increased number of good schools arises from the increased number of cases that need to be considered. The following section uses numerical simulations to demonstrate that our conclusions are not limited to the two-good-school model.

6. MORE GOOD SCHOOLS

6.1. *Numerical Simulations.* We compute numerical simulations[19] in which we consider four good schools $G = \{1, 2, 3, 4\}$ and one worst school w. Each school has 20 slots and there are 100 students. Students' valuations for schools have three components: (1) an extra for neighborhood school (a *neighborhood effect* caused by geographical proximity), (2) an independent value component u_i, and (3) a common value component c,

$$v_{is} = 1\{\pi(i) = s\} + \alpha u_{is} + \beta c_s, i \in I, s \in G.$$

All values v_{iw} are zeros. All c_s and u_{is}, $\forall i \in I, s \in G$, are independently drawn from the uniform distribution. The common vector c is then sorted so that $c_1 > c_2 \ldots$. Therefore, schools are numbered according to their popularity.

Simulations are programmed and computed with MATLAB R2022a, with default seed for random number generation. We consider a grid of scenarios varying according to:

1. The overprioritization of the two most popular schools compared to the other less popular good schools, combined with different levels of overall underprioritization of good schools (overprioritization of the worst school), resulting in 12 possibilities: [40/30/20 students prioritized by the worst school] × [large/small difference in number of prioritized students among good schools[20]] × [the two most popular schools are the least/ the most overprioritized].
2. Different values for α and β to change the importance of each component of v_i, resulting in nine possibilities: [balance between neighborhood effect and other sources, more weight to neighborhood effect, less weight to neighborhood effect] × [balance between common component and independent component, more weight to common component, less weight to independent component].

For each scenario, 50 valuation matrices and 50 single tiebreakers were computed for each of the 50 calculated valuation matrices, resulting in a total of 2,500 computed assignments per

[19] We thank Juan Sebastian Pereyra and Li Chen for the help with the simulations.
[20] We alternate values of 10, 15, 20, and 25 prioritized students.

	Percentiles	Smallest		
1%	.00065	.00065		
5%	.00101	.00092		
10%	.00139	.001	Obs	78
25%	.03049	.00101	Sum of Wgt.	78
50%	.05654		Mean	.0916532
		Largest	Std. Dev.	.0877828
75%	.15264	.27255		
90%	.21811	.30176	Variance	.0077058
95%	.27255	.32057	Skewness	1.000891
99%	.35134	.35134	Kurtosis	3.156208

FIGURE 1

$ABS(TTC) - ABS(DA)$ WHEN ALLOCATIONS ARE NOT IDENTICAL

scenario and mechanism (5,000 total assignments). Therefore, only 90 scenarios were considered,[21] for a total of 450,000 computed assignments.

Twelve of the scenarios resulted in DA and TTC both collapsing into Serial Dictatorship because the valuation generating formula gave too much weight to the neighborhood effect and common value component.[22] We will focus on the remaining cases, where differences are observed.

Results are shown along the next tables. A complete deploy of all calculations and the MATLAB code are shown in the Online Appendix. For each scenario, we compute ABS for all students and for each priority group $\Pi_1, \ldots, \Pi_4, \Pi_w$, both under DA and under TTC, measured as the fraction of the considered group who obtain a slot in a school better than the priority-giving school. We calculate the AWS for each student group, which represents the proportion of students that get assigned to a school worse than their preferred one. In addition, we determine the fraction of students in each priority group who prefer their allocation under DA over TTC and vice versa.

The overall observation is that ABS under TTC is *always* superior. ABS dominance tends to be minor when ordinal preferences are highly correlated among individuals and there is a high weight of the neighborhood effect (e.g., $\alpha = 0.5$, $\beta = 1.5$). In such a case, the allocations under TTC and under DA tend to coincide, and to collapse into a serial dictatorship allocation.

Figure 1 displays the comparison of ABS between TTC and DA in scenarios where the mechanisms produce different matchings. Figure 2 uses a simple linear regression to approximately link this difference to the parameters of each scenario, as presenting all the results would take up a lot of space (the complete results can be found in the Online Appendix).

The size of the ABS domination of TTC over DA becomes enormous in some cases. See, for instance, Table A10.c in the Appendix ($n_1 = n_2 = 10, n_3 = n_4 = 20$, particularly when $\alpha = 6$, $\beta = 2$.) We see differences of the order of 0.32 (from 12.6% under DA to 44.6% under TTC).

From the estimation, we see that increasing the level of overprioritization in either most popular or less popular good schools has a similar effect. By increasing by one the number

[21] Criteria (1) and (2) should yield 108 scenarios. In using criterion set (1), we considered two kinds of scenarios in which all good schools have the same number of prioritized students. In such cases, the criterion of giving the highest number of prioritized seats to popular/less popular good schools does not bite. This is the reason why we have 90 scenarios to consider.

[22] For a Serial Dictatorship, one needs to have a linear ordering of all students. We refer to the linear ordering in which students are first ordered according to the priority-giving school (being school 1 first, school 2 second, …) breaking ties thereafter with the lottery number. In those scenarios, all students from Π_1 preferred school 1 the most, then all students from Π_2 ranked school 2 first or immediately after school 1, etc.

Source	SS	df	MS
Model	.518584519	4	.12964613
Residual	.074763741	73	.001024161
Total	.593348261	77	.007705822

Number of obs	=	78
F(4, 73)	=	126.59
Prob > F	=	0.0000
R-squared	=	0.8740
Adj R-squared	=	0.8671
Root MSE	=	.032

diffabs	Coef.	Std. Err.	t	P>|t|	[95% Conf.	Interval]
n12	.0052519	.0010173	5.16	0.000	.0032244	.0072793
n34	.0049911	.0010173	4.91	0.000	.0029637	.0070185
alpha	.0485095	.0022583	21.48	0.000	.0440087	.0530103
beta	-.009126	.0020945	-4.36	0.000	-.0133003	-.0049516
_cons	-.1908442	.0337659	-5.65	0.000	-.2581395	-.1235488

FIGURE 2

REGRESSION OF $ABS(TTC) - ABS(DA)$ OVER THE PARAMETERS OF EACH SCENARIO, WHEN THE MECHANISMS DELIVER UNIDENTICAL ALLOCATIONS

NOTES: N12 STANDS FOR n_1 AND n_2, WHEREAS N34 STANDS FOR n_3 AND n_4.

of prioritized students, the difference between $ABS(TTC)$ and $ABS(DA)$ slightly increases 0.5%. This is in line with the idea that overprioritization is bad news for DA concerning ABS. The effect of α, the weight of i.i.d. preferences, is clear: an increase in one unit enlarges the difference regarding ABS by 4.85%. The coefficient associated to β, the weight of common values, has the expected negative sign. An increase in one unit shrinks the difference by 0.9%. Decreasing the weight of the neighborhood effect by increasing $\alpha + \beta$ while keeping α/β constant has the expected positive effect as long as $\alpha/\beta > 0.9/4.85 \approx 0.19$. As an example, when $\alpha = \beta$, increasing $\alpha + \beta$ by one unit enlarges the difference in ABS by approximately 2%.

We wish to stress that α is the parameter that represents taste variety. It is in the context of a high α that School Choice programs have a clearer purpose. And it is precisely in these scenarios that TTC dominates more strongly.

We notice that, as expected, ABS under DA collapses when all good schools are weakly overprioritized, given the highly likely appearance of an interrupter, since disadvantaged students obtain no chances at good schools (Table A6 in Online Appendix).

As for ABS for disadvantaged students (those in Π_w), it is clear that TTC is *superior in all cases*. The difference with DA tends to zero when either one is true: (1) ordinal preferences are perfectly correlated among individuals or there is a sufficiently high weight of the neighborhood effect, (2) all good schools are weakly overprioritized (no access to good schools for disadvantaged students), or (3) all good schools are weakly underprioritized (the amount of disadvantaged students who obtain placement at a good school is mechanism-invariant).

The domination of TTC over DA regarding ABS for disadvantaged students is sizable in the same cases in which the difference in overall ABS is high. In the example we were considering for the general case, disadvantaged students obtain ABS equal to 27.5% under DA and 49.7% under TTC.

Figure 3 presents a summary description of $ABS_w(TTC) - ABS_w(DA)$, exception made for the following cases: (1) when both TTC and DA deliver identical allocations coinciding to that of Serial Dictatorship, (2) when both good schools are weakly underprioritized. In the latter case, all students prioritized by a good school will certainly obtain a slot at a good school, under both mechanisms. The number of slots available for disadvantaged students is identical between mechanisms.

Figure 4 presents the results of a simple linear regression between $ABS_w(TTC) - ABS_w(DA)$ and the set of parameters characterizing each scenario, with a subsample that skips cases (1) and (2) above.

	Percentiles	Smallest		
1%	0	0		
5%	0	0		
10%	0	0	Obs	32
25%	.01184	0	Sum of Wgt.	32
50%	.04134		Mean	.0712959
		Largest	Std. Dev.	.0702518
75%	.12421	.16804		
90%	.16804	.18768	Variance	.0049353
95%	.22238	.22238	Skewness	.8130252
99%	.22258	.22258	Kurtosis	2.368389

FIGURE 3

SUMMARY DESCRIPTION OF $ABS_w(TTC) - ABS_w(DA)$ FOR CASES IN WHICH THIS VALUE IS NOT TRIVIALLY 0

Source	SS	df	MS			
				Number of obs	=	32
				F(4, 27)	=	47.71
Model	.134031542	4	.033507885	Prob > F	=	0.0000
Residual	.018963095	27	.000702337	R-squared	=	0.8761
				Adj R-squared	=	0.8577
Total	.152994636	31	.004935311	Root MSE	=	.0265

wdiffabs	Coef.	Std. Err.	t	P>|t|	[95% Conf.	Interval]
n12	.0001087	.0019096	0.06	0.955	-.0038096	.004027
n34	-.0000704	.0019096	-0.04	0.971	-.0039887	.0038478
alpha	.0388554	.0029206	13.30	0.000	.0328629	.0448479
beta	-.0130176	.0027296	-4.77	0.000	-.0186184	-.0074169
_cons	.0027422	.0710812	0.04	0.970	-.1431043	.1485886

FIGURE 4

REGRESSION OF $ABS_w(TTC) - ABS_w(DA)$ ON THE PARAMETERS CHARACTERIZING THE DIFFERENT SCENARIOS, FOR CASES IN WHICH THE DEPENDENT VARIABLE IS NOT TRIVIALLY 0

$ABS_w(TTC) - ABS_w(DA)$ behaves similarly to $ABS(TTC) - ABS(DA)$, although with lower intensities. The 95 percentile on ABS_w gain is around 22.2% (27.2% for general ABS). An increase in the number of prioritized students seems not to have a significant effect. An increase of one unit in α rises the difference in ABS_w by 3.88%, whereas an equivalent increase in β reduces this difference by 1.3%. Decreasing the weight of the neighborhood effect by increasing $\alpha + \beta$ while keeping α/β constant has a positive effect on $ABS_w(TTC) - ABS_w(DA)$ as long as $\alpha/\beta > 1.3/3.88 \approx 0.335$. As an example, when $\alpha = \beta$, increasing $\alpha + \beta$ by one unit enlarges the difference in ABS by approximately 1.29%.

TTC obtains worse results than DA regarding AWS in general, exception being the cases where the parameters forced one of these outcomes: (1) equivalent allocations in both mechanisms, (2) weakly overprioritized good schools (no prioritized student could bear a risk of obtaining a worse placement in either of the mechanisms).

It is however noticeable that, *in all cases considered*, the unfavorable difference for TTC in AWS is lower than the favorable difference in ABS. Figure 5 summarizes the ratio $\frac{AWS(TTC)-AWS(DA)}{ABS(TTC)-ABS(DA)}$ for cases in which the mechanisms do not deliver identical matchings.

In more than half of the cases considered, DA and TTC deliver identical AWS. This is not that surprising because of the scenarios in which all good schools are weakly overprioritized.

	Percentiles	Smallest		
1%	0	0		
5%	0	0		
10%	0	0	Obs	78
25%	0	0	Sum of Wgt.	78
50%	0		Mean	.0678943
		Largest	Std. Dev.	.104043
75%	.1391397	.269035		
90%	.2398644	.2837187	Variance	.0108249
95%	.269035	.3730301	Skewness	1.38615
99%	.4256759	.4256759	Kurtosis	4.108209

FIGURE 5

DESCRIPTION OF $\frac{AWS(TTC)-AWS(DA)}{ABS(TTC)-ABS(DA)}$ FOR CASES IN WHICH TTC AND DA DO NOT DELIVER IDENTICAL MATCHINGS

	Percentiles	Smallest		
1%	0	0		
5%	0	0		
10%	.00012	0	Obs	78
25%	.023224	0	Sum of Wgt.	78
50%	.057794		Mean	.090785
		Largest	Std. Dev.	.0893094
75%	.152402	.275814		
90%	.234668	.300872	Variance	.0079762
95%	.275814	.319954	Skewness	1.023942
99%	.354396	.354396	Kurtosis	3.146456

FIGURE 6

DIFFERENCE BETWEEN THE PERCENTAGE OF STUDENTS WHO PREFER TTC TO DA AND THE PERCENTAGE OF STUDENTS WHO PREFER DA TO TTC, FOR CASES IN WHICH TTC AND DA DO NOT TRIVIALLY YIELD IDENTICAL ALLOCATIONS

In such scenarios, no student prioritized by a good school could ever obtain a worse allocation, in either of the mechanisms considered. The 95 percentile of this ratio is roughly below 27%.

With a welfare aggregator of the type

$$W(\mu) = ABS(\mu) - \delta AWS(\mu) + \gamma ABS_w(\mu),$$

even if we suppress the bonus to ABS_w ($\gamma = 0$), TTC would welfare-dominate DA in more than 90% of the cases in which the allocations differ, even for penalties to AWS, δ, as big as 4. Such a value of δ indicates that in order to compensate for a student who obtains a worse placement than her priority-giving school, we need at least four students who are assigned to a better school than each one's priority-giving school. With $\delta = 2$, TTC would welfare-dominate DA in all the simulations we computed.

The comparison of students' preferences between TTC and DA assignments is presented in Figure 6, with the difference between the percentages shown. For the scenarios where TTC and DA do not match the Serial Dictatorship allocation, a linear regression of this difference on the parameters of each scenario is displayed in Figure 7. We do not extend on redundant comments, since the behavior of this difference is so similar to that of $ABS(TTC) - ABS(DA)$.

Source	SS	df	MS			
Model	.538481376	4	.134620344	Number of obs	=	78
Residual	.075682988	73	.001036753	F(4, 73)	=	129.85
				Prob > F	=	0.0000
Total	.614164364	77	.007976161	R-squared	=	0.8768
				Adj R-squared	=	0.8700
				Root MSE	=	.0322

diffbest	Coef.	Std. Err.	t	P>|t|	[95% Conf.	Interval]
n12	.0054053	.0010235	5.28	0.000	.0033654	.0074451
n34	.0057407	.0010235	5.61	0.000	.0037009	.0077806
alpha	.0491478	.0022721	21.63	0.000	.0446194	.0536762
beta	-.0095235	.0021074	-4.52	0.000	-.0137235	-.0053236
_cons	-.2081452	.0339729	-6.13	0.000	-.275853	-.1404374

FIGURE 7

REGRESSION OF THE VARIABLE DESCRIBED ON THE PREVIOUS FIGURE WITH RESPECT TO THE PARAMETERS CHARACTERIZING EACH SCENARIO

In all scenarios considered, the percentage of students who prefer the assignment under TTC is at least slightly higher than the percentage of students who prefer the assignment under DA.

This statement is not true for every priority group. Interestingly, we actually find cases in which more students from Π_w prefer their DA assignment to their TTC assignment is higher than the opposite. This finding is consistent with the observation that disadvantaged students have increased access to good schools in general under TTC. The reason for this is that under TTC, these students are unable to disrupt trades among other students, which decreases their chances of getting into highly sought-after schools. However, as a trade-off, they have greater access to good schools overall.

6.2. *Other Observations for Arbitrarily Many Good Schools.*

6.2.1. *A lower bound for ABS under DA.*

It is easy to compute a preference-independent upper bound for ABS under DA. It is clear that, from each priority group Π_s, $s \in G$, the minimum mass of students in the group obtaining a slot at a good school is $\min\{q_s, n_s\}$. So the maximum mass of good school slots available to disadvantaged students is $\sum_{s\in G} \max\{0, q_s - n_s\}$. The infimum lottery number among students from Π_w that are assigned to the worst school, namely, l_w, is not higher than $\frac{1}{n_w} \sum_{s\in G} \max\{0, q_s - n_s\}$.

Here is where stability imposes its cost to ABS. Since the allocation is stable, every student who obtains a better allocation than her priority-giving school must at the very least have a lottery number below l_w. Hence:

PROPOSITION 2. $ABS(DA) \leq \frac{N}{n_w} \sum_{s\in G} \max\{0, q_s - n_s\}$.

COROLLARY 1. $ABS(DA) = 0$ *if all schools are weakly overprioritized.*

We end this subsection by noting that the TTC allocation, as it is not necessarily stable, is not bound by this upper limit.

6.2.2. *Accessibility.*

Another way to look at ABS would be to determine how many nonprioritized students receive slots from a school they like more than their assigned priority school. We consider a school to be *accessible* in a particular match if it assigns at least one slot to these nonprioritized students. A school is considered more accessible in one match compared to another if it assigns a greater number of slots to these students in the former match. We have a few observations to make about accessibility.

PROPOSITION 3. *If the worst school is weakly underprioritized, there is at least one good school that is not accessible under DA.*

We provide an informal proof here. Suppose all good schools are accessible under DA. This implies all students prioritized by some good school must be allocated to a good school. Again, this is a feature of stability: a prioritized student that is assigned to the worst school blocks accessibility. Therefore, the number of slots at good schools remaining for disadvantaged students is $\sum_{s\in G}(q_s - n_s) = n_w - q_w$. This makes the infimum lottery number among students from Π_w that are assigned to the worst school, namely, l_w, equal to $1-\rho_w$. Since $ABS(DA) \leq Nl_w$, $\rho_w \geq 1$ gives us a contradiction.

PROPOSITION 4. *There is at least a school that is weakly more accessible under TTC than under DA.*

Here, the proof is based on the Pareto-optimality of the TTC assignment. By Pareto-optimality, there must be a school s that gives all its slots only to students regarding that school as favorite. This implies that, among students from Π_s, only students from Π_{ss} could obtain a slot at s. Now, suppose school s is accessible under DA, otherwise we would be trivially done. This means that $q_s > n_{ss}$, since all students from Π_{ss} must be assigned to s in order to allow for accessibility (again this being a consequence of stability). But then, the mass of slots that school s gives to nonprioritized students wishing to be assigned there under TTC is $q_s - n_{ss}$, a number that cannot be exceeded by DA.

7. DISCUSSION

7.1. *The Boston Mechanism.* In earlier drafts of this manuscript, there was a greater focus on contrasting the Boston Mechanism (BM) and DA, which were then the two most commonly used mechanisms in practice. The aim was to demonstrate that both mechanisms could fall short in providing ABS due to various reasons. However, in the present version, we only briefly mention ABS under BM and do not include formal proofs in the main text.[23]

BM operates similarly to DA, with a key difference in the assignment algorithm. Unlike DA where all acceptances are provisional, BM requires each acceptance to be final during the allocation process. This makes BM manipulable. In previous versions of this article, we argued that if students have a relatively positive view of their priority-giving school, ABS could decline. A single Nash equilibrium could arise in which all students would rank their priority-giving school first, even if they prefer another school. As a result, good schools would be accessible to only a limited extent, if at all. This reasoning is different from that of DA, which may fail to provide ABS due to stability. BM, on the other hand, may fail because it could lead risk-averse families to adopt safe strategies, as noted in the empirical findings of Calsamiglia and Güell (2018).

7.2. *The Need for Stratification.* A common question is whether the results would change if the underlying assumptions of the model were altered. Specifically, one might question if the stratification assumption is too strict. It is possible that some students from disadvantaged areas might prefer a nearby worst school.

However, it is important to note that the worst school does not necessarily need to be the least preferred by all, in order to keep the results unchanged. It just needs to be not better than the priority-giving school for anyone, for DA to perform poorly in terms of ABS.

[23] We are thankful to referees that suggested a change of focus of the article.

TABLE 3

PREFERENCES, PRIORITY, AND ALLOCATION UNDER TTC

	First	*Second*	*Third*
i	b	a^{π}_{TTC}	w
j	$\mathbf{w}_{TTC}$	b^{π}	a
k	b_{TTC}	a	w^{π}

If we were to relax this assumption, it would be possible to construct an example where a nonstratified district would yield opposite results to what was presented here, with DA outperforming TTC in terms of ABS. Consider a finite-economy example where DA outperforms TTC even when the economy is replicated many times.

EXAMPLE 1. There are three schools a, b, and w with one slot each. There are three students i, j, and k with schools ranked as in Table 3. The superscript π indicates where the student has priority. The subscript TTC indicates the allocated slot under such mechanism.

$ABS(TTC) = 2/3$ (replica-invariant).

In this example, school w is better than the priority-giving school for one student. DA yields the allocation $(i \to b, j \to w, k \to a)$ (ABS = 1) with probability 1/2 (i's lottery number better than k's) and the TTC allocation in all other cases. Hence, $ABS(DA) = 5/6$ (even with any number of replicas).

We nevertheless argue that it is rather easy to come up with a relaxation of stratification, and still find large economies in which ABS collapses under DA. We say that the school district is *weakly stratified* when there is a set G of "good" schools for which the union of all of its prioritized students prefer all schools in G to all other schools. Notice that hierarchies are not that clear here. We allow for (some) nonprioritized students to actually dislike some schools in G.

Consider a large economy (with a continuum of agents), not necessarily with a worst school. But there is a group of "good" schools G for which, for prioritized students, these are actually superior to a set of "bad schools" B. However, for students from B, other factors (e.g., geographical proximity) might make some schools outside G preferred to those on G.

We say that $s \in G$ is *chain-linked* to another set B with $G \cap B = \varnothing$ if there is an array of positive-measured sets of agents $I_0, I_1, \ldots, I_K$ and schools $\{s_1, \ldots, s_K\} \subset G$ such that: (1) $I_0 \subset \Pi_B = \cup_{s \in B} \Pi_s$, (2) $I_k \subset \Pi_{s_k}$ $\forall k = 1, \ldots, K$, (3) $s_1 \succ_i s'$ $\forall s' \notin G, \forall i \in I_0$, (4) $s_k \succ_i s_{k-1}$ $\forall i \in I_{k-1}, \forall k = 2, \ldots, K$, and (5) $s_K = s$.

PROPOSITION 5. *Consider a continuum, weakly stratified school district with a set of weakly overprioritized good schools G. Assume that every $s \in G$ is chain-linked to a set of schools B such that $G \cap B = \varnothing$. Then $ABS_s(DA) = 0$ for all $s \in G$.*

The proof is rather short. Stability and (1) imply that no student in Π_B can obtain a slot in G. But then, since $\inf\{l(i) : i \in I_0, DA(i) \neq s_1\} = 0$ and by stability, no slot at s_1 is occupied but for prioritized students. Recursively, we find that $DA^{-1}(s_k) = \Pi_{s_k}$ for all $k = 2, \ldots, K$.

8. CONCLUSIONS

Since Abdulkadiroğlu and Sönmez (2003), the BM has received widespread criticism in the school choice literature. As a result, many cities have replaced it with the Gale–Shapley DA mechanism.[24] DA was adopted as a better alternative because it is not manipulable, protects

[24] See Pathak and Sönmez (2013) for evidence on the number of cities around the word where the BM has been banned.

nonstrategic parents, and provides more efficient assignments in setups with strict priorities. However, the debate between these two mechanisms was based on models that did not consider important realities about the school system, such as the differentiation between good and bad schools.

In this article, we study a simple model of school choice with coarse residential priorities and vertical differentiation between good and bad schools. Our findings suggest that if the goal of school choice is to improve ABS, DA may perform poorly. The priority structure, in the presence of a stratified school system, can greatly influence the final allocation in both mechanisms.

We also analyze TTC, an alternative that was previously considered by Abdulkadiroğlu and Sönmez (2003), but discarded due to its lack of stability. Our article highlights how stability can limit the ability of DA to improve ABS in the presence of bad schools and neighborhood priorities. TTC is more resilient to the priority structure because students at good schools can trade their slots without interference from the catchment area of bad schools. Our results show that TTC provides disadvantaged students with better access to good schools than DA, making it both more efficient and fair under the measure of ABS.

Our model does not include minority reserves as suggested by Hafalir et al. (2013). Further research should examine the comparison between mechanisms with and without reserves.[25]

This article highlights the significant impact that neighborhood priorities can have on the final allocation of students to schools, limiting the influence of student preferences. Although these priorities are often considered exogenous, they are actually a crucial aspect of the final assignment that can be adjusted by administrators.

Future studies should incorporate the design of these priorities as a crucial part of the mechanism design problem.

A.1. *Access to Better Schools (ABS) under Deferred Acceptance (DA) in the Finite-Economy Example.* As said in the main text, students with priority at different good schools would like to "exchange" their guaranteed slots, yet then the students from the bad school catchment area may block this trade. We want to derive the chances of exactly a number x of exchanges occurring. To gain more understanding, we illustrate a simple case where $n = 2$ and $x = 1$. We calculate all the cases in which this event happens. It could be that the two top-ranked students in the tie-breaking lottery are one student of type i_1 and another one of type i_2, and the third-ranked student is i_3. We could have picked $\binom{2}{1} = 2$ students from each type, and the order between types i_1 and i_2 does not matter (there are $2! = 2$ ways to arrange them). There are also $(6-3)!$ ways to arrange the remaining students among themselves. Hence, we find $2 \cdot 2 \cdot 2 \cdot 2! \cdot 3! = 96$ lottery outcomes satisfying this condition. But we have not covered all cases. It could also be that two students of type i_1 and another one of type i_2 occupy the first three positions in the lottery ranking, whereas the fourth position is occupied by an i_3 student. In this case, there is only one way, or $\binom{2}{2}$, to pick two students out of the two existing i_1 students. We could still pick $\binom{2}{1} = 2$ students from each of the other types. The way we arrange the two i_1 students and the i_2 student does not matter (there are $3!$ combinations). There are $(6-4)!$ ways to arrange the remaining students. We have found other $1 \cdot 2 \cdot 2 \cdot 3! \cdot 2! = 48$ such lottery outcomes. This number has to be multiplied by 2, to cover the final yet symmetric case in which two students of type i_2 and another one of type i_1 occupy the first three positions in the lottery ranking, whereas the fourth position is occupied by an i_3 student. We obtain a total of 192 favorable cases out of $6! = 720$ possible lottery outcomes. The probability of exactly one exchange with two students per school is $P(1, 2) = \frac{4}{15}$. More

[25] Preliminary simulations show that the qualitative results (TTC provides more ABS in general and particularly for disadvantaged students) hold. Related material is posted on https://sites.google.com/view/antoniomiralles.

generally,

$$P(x,n) = \frac{1}{(3n)!}[\binom{n}{x}\binom{n}{x}n(2x)!(3n-2x-1)! + \\ +2\sum_{i=x+1}^{n}\binom{n}{x}\binom{n}{i}n(x+i)!(3n-x-i-1)!] \\ = \binom{n}{x}\left[\frac{n}{3n-2x}\frac{\binom{n}{x}}{\binom{3n}{2x}} + 2\sum_{i=x+1}^{n}\frac{n}{3n-x-i}\frac{\binom{n}{i}}{\binom{3n}{x+i}}\right].$$

Let $X(n)$ denote the expected percentage of students that obtain a slot in a school better than their catchment area school under DA, when each school has n slots and n prioritized students. Then

$$X(n) = \frac{2}{3}\frac{1}{n}\sum_{x=1}^{n} xP(x,n).$$

The $\frac{2}{3}$ fraction appears because one-third of students (those with priority at the bad school) have no chance to escape from the bad school. Values for $X(n)$ are reported in Table 1 (main text). It can be shown that $X(n) \to 0$, in fact quite fast (e.g., $X(20) = 0.0177$).[26]

A.2. *Proof of Proposition 1.* We will divide the main theoretical result of the current article into two parts. In the first one, we show that Top- Trading Cycles (TTC) ABS dominates DA for disadvantaged students. In the second part, we show that the same result follows for all priority groups. The proof of the latter part uses the former one. That is the reason why we split the proof.

We will use the notation $k_{ss'}^{s''} = \sup\{l(i) : i \in \Pi_{ss'}, TTC(i) = s''\}$ as the maximum (worst) lottery allowing for a student with priority at s and preferred school s' to obtain a slot at s''.

REMARK A.1. $k_{ss'}^{s''} \geq k_{ws'}^{s''}$ for every $s'' \in G$.

We do not provide a formal proof, since this remark stems from a trivial observation. Students from $\Pi_s, s \in G$ could get access to another good school though trading of preexisting priorities or either through being pointed by means of the lottery number. Students from Π_w only count on the latter source, if any, for being assigned to a good school.

We regard school $s \in G$ as *accessible* under matching μ if $\lambda(\{i \notin \Pi_s, s \succ_i \pi(i) : \mu(i) = s\}) > 0$. Otherwise we regard it as inaccessible under μ.

PROPOSITION A.1. *When* $|G| = 2$, $ABS_w(TTC) \geq ABS_w(DA)$.

PROOF. There are three cases to consider, of which the first two are immediate.

Case 1 No good schools are accessible under DA. It obviously yields $ABS_w(DA) = 0 \leq ABS_w(\mu)$ for every other μ.

Case 2 Both good schools are accessible under DA. For every accessible school, it must be the case that all of its prioritized students obtain a slot at a good school, by means of stability. This means that $q_1 + q_2 \geq n_1 + n_2$ and that the amount of good slots available to students from Π_w is simply $q_1 + q_2 - n_1 - n_2$. No other matching could yield less good slots available to such students when $q_1 + q_2 \geq n_1 + n_2$, therefore $ABS_w(DA) \leq ABS_w(\mu)$ for every other μ.

[26] In a previous version of this article, we show that if we fix a proportion of agents wishing to exchange good school slots, the probability they *all* do so shrinks to zero at factorial speed as n grows.

Case 3 One good school (say school 1) is inaccessible under DA whereas the other one (say school 2) is accessible. Notice first that it must be the case that $q_1 \leq n_1$ and that $q_2 > n_2$. If it where the case that $q_1 > n_1$, school 1 would be accessible. Provided that school 1 is not accessible and that school 2 is, all students prioritized by school 2 must end up assigned there and besides some other slots should be available to others, concluding that $q_2 > n_2$.

We calculate $c_2 \equiv \sup\{l(i) : i \notin \Pi_2,\ DA(i) = 2\}$. Let $\tilde{c}_1 \equiv \sup\{l(i) :\ i \in \Pi_1,\ DA(i) = 1\}$. If $\tilde{c}_1 \geq c_2$, students from Π_{11} that do not obtain a slot at school 1 would not have a chance at school 2. Therefore, $c_2 = \frac{q_2 - n_2}{n_{12} + n_w}$, where the numerator takes into account that a mass n_2 of slots at school 2 are assigned to prioritized students. If $\tilde{c}_1 \leq c_2$, then no students from Π_{12} that do not obtain a slot at school 2 could have a chance at school 1. This yields $\tilde{c}_1 = q_1/n_{11}$. Moreover, we have $q_2 - n_2 = c_2(n_{12} + n_w) + (c_2 - \tilde{c}_1)n_{11}$, or $c_2 = \frac{q_1 + q_2 - n_2}{n_1 + n_w}$. Indeed,

$$c_2 = \min\left\{\frac{q_1 + q_2 - n_2}{n_1 + n_w}, \frac{q_2 - n_2}{n_{12} + n_w}\right\}.$$

Note that $ABS_w(DA) = c_2\lambda(\Pi_w)$.

We proceed to calculate $ABS_w(TTC)$. We aim to obtain the value of k^2_{w2}. Since $ABS_w(TTC) \geq k^2_{w2}\lambda(\Pi_w)$, we just need to show that $k^2_{w2} \geq c_2$. We run the TTC algorithm by letting school 2 be the first in pointing a student, as long as it keeps available slots (such choice does not alter the final allocation, which is unique). This does not necessarily imply that school 2 assigns all slots before school 1 does so. School 2 could point at a student that points at school 1, creating a cycle in which school 1 gives a slot.

We denote with t a mass of slots of school 2 assigned using this approach, which we use as a measure of time. If slots at school 1 are still unassigned when school 2 has filled capacity, we continue our measure of time from then on by means of counting assigned slots of school 1.

We use the notation $l_s(t)$ for the infimum lottery number among students from Π_s that are still unassigned at time t. There is $t > 0$ for which $l_w(t) < l_s(t)\ \forall s \in G$, since the first cycles involve students from Π_1 and Π_2 only.

We finally use t_s for the moment at which school s has given all of its slots, also called school s termination time. Clearly, $t_2 = q_2$. We consider subcases depending on whether termination time for school 1 is below or above t_2.

Case 3.1 $t_1 < t_2$.

Suppose that $l_1(t) > l_w(t)$ for all $t \in (0, t_1)$. This means that students from Π_1 have been so far allocated though cycles in which school 1 pointed them, implying $l_1(t_1) = \rho_1$.

All the students from Π_2 that are still unassigned (if any) are now assigned to school 2. There remains a mass $q_2 - n_2$ of slots of school 2 to allocate, along with the slots of the bad school. Since school 2 is the last good school still assigning slots, and thus all unassigned students point at it, we have $k^2_{ws} = k^2_{1s'}$ for every $s, s' \in G$. In such a case, we have $q_2 - n_2 = k^2_{w2}n_w + \max\{0, k^2_{w2} - l_1\}n_1$, or

$$\begin{aligned} k^2_{w2} &= \min\left\{\frac{q_2 - n_2 + q_1}{n_1 + n_w}, \frac{q_2 - n_2}{n_w}\right\} \\ &\geq \min\left\{\frac{q_2 - n_2 + q_1}{n_1 + n_w}, \frac{q_2 - n_2}{n_{12} + n_w}\right\} = c_2. \end{aligned}$$

Suppose instead that we reach a point in time $t' < t_1$ in which $l_1(t') = l_w(t')$. Once this equality arises, it holds for the rest of the assignment algorithm, since school 2 does not discriminate among nonprioritized students other than by the lottery number. This implies that $k^2_{1s} = k^2_{ws'} = k$ for every $s, s' \in G$. Therefore, we have the simple feasibility equation

$$q_1 + q_2 = n_2 + k(n_1 + n_w)$$

yielding $k^2_{w2} = \frac{q_2-n_2+q_1}{n_1+n_w} \geq c_2$.

Case 3.2 $t_2 \leq t_1$.

Since $q_2 > n_2$, all students from Π_2 have been assigned already. At the point at which the last student from Π_2 has been assigned, a mass n_{21} of slots of school 2 have been assigned to students from Π_{12}, thus $l_1(n_2) = n_{21}/n_{12}$. Note $l_w(n_2) = 0$ (no disadvantaged students have been assigned so far). We study the continuation of the TTC algorithm from then on, with a mass $q_2 - n_2$ of pending slots of school 2 to be assigned.

By school 2 being the first in giving all slots, we have $k^2_{h1} = 0$ for all schools h. If a student from Π_{h1} were pointed by school 2, this student would point at school 1, which has available slots.

We calculate the value of k^2_{w2}. Let $\delta \in [0, q_2 - n_2]$ be a mass of pending slots from school 2. While $l_w(n_2 + \delta) < l_1(n_2 + \delta)$, we have

$$l_1(n_2+\delta) = \frac{n_{21} + \delta \frac{n_{w1}}{n_w}}{n_{12}},$$
$$l_w(n_2+\delta) = \frac{\delta}{n_w}.$$

While $l_w(n_2 + \delta) < l_1(n_2 + \delta)$, each remaining slot of school 2 is assigned through cycles where school 2 points at a student from Π_w, yet with probability $\frac{n_{w1}}{n_w}$ it points to a student from Π_{w1} thus yielding a cycle where the slot is assigned to a student from Π_1.

Since $l_w(n_2) < l_1(n_2)$, school 2 starts pointing at a student from Π_w, and keeps doing so until l_w and l_1 coincide (if it does). If $l_w(q_2) \leq l_1(q_2)$, then all $q_2 - n_2$ remaining slots are assigned through cycles in which a student from Π_w is involved, hence:

$$k^2_{w2} = \frac{q_2-n_2}{n_w} \geq \frac{q_2-n_2}{n_{12}+n_w} \geq c_2.$$

Let us instead suppose that there is $\delta^* \in [0, q_2 - n_2)$ such that $l_w(n_2 + \delta^*) = l_1(n_2 + \delta^*) = l^*$, or $\delta^* = \frac{n_{21}n_w}{n_{12}-n_{w1}}$ and $l^* = \frac{n_{21}}{n_{12}-n_{w1}}$. From then on, there remains a mass $q_2 - n_2 - \frac{n_{21}n_w}{n_{12}-n_{w1}}$ slots of school 2 to assign. Notice that, in this case, $k^2_{w2} = k^2_{12}$ (once lottery numbers are tied, they keep tied until school 2 fills capacity). Since $k^2_{h1} = 0$ for all h, and since all students from Π_2 are already assigned,

$$q_2 - n_2 - \frac{n_{21}n_w}{n_{12}-n_{w1}} = \left(k^2_{w2} - \frac{n_{21}}{n_{12}-n_{w1}}\right)(n_{w2}+n_{12})$$

giving

$$k^2_{w2} = \frac{q_2-n_{22}}{n_{w2}+n_{12}} \geq \frac{q_2-n_2}{n_{w2}+n_{12}} \geq c_2.$$

□

PROPOSITION A.2. *When $|G| = 2$, $ABS_s(TTC) \geq ABS_s(DA)$ for every $s \in G$.*

PROOF. Also here we consider three cases, of which two of them are immediate:

Case 1 No good school is accessible under DA. Trivially, TTC cannot provide less ABS.

Case 2 One good school (say school 1) is not accessible and the other (school 2) is, under DA. In such a case, we have $c_1 = 0$ so obviously TTC cannot provide less access to school 1. As for school 2, we use the result of Proposition 6 ($k^2_{w2} \geq c_2$). Since $k^2_{12} \geq k^2_{w2}$ (Remark 1), TTC provides more access to school 2 from any other school.

Case 3 Both good schools are accessible under DA. Let us assume without loss of generality that $c_1 \leq c_2$. Since all good schools are accessible, all students prioritized there obtain a slot at a good school, leaving exactly a mass $q_1 + q_2 - n_1 - n_2 = n_w - q_w$ of slots for students from the bad school. We then have

$$c_2 = \frac{n_w - q_w}{n_w} = 1 - \rho_{w.}$$

As for c_1, it is calculated using

$$q_1 - n_{11} - (1 - c_2)n_{12} = c_1(n_{21} + n_{w1})$$

or

$$c_1 = \frac{q_1 - n_{11} - \rho_w n_{12}}{n_{21} + n_{w1}}.$$

Notice that

$$c_2 \leq \frac{q_2 - n_{22} - \rho_w n_{21}}{n_{12} + n_{w2}}$$

as an implication of $c_1 \leq c_2 = 1 - \rho_w$. In other words, for each $s \in G$:

$$c_s = \min\left\{1 - \rho_w, \frac{q_s - n_{ss} - \rho_w n_{ss'}}{n_{s's} + n_{ws}}\right\}.$$

As for the comparison to TTC: We analyze $k^s_{s's}$, where s and s' are both good schools. Along the proof, we assume that $k^s_{s's} < 1$, otherwise we would be trivially done.

We model the TTC algorithm as always making school s be the first in pointing at a student, as long as it has available slots. This is innocuous in that it does not affect the final allocation.

We denote with t a mass of slots of school s assigned using this approach, which we use as a measure of time. If some slots at school s' are still unassigned when school s has filled capacity, we continue our measure of time from then on by means of counting assigned slots of school s'.

We finally use t_s for the moment at which school s has given all of its slots, also called school s termination time. Clearly, $t_s = q_s$. We consider subcases depending on whether termination time for school s' is below or above t_s.

Case 3.1 $t_s < t_{s'}$.

This implies that $k^s_{hs'} = 0$ for all schools h. If school s points at a student from $\Pi_{hs'}$, she will point at school s', since it still has available slots.

But then, $q_s = n_{ss} + k^s_{s's} n_{s's} + k^s_{ws} n_{ws}$, and $k^s_{s's} \geq k^s_{ws}$ implies

$$k^s_{s's} \geq \frac{q_s - n_{ss}}{n_{s's} + n_{ws}} \geq \frac{q_s - n_{ss} - \rho_w n_{ss'}}{n_{s's} + n_{ws}} \geq c_s.$$

Case 3.2 $t_{s'} < t_s$.

This implies $q_{s'} < n_{s'}$ provided $k^s_{s's} < 1$. (If $q_{s'} \geq n_{s'}$ we have that every student from $\Pi_{s'}$ is assigned to a good school, and since $t_{s'} < t_s$ we have $k^{s'}_{s's} = 0$—a student from $\Pi_{s's}$ cannot point at s' since she prefers s which has pending unassigned slots—and then $k^s_{s's} = 1$). Since $\max\{c_1, c_2\} = 1 - \rho_w > 0$, or $q_w < n_w$, we must have $q_s > n_s$ (provided $\sum_h q_h = \sum_h n_h$).

Because s is the last good school in assigning remaining slots to students, $q_s > n_s$ implies that school s eventually points at students according only to their lottery number. Consequently, $k^s_{wh} \geq k^{s'}_{wh}$, for both $h \in G$. Because of the lottery number criterion, $k^s_{ws} = k^s_{ws'}$. With this and Proposition 6 implying $ABS_w(TTC) \geq ABS_w(DA) = n_w(1-\rho_w)$, we must have that $k^s_{ws} = \max_{g,h \in G} k^g_{wh} \geq 1-\rho_w$. But then, by Remark 1, $k^s_{s's} \geq k^s_{ws} \geq 1-\rho_w \geq c_s$.

This completes all relevant cases (the remaining knife-edge case in which slots from schools 1 and 2 are exhausted simultaneously can be treated as part of subcases 3.1 or 3.2, indistinctively). $\square$

Supporting Information

Additional supporting information may be found online in the Supporting Information section at the end of the article.

Table A1: n1 = n2 = 10; n3 = n4 = 20
Table A2: n1 = n2 = 15; n3 = n4 = 15
Table A3: n1 = n2 = 10; n3 = n4 = 25
Table A4: n1 = n2 = 15; n3 = n4 = 20
Table A5: n1 = n2 = 15; n3 = n4 = 25
Table A6: n1 = n2 = 20; n3 = n4 = 20
Table A7: n1 = n2 = 20; n3 = n4 = 10
Table A8: n1 = n2 = 25; n3 = n4 = 10
Table A9: n1 = n2 = 20; n3 = n4 = 15
Table A10: n1 = n2 = 25; n3 = n4 = 15
Data S1

REFERENCES

ABDULKADIROĞLU, A., Y. CHE, P. PATHAK, A. ROTH., AND O. TERCIEUX, "Efficiency, Justified Envy, and Incentives in Priority-Based Matching," *American Economic Review: Insights* 2 (2020), 425–42.

———, ———, and Y. YASUDA, "RESOLVING CONFLICTING PREFERENCES IN SCHOOL CHOICE: THE, 'Boston' Mechanism Reconsidered," *American Economic Review* 101 (2011), 399–410.

———, ———, and ———, " EXPANDING, 'Choice' in School Choice," *American Economic Journal: Microeconomics* 7 (2015), 1–42.

———, P. A. PATHAK, A. E. ROTH, and T. SÖNMEZ, "Changing the Boston School Choice Mechanism," Unpublished Manuscript, 2006.

———, and T. SÖNMEZ, "School Choice: A Mechanism Design Approach," *American Economic Review* 93 (2003), 729–47.

AUMANN, R. J., "Markets with a Continuum of Traders," *Econometrica* 32 (1964), 39–50.

CALSAMIGLIA, C., and M. GÜELL, "Priorities in School Choice: The Case of the Boston Mechanism in Barcelona," *Journal of Public Economics* 126 (2018), 20–36.

———, C. FU, and M. GÜELL, "Structural Estimation of a Model of School Choices: The Boston Mechanism vs Its Alternatives," *Journal of Political Economy* 128 (2020), 642–80.

CHEN, Y., and T. SÖNMEZ, "School Choice: An Experimental Study," *Journal of Economic Theory* 127 (2006), 202–31.

———, and O. KESTEN, "From Boston to Chinese Parallel to Deferred Acceptance: Theory and Experiments on a Family of School Choice Mechanisms," Mimeo, Carnegie Mellon University, 2013.

COMBE, J., O. TERCIEUX, and C. TERRIER, "The Design of Teacher Assignment: Theory and Evidence," *Review of Economic Studies* 89 (2022), 3154–222.

EHLERS, L., and M. MORRILL, "(Il)legal Assignments in School Choice," *Review of Economic Studies* 87 (2019), 1837–75.

ERGIN, H., "Efficient Resource Allocation on the Basis of Priorities," *Econometrica* 70 (2002), 2489–97.

———, and A. ERDIL, "What's the Matter with Tie-Breaking? Improving Efficiency in School Choice," *American Economic Review* 98 (2008), 669–89.

———, and T. SÖNMEZ, "Games of School Choice under the Boston Mechanism," *Journal of Public Economics* 90 (2006), 215–37.

FEATHERSTONE, C. R., "Rank Efficiency: Investigating a Widespread Ordinal Welfare Criterion," Unpublished Manuscript, 2014.

HAFALIR, I., M. B. YENMEZ, and M. A. YILDIRIM,"Effective Affirmative Action in School Choice," *Theoretical Economics* 8 (2013), 325–63.

Hastings, J., T. Kane, and D. Staiger, "Heterogeneous Preferences and the Efficacy of School Choice," combines and replaces National Bureau of Economic Research Paper Working Papers No. 12145 and No. 11805, 2010.

———, and J. Weinstein, "Information, School Choice, and Academic Achievement: Evidence from Two Experiments," *Quarterly Journal of Economics* 123 (2008), 1373–414.

He, Y., "Gaming the Boston School Choice Mechanism in Beijing," Mimeo, Toulouse School of Economics, 2017.

Kesten, O., "School Choice with Consent," *Quarterly Journal of Economics* 125 (2010), 1297–348.

Lavy, V., "Effects of Free Choice among Public Schools," *Review of Economic Studies* 77 (2010), 1164–91.

Leshno, J., and I. Lo, "The Cutoff Structure of Top Trading Cycles in School Choice," *Review of Economic Studies* 88 (2021), 1582–623.

Miralles, A., "School Choice: The Case for the Boston Mechanism," Unpublished Manuscript, 2008.

Pathak, P. A., and T. Sönmez, "Leveling the Playing Field: Sincere and Strategic Players in the Boston Mechanism," *American Economic Review* 98 (2008), 1636–52.

———, and ———, "School Admissions Reform in Chicago and England: Comparing Mechanisms by Their Vulnerability to Manipulation," *American Economic Review* 103 (2013), 80–106.

Roth, A., "What have We Learned from Market Design?" *Economic Journal* 118 (2008), 285–310.

Terrier, C., P. A. Pathak, and K. Ren, "From Immediate Acceptance to Deferred Acceptance: Effects on School Admissions and Achievement in England," NBER Working Paper 29600, 2021.

Troyan, P., "Comparing School Choice Mechanisms by Interim and Ex-Ante Welfare," *Games and Economic Behavior* 75 (2012), 936–47.

———, D. Delacrétaz, and A. Kloosterman, "Essentially Stable Mechanisms," *Games and Economic Behavior* 120 (2020), 370–90.