

NBER WORKING PAPER SERIES

SCHOOL CHOICE

Atila Abdulkadiroglu
Tommy Andersson

Working Paper 29822
<http://www.nber.org/papers/w29822>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
March 2022

We would like to thank Alan Jeske for comments. Andersson gratefully acknowledge financial support from the Jan Wallander and Tom Hedelius Foundation (research grant P2018–0100). The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2022 by Atila Abdulkadiroglu and Tommy Andersson. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

School Choice
Atila Abdulkadiroglu and Tommy Andersson
NBER Working Paper No. 29822
March 2022
JEL No. D02,D47,I26

ABSTRACT

School districts in the US and around the world are increasingly moving away from traditional neighborhood school assignment, in which pupils attend closest schools to their homes. Instead, they allow families to choose from schools within district boundaries. This creates a market with parental demand over publicly-supplied school seats. More frequently than ever, this market for school seats is cleared via market design solutions grounded in recent advances in matching and mechanism design theory. The literature on school choice is reviewed with emphasis placed on the trade-offs among policy objectives and best practices in the design of admissions processes. It is concluded with a brief discussion about how data generated by assignment algorithms can be used to answer contemporary empirical questions about school effectiveness and policy interventions.

Atila Abdulkadiroglu
Department of Economics
Duke University
213 Social Sciences Building
Durham, NC 27708
and NBER
atila.abdulkadiroglu@duke.edu

Tommy Andersson
Lund University
Department of Economics
Sweden
tommy.andersson@nek.lu.se

School Choice

Atila Abdulkadirođlu¹ and Tommy Andersson²

¹Duke University and NBER

²Lund University and Stockholm School of Economics

Abstract

School districts in the US and around the world are increasingly moving away from traditional neighborhood school assignment, in which pupils attend closest schools to their homes. Instead, they allow families to choose from schools within district boundaries. This creates a market with parental demand over publicly-supplied school seats. More frequently than ever, this market for school seats is cleared via market design solutions grounded in recent advances in matching and mechanism design theory. The literature on school choice is reviewed with emphasis placed on the trade-offs among policy objectives and best practices in the design of admissions processes. It is concluded with a brief discussion about how data generated by assignment algorithms can be used to answer contemporary empirical questions about school effectiveness and policy interventions.

1 Introduction

School choice means expanding families' schooling options. This may come in the form of financial support for private schooling or in the form of expanding options available to parents. In the US, the former includes tax-credit education savings accounts,¹ school vouchers,² education savings accounts,³ and individual tax credits

¹These accounts allow taxpayers to receive full or partial tax credits on education-related expenses, such as private school tuition, private tutoring, etc.

²Districts allocate public funds via vouchers to qualifying families to partially or fully cover their child's private school tuition.

³Education savings accounts allow parents who withdraw their children from public schools to use public funds to cover expenses for private school, online learning programs, private tutoring, etc.

and deductions.⁴ The latter includes inter/intra-district public school choice that allows parents to choose non-local schools, charter schools,⁵ magnet schools,⁶ home schooling and several other options.⁷

This chapter concerns choice over public schools, also referred to as inter/intra-district public school choice in the education literature. School districts in the US and around the world are increasingly moving away from traditional neighborhood school assignment, in which pupils attend closest schools to their homes. Instead, they allow families to choose from schools within district boundaries. This creates a market with parental demand over publicly-supplied school seats. More frequently than ever, this market for school seats is cleared via market design solutions grounded in recent advances in matching and mechanism design theory. This chapter reviews the literature on school choice, with emphasis placed on the trade-offs among policy objectives and best practices in the design of admissions processes.

The roots of school choice in the US trace back to historic educational inequality in urban districts. In traditional neighborhood-based assignment, pupils are assigned to schools located in their residential area. Consequently, families can choose schools to the extent they can afford to live in the neighborhoods. As a result, housing markets and public-school financing regimes foster segregation across neighborhoods along income and other correlated factors, such as race and ethnicity. This limits low-income families' access to good schools, which are usually located in wealthy neighborhoods. Parental choice over public schools has therefore become a major policy tool to combat inequality in access to schools.

Today, another reason for parental choice is that school districts serve more heterogeneous populations. To meet their students' varying needs, school districts have been moving away from one-size-fits-all models of schooling, as well as adopting alternative curricula and pedagogical approaches. They have introduced alternative school management models, such as charter schools, to create competitive forces on schools and teachers. Furthermore, to increase diversity, school districts have opened schools with promising features in under-served areas to attract students from wealthier neighborhoods. Parental choice and preference-based school assignment become an integral part of enrollment planning, as neighborhood-based assignment can no longer clear the market with highly heterogeneous student populations and schooling options.

⁴Under such plans, individuals can claim state tax credits and deductions for approved educational expenses, such as private school tuition, private tutoring, etc.

⁵Charter schools are publicly funded, privately managed, and freely available schools.

⁶Magnet schools are public schools with specialized curricula and programs that are not available in traditional public schools.

⁷For a more detailed list and description of the parents' options, see: <https://www.edchoice.org/school-choice/types-of-school-choice/>.

The question of how to design admissions in school choice programs was introduced by Abdulkadiroğlu and Sönmez (2003) as an application of matching theory (Gale and Shapley, 1962). Since the introduction of the problem, economists have been deeply involved in the study and design of student assignment systems,⁸ starting with the redesign of student assignment systems in Boston and New York City.

The Boston Public Schools (BPS) has a long history with school choice. BPS used to grant students priorities in admissions at schools based on distance between school and home address, siblings' enrollment status, etc. Accordingly, a student within walking distance of a school and with a sibling enrolled at the school would receive high priority in admissions. However, if the school is over-subscribed, a student could take advantage of such high priority only if she ranked it as her first choice in the application form. That was because the BPS assignment algorithm favored students at a school if they ranked it high in their application forms. This feature of the BPS admissions process did not go unnoticed by Boston families. Parents that lived in an affluent part of the city, designated as the West Zone, had a Yahoo email list to educate newcomers about how to most strategically fill out application forms. The strategic aspect of the process was also being advertised in the official BPS school guide. Lack of ability to give parents straightforward advice about applications, such as listing schools in preference order in the application form, created a burden on families as well as district officials. That difficulty formed the basis for the redesign of the BPS school match (Abdulkadiroğlu and Sönmez, 2003; Abdulkadiroğlu et al., 2005b).

The redesign of the New York City High School Match was initiated by a congestion problem. Incoming high school students could apply to at most five schools in five boroughs in the City. Each school could extend offers to students, but only in an uncoordinated manner. A student could get multiple offers, a single offer, or no offer. In the latter case, the student would be put on the waiting list of the schools she ranked. Students with offers could accept at most one offer. Unassigned students would go through two more rounds of offers/acceptances. In 2003, only about half of the applicants received offers initially, and about 34 percent of them received multiple offers. When the admissions process concluded, approximately 30 percent of applicants were assigned to a school outside of their choice list (Abdulkadiroğlu et al., 2005a).

Both cities replaced their admissions processes with assignment algorithms based on the celebrated deferred acceptance algorithm of Gale and Shapley (1962). Although variants of the deferred acceptance algorithm now dominate school choice, other algorithms have been adopted in the field.

⁸In particular, the authors of this chapter have assisted major school districts in the US and Europe in introducing and redesigning school admissions processes.

The rest of this chapter describes the school choice problem using matching and mechanism design theory. Section 2 sets up the theoretical framework and introduces a fairly general matching model with a set of students and schools. Students rank schools in a strict preference order in their applications, while schools may rank students using admissions priorities, entrance exam scores, academic records and preferences of their admissions office, etc. Importantly, a school may rank multiple students equally. A student can be matched with at most one school, whereas a school can be matched with as many students as the number of available seats at the school. In the literature, such a model is referred to as a two-sided many-to-one matching model.⁹ Section 3 formalizes common policy objectives that have been critical in the design of real-life admissions procedures. Section 4 introduces matching algorithms that have been central to the theory and applied in school choice programs around the world. These algorithms are analyzed in Section 5 for the commonly analyzed case in which schools rank students in a strict order. Section 6 focus on a more general case in which some schools may rank students equally.

When districts have preferences over diversity of students at schools, a simple ranking of students may not capture district’s policy objectives. Section 7 generalizes the model to account for more complex objectives in admissions. This generalization lays a foundation for the analysis of school choice under various distributional constraints studied in Section 8, which also discusses the common practice of allocating school seats via multiple separate application and registration processes. Section 9 contains a brief discussion about how data generated by assignment algorithms can be used to answer contemporary empirical questions about school effectiveness and policy interventions. Finally, Section 10 provides some concluding remarks and directions for future work.

2 The Model

Matching theory constitutes the foundation for school choice problems (Abdulkadiroğlu and Sönmez, 2003). In their seminal paper, Gale and Shapley (1962) introduced a basic matching problem in the context of college admissions. In their model, students rank colleges and colleges rank students in strict preference order. Each college has a certain capacity to be filled with students. A matching of students and colleges is said to be feasible if each student is matched with at most one college

⁹In a one-to-one matching model, by contrast, each agent can be matched with at most one agent from the other side of the market. Although this model is important in the development of the matching theory, it does not have direct applications for school choice. For a comprehensive and recent survey of matching theory, including one-to-one matching, see Echenique et al. (2022). See also Roth (2018) and Sönmez and Ünver (2011).

and the number of students matched with a college does not exceed the capacity of the college. The so-called stability notion of Gale-Shapley is central in finding a matching of students and colleges: if a student and a college are matched, the student prefers the college to remaining unassigned and the college prefers filling one of its seats with the student to leaving the seat empty. Moreover, whenever a student prefers a college to her match, either the college does not rank her or it is already fully matched with students that the college ranks higher than the student.

The basic Gale-Shapley model resembles many aspects of a real-life school choice programs. For example, when families choose schools for their children through a centralized admissions system, they are asked to rank schools in strict preference order. At the same time, schools rank students. The rankings of students at schools may be determined by entrance exams, certain admissions priorities, e.g. proximity of residential location to school, enrollment status of siblings at schools, preferences of the admissions office, lottery numbers, or any combination of these criteria. Moreover, the notion of stability applies naturally when school rankings reflect school preferences. When rankings do not involve school preferences, the stability notion has a straightforward interpretation: A matching of students and schools is said to be free of justified envy if whenever a student prefers a school to her match, either she is not eligible (i.e., the school does not rank her), or the school is already fully matched with students that the school ranks higher (Abdulkadiroğlu and Sönmez, 2003).

In practice, however, real-world school choice problems differ from the basic Gale-Shapley model in several important ways. First, depending on the institutional setting, stability may not be a cause of concern. In that case, stability has negative welfare consequences on the matching. Second, unlike in the early applications of Gale and Shapley (1962), strict ranking of students by schools is a knife-edge phenomenon in practice. In real-world choice programs, schools often sort students into thick priority groups. For example, a student living within a certain distance from a school may be granted neighborhood priority at the school, and ties among equal priority students are often broken by lottery. Treating that final strict ranking via lotteries has welfare implications. Third, the school choice problem is often bounded by distributional constraints such as imposing requirements on the demographic distribution of students at each school in the district. Fourth, a district may have multiple school sectors, such as private schools, traditional public schools, alternative schools such as magnets and charters, and the admission processes may not be coordinated across sectors in the district.

A school choice problem that captures many aspects of a typical school choice program consists of a set of students $N = \{1, \dots, |N|\}$ that are applying to schools in the set $S = \{s_1, \dots, s_{|S|}\}$. Each school s has a certain number of seats to be

allocated to students. The number of available seats at school s is denoted by q_s , also referred to as the capacity of school s .

Each student $i \in N$ has strict preferences P_i over schools and being unmatched, denoted $S \cup \{i\}$, where $\{i\}$ represents being unmatched for student i . For any two schools s and s' , the notion $sP_i s'$ means that student i prefers s to s' . School s is said to be acceptable for student i if and only if $sP_i i$. The preference P_i can also be viewed as the preferences of the family or legal guardian of student i , as they are usually the ones who submit application forms on behalf of the student. For convenience, the “at least as good as relation” R_i is defined as follows:

$$sR_i s' \iff sP_i s' \text{ or } s = s'.$$

In words, school s is at least as good as school s' from the perspective of student i if (i) she strictly prefers s to s' , or (ii) s and s' are the same school.

In the model, it is assumed that each student ranks schools without any regard to enrolled students at schools. This is a rather simplified version of reality. In general, families care about schools as well as peers of their children at schools. However, it is difficult, and most of the time impossible, to extend the theory of matching by generalizing student preferences over sets of students enrolled at schools. Despite that observation, P_i represents a typical application form in which families are asked to rank schools in a strict order.

Each school $s \in S$ has a weak relation \succsim_s over $N \cup \{s\}$, where $\{s\}$ represents keeping a seat empty. For any two students i and j , the notion $i \succsim_s j$ means that school s ranks student i higher than or equally as student j . When these rankings reflect admissions priorities, this means that i has a weakly higher priority at s than j . The strict ranking \succ_s is obtained in the usual manner:

$$i \succ_s j \iff i \succsim_s j \text{ and not } j \succsim_s i.$$

Student i is said to be eligible at school s if and only if $i \succ_s s$.

Unlike family preferences over schools, schools may use various criteria to rank students. For instance, a student whose sibling is already enrolled at school s may be granted sibling priority, which may place her higher in \succsim_s . Likewise, students living within a certain distance from a school may be given neighborhood or walk-zone priority. These two priorities are very common in practice. In general, however, priorities can be based on any measurable criteria, such as priority for students coming from failing or closing schools, for military children, and for children of staff at a school. Such priorities sort students into thick priority groups.

The relation \succsim_s may also represent a ranking of students according an entrance examination score. For example, the seats in eight high schools in New York City

are allocated based on student scores on the Specialized High School Admissions Test. Chicago Public Schools, by comparison, use a composite score for admissions to their selective high schools. This score is composed of the 7th grade standardized test, the 7th grade final grades, and the Selective Enrollment Test with equal weights. In addition, the relation \succsim_s may also reflect preferences of the school management over students. For instance, some audition schools in New York City rank students based on priorities and also interviews and arts performances but they do not use any academic records. While \succsim_s may be due to a variety of factors, such as those described above, This chapter abstracts from these considerations and simply refers to \succsim_s as a school ranking.

As in the case of family preferences, schools may have more complicated preferences over sets of students that may not be captured by a simple ranking of individual students. For example, Ed opt schools in New York City use a standardized English Language Arts test scores from 7th grade and aim to enroll 16 percent from top, 68 percent from middle and 16 percent from lowest score students. The theory with such preferences will be extended in Section 7. However, it is worth noting here that school preferences over groups of students can be accommodated to the extent that they do not introduce complementarities among students”, as would be the case, for example, if a school prefers to admit a quarterback for their football team only if they can also admit an offensive guard at the same time and vice versa. Extending the theory to accommodate such preferences is possible, but would require more restrictive assumptions.

To summarize, a school choice problem consists of:

- a set of students $N = \{1, \dots, |N|\}$,
- a set schools $S = \{s_1, \dots, s_{|S|}\}$,
- each school $s \in S$ has a capacity of q_s ,
- each student $i \in N$ has a strict preference relation P_i over $S \cup \{i\}$,
- each school $s \in S$ has a weak ordering \succsim_s over $N \cup \{s\}$.

Let $q = (q_s)_{s \in S}$ denote capacities of the schools, $P = (P_i)_{i \in N}$ denote the preference profile of the students, and $\succsim = (\succsim_s)_{s \in S}$ denote the profile of school rankings. Then a problem is given by (N, S, q, P, \succsim) .

A typical school choice program consists of several types of schools with different admissions criteria. The model captures a wide variety of real-world admissions criteria. When school rankings are strict and reflects school preferences, the model

reduces to that of Gale and Shapley (1962). When school rankings are strict and determined by college entrance exams, the model reduces to the model of Balinski and Sönmez (1999). When the ranking at each school is determined by some exogenous and verifiable rules, such as exam scores, sibling and neighborhood priority, the school will be referred to as non-strategic. If rankings reflect school preferences, it will be referred to as strategic. Throughout, students are assumed to be strategic since student preferences are private information and cannot be verified.

If μ denotes a matching of students and schools, $\mu(a)$ is the match of $a \in N \cup S$, such that each student $i \in N$ is assigned a school or remains unmatched, i.e. $\mu(i) \in S \cup \{i\}$, each school $s \in S$ is matched with a set of students up to its capacity, i.e. $\mu(s) \subset N$ and $|\mu(s)| \leq q_s$, and $\mu(i) = s \in S$ if and only if $i \in \mu(s)$.

Matchings will frequently be compared throughout the chapter. In such situations, a student is indifferent between two distinct matchings whenever she is assigned the same school at these matchings; otherwise she prefers the one that matches her with a school that she ranks higher in her preference list. This is a consequence of the assumption that students only care about schools but not who else is enrolled at a school. Define

$$\mu R_i \mu' \iff \mu(i) R_i \mu'(i).$$

Finally, a mechanism φ determines a matching for any given problem (N, S, q, P, \succsim) . Some parts of a problem will be fixed throughout the analysis below. For example, when only alternative student preferences are considered, (N, S, q, \succsim) will be fixed and a problem will be given by P . In that case, the matching produced by φ is denoted by $\varphi(P)$. Student i 's match is $\varphi_i(P) \in S \cup i$ and the set of students matched with school s is $\varphi_s(P) \subset N$.

3 Policy Objectives

How should students and schools be matched? Prior to the intervention of market design in the early 2000s, school districts came up with their own ad hoc solutions. For instance, a student in New York City could be simultaneously admitted to multiple high schools, leading to congestion in the system. Boston Public Schools matched students to schools via a process that forced parents to make complicated decisions and submit strategically constructed preference lists.

Economic theory takes a positive approach to study such problems. Specifically, it identifies various policy objectives, e.g., efficient utilization of school capacities. It formulates them as axioms and studies solutions around these axioms. More

importantly, it also characterizes the set of solutions that satisfy various axioms. The economic approach also clarifies the set of solutions that satisfy these axioms, thereby clarifying what goals can be achieved and how. Furthermore, it identifies trade-offs when multiple objectives cannot be simultaneously achieved.

This section introduces three major policy objectives that have been critical in the design of real-life school admissions procedures, namely efficiency, stability and strategy-proofness. Efficiency is the most obvious of the three desiderata. It concerns the welfare of economic agents in the matching process. Stability, as discussed above, concerns meeting student preferences and school rankings simultaneously. Strategy-proofness aims to eliminate gaming of the system by strategic agents via manipulation of preferences during the application process. These three objectives is formalized as axioms below.

3.1 Pareto Efficiency

Efficiency can be defined in various ways. The notion of Pareto efficiency eliminates the possibility of improving an agent's welfare without harming others. In the context of school choice, a matching μ Pareto dominates another matching μ' if every student weakly prefers μ to μ' , i.e.

$$\mu(i)R_i\mu'(i) \text{ for all } i \in N,$$

and at least one student strictly prefers μ to μ' ,

$$\mu(i)P_i\mu'(i) \text{ for some } i \in N.$$

A matching μ is Pareto efficient if there does not exist another matching μ' that Pareto dominates μ . In other words, a matching is Pareto efficient whenever it is impossible to find another matching that matches one student to one of her more preferred schools without hurting the match of another student. Pareto efficiency is a natural policy objective.

Notice that the above defines Pareto efficiency only from the perspective of students. In general, when school rankings represent preferences for some schools, the notion can be generalized from the perspective of students and such schools. However, there is a close connection with efficiency and stability, which is defined next, so that Pareto efficiency from the perspective of students becomes a natural objective even in the presence of schools with preferences.

3.2 Stability

At the most basic level, a student's being ranked by a school may be interpreted as an indication of her eligibility, while not being ranked may be interpreted as ineligibility. However, school rankings can more generally be thought of as reflecting district preferences over admissions policies. For instance, assigning siblings the same school simplifies the task of transportation for families; it may also create positive spillover effects among siblings, as they share their school experience. Assigning pupils closer to home may also encourage families to engage with schools more.

When students are ranked according to an entrance exam, GPA or a composite score, favoring students with lower scores may cause litigation.¹⁰ School rankings may as well reflect school preferences over students. For these reasons, a district may prefer to assign seats to students who are higher in school rankings before considering other students.

The stability notion (Gale and Shapley, 1962) captures such preferences in assignment. A matching μ is individually rational if every student is matched with an acceptable school at which she eligible or remains unmatched. A student-school pair blocks μ if they mutually prefer to be matched to each other. Matching μ is stable if it is individually rational and cannot be blocked by any student-school pair. Formally, μ is stable if there does not exist any student-school pair $(i, s) \in N \times S$ such that $i \succ_s s$, $s \succ_i \mu(i)$, and either there are available seats at the school, $|\mu(s)| < q_s$ or there is i' who is assigned s and ranked lower by s , i.e. $i' \in \mu(s)$ and $i P_s i'$. A matching that is not stable is said to be unstable.

In matching models, there is a fundamental trade-off between stability and Pareto efficiency. In some cases, it may be impossible to simultaneously meet both criteria, as the following example taken from Roth (1982) demonstrates:

Example 1. Suppose that there are three students, 1, 2 and 3, and three schools, s_1 , s_2 and s_3 . Each school has one available seat. The student preferences and

¹⁰Probably the most famous legal case is *Wessmann v. Boston School Committee*, in which the plaintiff were denied admissions to the prestigious Boston Latin School because of race-conscious admissions at the time. See: <https://law.justia.com/cases/federal/district-courts/FSupp/996/120/1626115/>.

school rankings are given by:¹¹

$$\begin{aligned}
1 &: s_2 P_1 s_1 \\
2 &: s_1 P_2 s_2 P_2 s_3 \\
3 &: s_1 P_3 s_2 P_3 s_3 \\
s_1 &: 1 \succ_{s_1} 3 \succ_{s_1} 2 \\
s_2 &: 2 \succ_{s_2} 1 \succ_{s_2} 3 \\
s_3 &: 2 \succ_{s_3} 3
\end{aligned}$$

In this problem, there is a unique stable matching, namely

$$\mu = ((1, s_1), (2, s_2), (3, s_3)),$$

where (i, s) means student i is matched with school s . The stable matching μ is Pareto dominated by the following Pareto efficient matching:

$$\mu' = ((1, s_2), (2, s_1), (3, s_3)).$$

Note that students 1 and 2 are matched to their first choices under μ' . However, $(3, s_1)$ forms a blocking pair because 3 prefers s_1 to s_3 , and s_1 ranks 3 higher than 2. \square

Stability is an appealing property for several reasons. Roth (2002) shows that stable matching is key for long term survival of centralized markets in the entry level labor markets.¹² In that study, most markets that selected stable matchings kept operating, while most markets that selected unstable matchings were abolished or replaced.

When schools can act strategically, an unstable matching yields a blocking student-school pair that prefers to be matched together. Presence of blocking pairs undermines the entire matching process, as they have incentive to circumvent the process. Unstability was a defining feature of the old unstable matching system in New York City. Indeed, it was reported that some schools concealed capacity in an effort to be matched later with preferable students.¹³

¹¹Recall that each student has a preference relation over the set of schools and herself. In most examples in this chapter, student preferences are truncated below the last acceptable school for notational simplicity. The same simplification applies to school rankings. In this example, only schools s_1 and s_2 are acceptable for student 1, while only students 2 and 3 are eligible at school s_3 .

¹²See Roth (1984a) and Roth and Peranson (1997).

¹³Joel Klein, the deputy chancellor, quoted in the New York Times (November 19, 2004): “Before you might have a situation where a school was going to take 100 new children for 9th grade, they might have declared only 40 seats, and then placed the other 60 outside the process.”

School rankings are a major tool to implement district policies via admissions priorities such as sibling and neighborhood priority. Any unstable matching, therefore, can be deemed as violating admission rules. When a student i prefers school s to her own match, the match can be justified if all the seats at s are assigned to students with higher rank or higher priority. In other words, i 's priority at s is violated if i prefers s to her match and another student with lower priority is assigned to s . Such priority violation is also referred to as justified envy (Abdulkadiroğlu and Sönmez, 2003). Justified envy is concerned with eliminating priority violations, whereas the notion of blocking pairs is concerned with the agents' incentives to mutually circumvent a match. This chapter will refer to both forms as stability.

3.3 Strategy-Proofness

Poorly designed admissions procedures may incentivize families to circumvent a match. A well-known such procedure is “immediate acceptance” or “first-preference-first,” which is formally defined later in the chapter. This procedure gives higher priority to students at a school who rank the school higher in their choice list. Consequently, a family may improve their odds of assignment to a school by elevating the ranking of the school in their choice list. This aspect is usually recognized easily by district officials and families. In Boston, where an immediate acceptance process was in place until 2005, the school guide for Boston public schools (BPS, 2004) used to inform parents about strategically listing schools (quotes in original):

“For a better chance of your ‘first choice school’ ... consider choosing less popular schools. Ask Family Resource Center staff for information on ‘underchosen’ schools.”

And a family email list known as The West Zone Parent Group (WZPG) in Boston used to recommend alternative strategies, such as:

“One school choice strategy is to find a school you like that is undersubscribed and put it as a top choice, OR, find a school that you like that is popular and put it as a first choice and find a school that is less popular for a ‘safe’ second choice.”

These strategies are fairly common among experimental subjects in controlled laboratory experiments with immediate acceptance (Chen and Sönmez, 2006),¹⁴ and

¹⁴See Hakimov and Kübler (2021) for a recent survey on experiments in centralized school choice and college admissions.

were the major reason for abandoning such admissions procedures in Boston and in the UK (Abdulkadiroğlu and Sönmez, 2003; Abdulkadiroğlu et al., 2005a,b; Pathak and Sönmez, 2008, 2013).

Mechanism design theory provides a foundation to study strategic incentives. Formally, a mechanism consists of a message space for each strategic agent and determines the outcome as a function of messages sent by the agents. This induces a game among agents, the outcome of which can be predicted by various equilibrium notions. In most equilibrium concepts, how much each agent knows about other agents (i.e. the information structure) determines whether a proposed equilibrium exist. Consequently, solution concepts are dictated by the information structure, as well as policy goals, such as strategy-proofness, which will be defined formally below.

A mechanism can have a fairly general message space. A mechanism is said to be direct if the message space for each agent consists of the parameters that determine the agent’s preference. In school choice, this is typically the set of all possible preference relations for a student. In practice, however, it is also part of the information that districts collect from parents but cannot verify. All other information, such as siblings’ enrollment status, is either readily available or can be verified. Therefore, direct mechanisms are a natural choice to study the school choice problem. In addition, the celebrated revelation principle states that there is no loss of generality in restricting attention to direct mechanisms; any outcome of any mechanism can also be obtained by a direct mechanism via the equilibrium solution of choice (Gibbard, 1973; Myerson, 1979).

A direct matching mechanism, or simply a mechanism, φ is a function that maps every pair (P_N, \succ_S) to a matching μ . For student i and school s , let $\varphi_i(P_N, \succ_S)$ be student i ’s match and $\varphi_s(P_N, \succ_S)$ be the set of students assigned to school s . Since a mechanism determines the outcome as a function of preferences, it induces a “preference revelation game” among students and strategic schools. Dominant strategy incentive compatibility ensures that each strategic agent finds reporting true preferences to the mechanism as best strategy regardless of what the agent knows about the game and regardless of how other agents act in the game. Formally, a matching mechanism φ is dominant strategy incentive compatible for student $i \in N$ if, for every profile of student preferences and school rankings (P_N, \succ_S) and every alternative preference relation P'_i for student i ,

$$\varphi_i(P_N, \succ_S) R_i \varphi_i(P'_i, P_{-i}, \succ_S),$$

where $P_{-i} = (P_j)_{j \in N \setminus \{i\}}$ is the preference profile of all students in P_N except i .

It is worth to take a close look at this solution concept. $\varphi_i(P_N, \succ_S)$ indicates that student i ’s assignment is a function of her submitted preference list, other students’ preference lists and school rankings, as well as school capacities, which are

suppressed here for simplicity. To compute her assignment under each alternative choice list, student i needs to know what preference list other students submit, P_{-i} , as well as the full list of school rankings, \succsim_S . In addition, if other students and some schools act strategically, these lists may not be the true lists. So, student i needs to predict what other strategic agents would do in equilibrium. Dominant-strategy incentive-compatibility frees the student from the burden of needing to know this immense amount of information. In particular, she does not need to know anything about other students’ preferences or a school’s rankings, or how strategically they act. By submitting her true preferences, she always achieves the best outcome she can ever hope to achieve by submitting a different preference relation.

A matching mechanism is strategy-proof if it is dominant strategy incentive compatible for all strategic participants, which includes all students in N and all strategic schools in S . The strategy-proofness property becomes less demanding as the set of strategic agents gets smaller. For example, no strategy-proof and stable mechanism exists when schools are strategic (Roth, 1982). In contrast, a strategy-proof solution exists when the set contains only students (Dubins and Freedman, 1981; Roth, 1982). Therefore, in this chapter, the notion of strategy-proofness will be defined for students only.¹⁵

Strategic behavior and violations of strategy-proofness are observed in real-life applications and in laboratory experiments. Families often find it “unsafe” to report their true preferences as they believe that being honest may hurt them in the matching process. As discussed above, this was a major concern of families in both Boston and New York in the early 2000s. Potential gains from strategically misrepresenting preferences are also reported in the Wake County Public School System in North Carolina. Dur et al. (2018a) show that students who systematically avoid applying to over-subscribed schools are 10 percentage points more likely to be assigned to one of their preferred schools. Similar evidence have been found in other studies, see, e.g., Agarwal and Somaini (2018), Burgess et al. (2015), Fack et al. (2019), Hastings et al. (2009), and He (2017).

Strategy-proofness is a demanding concept both in theory and practice. This motivates an interest for matching mechanisms that are “least manipulable” according to some criteria. There is a growing literature that investigates this problem in various contexts (see, e.g., Andersson et al., 2013; Azevedo and Budish, 2019; Maus et al., 2007; Kelly, 1988). For more recent work on school choice, see Bonkougou and Nesterov (2021), Chen and Kesten (2017), Chen et al. (2017), Dur et al. (2021), and Pathak and Sönmez (2013).

¹⁵Note that schools may gain both from misrepresenting their preferences over the students (if possible) as well as their capacities. See, e.g., Abdulkadiroğlu et al. (2005a,b), Kesten (2012), Kojima (2006), Konishi and Ünver (2006), or Sönmez (1997).

4 Matching Algorithms

Section 3 defines a mechanism as a function that produces a matching for every profile of student preferences and schools rankings. That function is also referred to as an algorithm. The chapter refers to it as mechanism when incentives matter, otherwise it refers to it as algorithm. This section will introduce matching algorithms that are not only central to the theory, but also commonly adapted all over the world. Often, stylized examples will be used to compare the algorithms. These simple examples will help highlighting features of and trade-offs among algorithms. Sections 5 and 7 then formalize and generalize these observations via the axioms introduced in Section 3.

Recall that school rankings may involve ties. A method to break ties (whenever they occur) is needed when an algorithm works with strict school rankings. Following Abdulkadiroğlu et al. (2009), a tie-breaker is defined as a bijection $t : N \rightarrow \mathbb{N}$ that breaks ties at school $s \in S$ by associating \succsim_s with a strict ranking \succ_s^t as follows: $i \succ_s^t j$ if and only if $i \succ_s j$, or i and j have the same ranking¹⁶ and $t(i) < t(j)$. A single tie-breaker uses the same tie-breaker at each school in S , whereas a multiple tie-breaker may use a different tie-breaker at each school in S . In the latter case, the tie-breakers are given by $\tau = (t_{s_1}, \dots, t_{s_{|S|}})$. This section describes the algorithms for strict school rankings possibly after tie-breaking. Section 6 discusses the welfare consequences of tie breaking in more detail.

4.1 The Deferred Acceptance Algorithm

The deferred acceptance algorithm (DA) was developed in the early 1950s by a consortium of medical schools to solve their problem of matching medical interns to residency programs. Later it was formulated independently by Gale and Shapley (1962), which established the literature on the theory of matching. Balinski and Sönmez (1999) also discovered that the deferred acceptance algorithm had been independently developed for college admissions in Turkey. The algorithm has a long history in entry level labor markets in medicine and law (see Roth, 2008, for a historical perspective). It is more broadly applied in assigning students to schools at all levels across the world. Versions of DA have been adopted by major school districts in the US to assign pupils to public K-12 schools, including Boston Public Schools (MA), Chicago Public Schools (IL), Denver Public Schools (CO), Indianapolis Public Schools (IN), Newark Public Schools (NJ), Recovery School District, New Orleans (LA), New York City High Schools (NY), Tulsa Public Schools (OK), and Washington DC Public Schools. It is used in many countries at different levels of

¹⁶i.e. $i \succsim_s j$ and $j \succsim_s i$

the education systems, for example, in Chile, Finland, France, Ghana, Hungary, Ireland, Norway, Romania, Spain, Taiwan, Tunisia, Turkey, and Sweden (Fack et al., 2019).¹⁷

Given a school choice problem (N, S, q, P, \succ) , the deferred acceptance algorithm finds the final matching as follows:

Step 1. Every student i applies to her most preferred school according to her preferences \succ_i . Every school s considers the students applying to it, and rejects ineligible students and provisionally assigns its seats to the remaining applicants in the order of its ranking \succ_s . When all seats at s are provisionally assigned, the school rejects all the remaining students.

Step k . Every student i that is rejected in the previous step applies to her next preferred school in \succ_i . Every school s considers students that it has provisionally assigned a seat in the previous step and students that apply in this step. From this set, school s rejects ineligible students and provisionally assigns its seats to the remaining students in the order of its ranking \succ_s . When all seats at s are provisionally assigned, the school rejects the remaining students.

The algorithm terminates when no student is rejected. The provisional assignments in the last step are finalized. Students without a provisional assignment in the last step remain unassigned.

Example 2. Suppose that there are three students, 1, 2 and 3, and three schools, s_1 , s_2 and s_3 . Each school has one available seat. Student preferences and school rankings are given by

$$\begin{aligned}
 1: & s_2 P_1 s_1 P_1 s_3 \\
 2: & s_1 P_2 s_2 P_2 s_3 \\
 3: & s_1 P_3 s_2 P_3 s_3 \\
 s_1 : & 1 \succ_{s_1} 3 \succ_{s_1} 2 \\
 s_2 : & 3 \succ_{s_2} 1 \succ_{s_2} 2 \\
 s_3 : & 1 \succ_{s_3} 2 \succ_{s_3} 3
 \end{aligned}$$

The deferred acceptance algorithm proceeds as follows. In Step 1, students 1, 2 and 3 apply to schools s_2 , s_1 and s_1 , respectively. Because 1 is the only student

¹⁷The Matching in Practice Network describes several of the European systems in detail. See <https://www.matching-in-practice.eu>.

who applies to s_1 , 1 is provisionally assigned the single seat at s_1 . Both 2 and 3 apply to s_1 . Since a s_1 has single seat and $3 \succ_{s_1} 2$, s_2 provisionally assigns its single seat to 3 and rejects 2. In the second step, 2, the only student that was rejected in Step 1, applies to her next most preferred school, s_2 . In this step, s_2 considers 1, who was provisionally assigned to the school in Step 1, and student 2. Since the school has a single seat and $1 \succ_{s_2} 2$, s_2 provisionally assigns student 1 and rejects student 2. In third step, 2 applies to her next most preferred school, s_3 . Because no student is provisionally assigned to s_3 and the school has unfilled seats, student 2 is provisionally assigned to s_3 . There are no more rejections. The algorithm terminates and the provisional assignments are finalized. The matching is given by

$$\mu = ((1, s_2), (2, s_3), (3, s_1)).$$

The roles of students and schools may be swapped. In that version of the algorithm, schools make offers to students, each student keeps the best among all offers she receives and rejects the remaining offers. That yields

$$\mu' = ((1, s_1), (2, s_3), (3, s_2)).$$

Note that all students prefer μ to μ' . Students 1 and 3 strictly prefer μ to μ' and 2 is indifferent between the two. This observation holds generally and will be discussed in Section 5. Therefore, the former version of the deferred acceptance algorithm is called “student-optimal” and the latter is called “school-optimal.” \square

The deferred acceptance algorithm can be used as a constructive proof to show that the set of stable matchings is non-empty.

Theorem 1. (*Gale and Shapley, 1962*) *The deferred acceptance algorithm converges to a stable matching in a finite number of steps.*

Proof. Convergence follows directly from the observation that there is a finite number of schools and each student applies to each school at most once. To prove stability, suppose that the matching μ is the outcome of the deferred acceptance algorithm for problem (P_N, \succ_S) . Recall that a matching is stable if it is individually rational and it cannot be blocked by any pair (i, s) . Individually rationality of μ follows directly from the construction of the algorithm because no student proposes to an unacceptable school and because schools reject ineligible students.

Next, suppose that there exists a blocking pair (i, s) at matching μ . Then, $sP_i\mu(i)$, $iP_s\mu(s)$ and either (a) $|\mu(s)| < q_s$ or (b) $i \succ_s j$ for some $j \in \mu(s)$. Because $sP_i\mu(i)$, student i must have applied to and been rejected by s at some Step k of the algorithm. Since i is eligible but rejected by s , all the seats at s must be provisionally

assigned at Step k . Then, by construction, there is no available seat at school s at any Step $k' > k$. Thus, case (a) cannot hold. From the fact that i was rejected by s at Step k , it follows that $j \succ_s i$ for every student j that was provisionally assigned to school s at Step k . Such j can be rejected by school s at some Step $k' > k$ only if some other student l applies school s at that step, where $l \succ_s j$. But this implies $l \succ_s i$. So, case (b) cannot prevail, either. Thus, there exists no blocking pair. \square

The algorithm described above is referred to in the literature as “student-proposing” deferred acceptance algorithm. The version in which schools extend offers is known as “school-proposing” deferred acceptance algorithm. This chapter mainly concerns the former, and therefore will refer to the student-proposing deferred acceptance simply as deferred acceptance.

4.2 The Immediate Acceptance Algorithm

The immediate acceptance algorithm, also known as the Boston mechanism, was invented independently by practitioners around the world. It is popular in school choice and college admissions in, e.g., China, Germany, Spain, and the US (see, e.g., Abdulkadiroğlu and Sönmez, 2003; Alcalde, 1996; Kojima and Ünver, 2014; Roth, 1991). The algorithm is easy to explain and implement. It processes students at schools in the order of choice, therefore it places more students to their most preferred schools than other mechanisms. This may explain its easy discovery and appeal in field.

The immediate acceptance algorithm produces a matching as follows:

Step 1. Every student i applies to her most preferred acceptable school in \succ_i . Every school s permanently assigns its seats to its eligible applicants in the order of its ranking \succ_s . It rejects ineligible applicants. When all seats are permanently assigned, the school rejects the remaining applicants.

Step k . Every student i who was rejected in the previous step applies to her k^{th} most preferred acceptable school. Every school s with available seats permanently assigns its remaining seats to its new eligible applicants in the order of \succ_s . It rejects ineligible applicants. When all seats are permanently assigned, the school rejects all the remaining applicants.

The algorithm terminates when no student application is rejected. The students that are rejected by all of their choice schools remain unassigned. Note the difference between the deferred acceptance algorithm where the students are provisionally

assigned to schools in all steps until the algorithm converges. In the immediate acceptance algorithm, all assignments in all steps are permanent.

Example 3. Suppose that there are three students, 1, 2 and 3, and three schools, s_1 , s_2 and s_3 . Student preferences and school rankings are given by

$$\begin{aligned}
1 &: s_2 P_1 s_1 P_1 s_3 \\
2 &: s_1 P_2 s_2 P_2 s_3 \\
3 &: s_1 P_3 s_2 P_3 s_3 \\
s_1 &: 1 \succ_{s_1} 3 \succ_{s_1} 2 \\
s_2 &: 2 \succ_{s_2} 1 \succ_{s_2} 3 \\
s_3 &: 2 \succ_{s_3} 1 \succ_{s_3} 3
\end{aligned}$$

In Step 1 of the algorithm, students 1, 2 and 3 propose to schools s_2 , s_1 and s_1 , respectively. Because 1 is the only applicant at s_2 , 1 is permanently assigned the single seat at s_2 . 2 and 3 apply to s_1 in this step. Since $3 \succ_{s_1} 2$, student 3 is assigned the single seat at s_1 and 2 is rejected by the school. In the second step, 2 applies to her second most preferred school s_2 . There are no remaining seats at s_2 , so 2 is rejected by s_2 . Then 2 applies to her third most preferred choice s_3 in Step 3. She is permanently assigned the single seat at s_3 . The outcome of the immediate acceptance algorithm is given by

$$\mu^{\text{immediate}} = ((1, s_2), (2, s_3), (3, s_1)).$$

The unique stable matching, and therefore the outcome of both versions of the deferred acceptance algorithm, is

$$\mu = ((1, s_1), (2, s_2), (3, s_3)),$$

illustrating that the deferred acceptance and the immediate acceptance algorithm need not recommend the same matching. \square

Theorem 2. (*Abdulkadiroğlu and Sönmez, 2003*) *The immediate acceptance algorithm converges in a finite number of steps to a matching that is Pareto efficient with respect to P_N .*

Proof. Convergence follows directly from the observation that there is a finite number of schools and each student applies to each school at most once.

Let matching μ be the outcome of the immediate acceptance algorithm for problem (P_N, \succ_S) . To prove Pareto efficiency by contradiction, suppose that the matching μ' Pareto dominates matching μ , and let $I = \{i \in N : \mu'(i) P_i \mu(i)\}$ be the set of

students that strictly prefer μ' to μ . Note first that every student $i \in I$ must have been rejected by school $\mu'(i)$ at some step of the immediate acceptance algorithm. Suppose that $i \in I$ is one of the students that is rejected by $\mu'(i)$ earliest, say in Step k , in the immediate acceptance algorithm. Let $s = \mu'(i)$. Next, note that all seats at school s must be assigned in Step k of the immediate acceptance algorithm to other students. Then, for matching μ' to assign i to s , there must exist some student $j \in \mu(s) \setminus \mu'(s)$. That is, j is assigned s by the immediate acceptance algorithm and she is assigned a different school at matching μ' . Since μ' Pareto dominates $\mu(j) \neq \mu'(j)$, it must be that j prefers $\mu'(j)$ to $\mu(j)$. Since assignments are permanent in the immediate acceptance algorithm, there must exist one such student j such that j is rejected by $\mu'(j)$ earlier than student i is rejected by $\mu'(i)$ by the immediate acceptance algorithm. Otherwise there would be available seats for i and i would be assigned $\mu'(i)$ by the the algorithm. This contradicts that i is one of the students that was rejected by $\mu'(i)$ earliest. Hence, matching μ is Pareto efficient with respect to P_N . \square

Pareto efficiency of the immediate acceptance algorithm is with respect to submitted preferences, as is the stability of the deferred acceptance algorithm. However, families may have incentives to misrepresent their preferences and submit a choice list that is different than their actual preference list. This issue is examined in more detail in Section 5.2. Example 3 shows one way the algorithm may punish families that submit their preferences truthfully. Notice that student 2 applies to school s_2 in Step 2 of the algorithm even if no seats are available at the school. By the time she applies to her third choice in Step 3, all the seats at her third choice might have been assigned in Step 2. In general, this means that students face risk of being assigned very low in their choice list simply because they may waste too many steps of the algorithm to apply to schools with no available seats (this point is later illustrated in Example 6).

To remedy this problem, algorithms may allow students to skip applying to schools, or report one school that they wish to “neutraliz”, or “secure” (Dur, 2019; Harless, 2017; Mennle and Seuken, 2015; Miralles, 2008; Decerf, 2021; Dur et al., 2018b).¹⁸ However, the popularity of the immediate acceptance algorithm has decreased in the last decade due to experimental evidence, theoretical findings, and empirical

¹⁸Many countries use alternative versions of the immediate acceptance algorithm or hybrids between the immediate acceptance and the deferred acceptance. For example, the immediate acceptance algorithm used to be the only mechanism to assign students to high schools and colleges in China, but some provinces have recently adopted the so-called parallel mechanism. Chen and Kesten (2017) propose an algorithm, called the application-rejection mechanism, which includes the immediate acceptance algorithm, the Chinese parallel mechanisms, and the deferred acceptance algorithm as special cases. Similarly, Dur et al. (2021) investigated the nationwide mechanism that was implemented recently in Taiwan for high school assignments. This mechanism is a hybrid of the immediate acceptance and the deferred acceptance algorithms.

observations in the field (e.g., Abdulkadiroğlu et al., 2005b,a, 2009; Abdulkadiroğlu and Sönmez, 2003; Chen and Sönmez, 2006). For example, the immediate acceptance algorithm has been abandoned in cities like Boston and Chicago, and it is now forbidden by law in the UK (Pathak and Sönmez, 2008, 2013).

4.3 The Top Trading Cycles Algorithm

The top trading cycles algorithm was first investigated by Shapley and Scarf (1974) in the context of a house allocation problem in which each economic agent owns a house and would like to swap it for a more preferred option.¹⁹ It was generalized to be applicable to school choice by Abdulkadiroğlu and Sönmez (2003). In the context of school choice, stability can be viewed as a constraint on the problem of maximizing student welfare in admissions and, therefore, may preclude Pareto efficiency (Example 1). The top trading cycles algorithm, by comparison, always produces a Pareto efficient matching. Because of the trade-off between stability and efficiency, the top trading algorithm has become a plausible alternative for design of admissions processes. The top trading cycles algorithm was adopted by the Recovery School District in 2012, and it has also been used by the Swedish municipality Oxelösund. It was one of the algorithms in consideration and also the recommendation of a task force during the redesign of admissions in Boston Public Schools in 2003-2005.

The top trading cycles algorithm works as follows:

Step 0. Every student and every school are initially available.

Step k. An available student becomes unavailable when she is assigned, or if none of the available schools at which she is eligible are acceptable for her. In the latter case, she remains unassigned. An available school becomes unavailable when all of its seats are assigned, or if none of the available students that find the school acceptable are eligible at the school. In the latter case, the remaining seats at the school remain unfilled.

Every available student i points to her most preferred acceptable school among all available ones. Every available school s points to the eligible student that is highest ranked in \succ_s among all available students.

A cycle is an ordered list of students and schools $(i_1, s_1, i_2, s_2, \dots, i_n, s_n)$ such that, for each $k = 1, \dots, n$, student i_k points to school s_k and school s_k points to student i_{k+1} , where $n + 1$ is replaced by 1.

¹⁹Shapley and Scarf attributed the algorithm to David Gale, so it is sometimes referred to as “Gale’s top trading cycles algorithm.”

For each cycle, assign each student in the cycle to a seat at the school that she points to.

The algorithm terminates if there are no available students or no available schools.

The following example demonstrates the algorithm.

Example 4. Consider now the same setting as in Example 3 with student preferences and school rankings as given by

$$\begin{aligned}
 1 &: s_2 P_1 s_1 P_1 s_3 \\
 2 &: s_1 P_2 s_2 P_2 s_3 \\
 3 &: s_1 P_3 s_2 P_3 s_3 \\
 s_1 &: 1 \succ_{s_1} 3 \succ_{s_1} 2 \\
 s_2 &: 2 \succ_{s_2} 1 \succ_{s_2} 3 \\
 s_3 &: 2 \succ_{s_3} 1 \succ_{s_3} 3
 \end{aligned}$$

In this example, every school is acceptable for every student, and every student is eligible at every school. In the beginning of Step 1, all students and school are available.

Student 1 points to s_2 , 2 and 3 point to s_1 . School s_1 points to 1, s_2 and s_3 point to 2. There is a cycle in which 1 points to s_2 , s_2 points to 2, 2 points to s_1 and s_1 points to 1. Student 1 is assigned the single seat at her first choice s_2 and 2 is assigned the single seat at his first choice s_1 . Students 1 and 2, as well as schools s_1 and s_2 , become unavailable. In Step 2, the only available student is 3 and the only available school is s_3 . They point to each other, forming a cycle, and student 3 is assigned her last choice s_3 . The outcome of the top trading cycles algorithm is

$$\mu^{\text{top}} = ((1, s_2), (2, s_1), (3, s_3)).$$

Note also that for the preferences given in Example 3, this solution is different from both the outcome of the deferred acceptance algorithm and the immediate acceptance algorithm. Hence, even in the simplest possible case when there are only three students and three schools, these three commonly used algorithms can produce three distinct matchings. \square

With finite numbers of students and schools, there exists a cycle in every step. Assigning students to schools reduces the set of available students and schools. Therefore, the top trading cycles algorithm converges in a finite number of steps. What is a bit more surprising is that the outcome is always a Pareto efficient matching.

Theorem 3. (*Abdulkadiroğlu and Sönmez, 2003*) *The top trading cycles algorithm converges in a finite number of steps to a Pareto efficient matching with respect to P_N .*

The proof of Pareto efficiency in Theorem 3 follows from a simple observation. No student and no school can be involved in multiple cycles simultaneously. Students that are assigned in Step 1 are assigned their first choices, so they cannot be made better off. A student that is assigned in a later step to a lower choice could not form a cycle at her more preferred choices earlier as her better choices already are fully assigned to other students. Therefore, it is not possible to assign her to a better choice without assigning one of those students to a school that they consider to be worse than their current match.

4.4 The Serial Dictatorship Mechanism

The serial dictatorship algorithm was first analyzed for queuing problems, such as assigning individuals to offices or public housing depending on their positions in a queue or waiting list (see, e.g., Abdulkadiroğlu and Sönmez, 1998; Hylland and Zeckhauser, 1979; Svensson, 1994; Zhou, 1990). It is easy to describe and implement the serial dictatorship algorithm even in a decentralized manner, individuals select, in the order of their position in the queue or waiting list, their most preferred object among the remaining ones. Serial dictatorship can be defined more formally in the context of school choice as follows.

Given a strict ordering of students,

Step k . The k -th student in the ordering is assigned a seat at her most preferred acceptable school among all schools with available seats, at which she is also eligible. If no such school exists, she remains unassigned.

The algorithm terminates when all students in the ordering are processed.

When it is implemented in centralized manner, the information requirement for the serial dictatorship is minimal. It only needs student preferences, and a queue or ordering of students, which can also be generated randomly. For that reason, The New York City Department of Education uses it to assign students to schools in the third and last round of their admissions process due to lack of time required for collecting school rankings at the end of the process.

The serial dictatorship algorithm can also be implemented via the deferred acceptance and top trading cycles algorithms.

Theorem 4. *If all school rankings are set to the ordering used in the serial dictatorship, then the outcomes of deferred acceptance, top trading cycles and serial dictatorship algorithms are the same for every P_N .*

The proof of this result follows from the following observations. Suppose that every student only ranks acceptable schools at which she is eligible. The first student in the ordering is assigned a seat at her most preferred school with an available seat. Under deferred acceptance, she applies to the same school. Since she has the highest ranking at the school, she is provisionally assigned a seat at the school at the first step and never rejected by the school later. Similarly, under top trading cycles, she points to the same school, which points back to her since she has the highest ranking. So she is assigned a seat at her top choice. In a serial dictatorship, the next student is assigned at their most preferred acceptable school with an available seat. So the result follows by induction.

Since the top trading cycles algorithm is Pareto efficient, the efficiency of serial dictatorship follows immediately.

Corollary 1. *The serial dictatorship mechanism is Pareto efficient with respect to every P_N .*

Notice that the deferred acceptance algorithm also becomes Pareto efficient when schools share the same ranking.

5 Results with Strict Preferences and Priorities

Most of matching theory has been developed under the assumption that agents on both sides of the market have strict preference rankings. That assumption fits earlier applications in entry-level labor markets (Roth, 1984a, 1991). It also fits in the context of school choice when school rankings are strict. Cities in the US, e.g. Boston, Chicago, and New York²⁰, and countries around the world use entrance exams for admissions to selective middle and high schools. Entrance examination is common practice in college admissions as well. This section focuses on results in the case when school rankings are strict.

²⁰New York City plans to eliminate entrance examination for admissions to their selective schools for concerns over segregation, see:
<https://www.nytimes.com/2020/12/18/nyregion/nyc-schools-admissions-segregation.html>.

5.1 Stable Matchings and The Deferred Acceptance Algorithm

A stable matching utilizes all seats at over-subscribed schools. However, it may leave some seats empty at under-demanded schools. As demonstrated in Example 2, there may be multiple stable matchings. To select among stable matchings, a natural question regarding seat utilization is whether some stable matchings can fill more seats at under-demanded schools. This question was originated in medical residency match programs, concerning hospitals in rural areas that had difficulties attracting interns. The famous rural hospitals theorem (McVitie and Wilson, 1970; Roth, 1986), when applied to school choice, states that any school that does not fill all its seats at some stable matching will be assigned precisely the same set of students at every stable matching.

Theorem 1 establishes that the set of stable matchings is nonempty. Equally important, the set of stable matchings is a partially ordered set, known as a lattice (Knuth, 1976). In a lattice, every pair of elements has a unique supremum and a unique infimum.²¹ In the set of stable matchings, the supremum and infimum are found by comparing two matchings, agent by agent, according to their preferences. Consider two arbitrary stable matchings μ and μ' . Construct a supremum μ'' by assigning each student i to the better of $\mu(i)$ and $\mu'(i)$, i.e.:

$$\mu''(i) = \begin{cases} \mu(i) & \text{if } \mu(i) \succ_i \mu'(i) \\ \mu'(i) & \text{otherwise} \end{cases}$$

The supremum, μ'' , is a stable matching (Roth, 1984b, 1985b).²² This is because every student weakly prefers μ'' to both μ and μ' . This implies that there exists a stable matching that every student weakly prefers to any other stable matching. The argument is inductive. More precisely, since there are finitely many students and schools, there are also finitely many stable matchings. Now, take any two stable matchings μ and μ' and find the supremum μ'' of the two. Then, select another matching $\tilde{\mu}$ that has not been selected before, find the supremum of μ'' and $\tilde{\mu}$. Repeat this process until all of the stable matchings are exhausted. The resulting matching is referred to as the student-optimal stable matching. The (student-proposing) deferred acceptance algorithm finds that solution.

Theorem 5. (*Gale and Shapley, 1962*) *The deferred acceptance algorithm selects the student-optimal stable matching for every (P_N, \succ_S) .*

²¹These are also called join and meet, respectively.

²²This follows from a result in Knuth (1976) referred to as the Pointing Lemma (attributed to John Conway). See also Roth and Sotomayor (1990).

In addition to being student-optimal stable, the deferred acceptance algorithm also induces straightforward incentives for students.

Theorem 6. *(Dubins and Freedman, 1981; Roth, 1982) The deferred acceptance algorithm is strategy-proof for students.*

Note that this result assumes that students are allowed to report their preferences over all schools. In reality, districts may impose limits on the number of schools that a student can rank in their application file. For example, in New York City, student can rank up to twelve programs. The Matching in Practice Network²³ reports that students in elementary schools in, for example, Estonia, Germany, Italy, Spain and Sweden can only rank a restricted number of schools (typically only three). In such cases, families and students are forced to choose strategically among schools they deem to be likely to get in to, which may not necessarily contain their most preferred schools. However, once a student determines the limited number of schools to apply for, ranking these schools in true preference order weakly dominates any other ranking of the same set of schools under deferred acceptance (Haeringer and Klijn, 2009; Romero-Medina, 1998).

A school-optimal stable matching also exists, one in which every school is assigned weakly higher ranked students in comparison to any other stable matching. The proof of this result is similar to the proof the existence of student-optimal stable matching. Instead of assigning each student i to the better of the two matchings, one needs to assign each school s to the better of two sets of students $\mu(s)$ and $\mu'(s)$. The school-proposing deferred acceptance algorithm then finds the school-optimal stable matchings.

The student-optimal stable matching is the most preferred stable matching for students. It is also the least preferred stable matching for the schools in the sense that it assigns weakly lower ranked students to each school compared to any other stable matching (see Roth and Sotomayor, 1990). This was also illustrated in Example 2. The converse is true for the school-optimal stable matching. When school rankings reflect school preferences, this represents a trade-off, which normally is resolved in favor of students. Although that choice can be a result of aiming at higher student welfare, there is also a theoretical justification for it, which is connected to incentives to manipulate preferences.

Theorem 7. *(Roth, 1982) When schools are strategic and can manipulate their rankings, there is no stable and strategy-proof mechanism.*

This result can be proved via the following simple counter example.

²³See <https://www.matching-in-practice.eu/elementary-schools>.

Example 5. Consider Example 2. There are only two stable matchings, μ^{student} and μ^{school} . Therefore, any stable matching mechanism must select one of these two matchings for this problem. First, suppose that a stable matching mechanism selects the former. If school s_1 reports the following preference relation,

$$1 \succ'_{s_1} 2 \succ'_{s_1} 3,$$

then μ^{school} is the only stable matching of the problem $(P_N, \succ'_{s_1}, \succ_{s_2}, \succ_{s_2})$, so it must be selected by the mechanism.²⁴ Notice that, according to \succ_{s_1} , school s_1 is matched to her most preferred student at this matching and the least preferred student at μ^{student} , so school s_1 is better off by reporting its rankings as \succ'_{s_1} when its actual ranking is \succ_{s_1} . Similarly, if the mechanism instead selects matching μ^{school} , then student 1 can successfully manipulate the mechanism by reporting \hat{P}_1 :

$$s_2 \hat{P}_1 s_3 \hat{P}_1 s_1.$$

For the reported preferences $(\hat{P}_1, P_{-1}, \succ_S)$, matching μ^{student} is the only stable matching, so it must be selected by the mechanism. The manipulation is successful because, according to \hat{P}_1 , student 1 is matched to her most preferred school at this matching, but only the second most preferred school at matching μ^{school} .

This demonstrates that either a school or a student can manipulate the outcome of any stable matching. Consequently, there is no stable and strategy-proof mechanism when both students and schools are strategic. \square

When an impossibility result is identified in economic theory, a common next step is to relax requirements and assumptions to investigate whether the impossibility somehow can be escaped. Recall that strategy-proofness requires dominant strategy incentive compatibility for students and schools. Therefore, a natural relaxation is to impose incentive compatibility either on students or on schools but not on both. In practice, schools offer more than one seat, which leads to the following negative result.

Theorem 8. (*Roth, 1984b*) *When schools are strategic and can manipulate their rankings and there are schools with multiple seats, there exists no stable mechanism where truth telling is a weakly dominant strategy for all schools.*

This negative result is a consequence of schools being assigned multiple students. In contrast, each student can be assigned at most one school. That restriction recovers incentive compatibility for students (Theorem 6).

²⁴One way to verify this is to check that both the student-optimal and the school optimal matchings yield the same outcome for problem $(P_N, \succ'_{s_1}, \succ_{s_2}, \succ_{s_2})$. The result then follows directly since the set of stable matchings has the structure of a lattice.

A similar positive result holds for the schools if, unrealistically, each school has a single seat. This is referred to as the marriage problem in the literature, where men and women are matched in heterosexual pairs in a one-to-one fashion.

The following result strengthens the negative findings in Theorem 7 by introducing another problematic trade-off among individual rationality, Pareto efficiency and strategy-proofness.

Theorem 9. *(Alcalde and Barbera, 1994) When schools are strategic and can manipulate their rankings, there is no individually rational, Pareto efficient and strategy-proof mechanism.*

Alcalde and Barbera (1994) also identify a non-trivial restriction on the preference domain, called top-dominance, under which there exist strategy-proof and stable mechanisms.

Example 1 demonstrates that stability and Pareto efficiency are not compatible, either. Given the impossibility of meeting stability and Pareto efficiency simultaneously, the student-optimal stable matching can be viewed as a constrained optimal solution for students subject to the stability constraint. Similarly, one can search for a Pareto efficient mechanism that minimizes priority violations (Abdulkadiroğlu et al., 2020a). Minimizing priority violations can be defined in various ways, all yielding similar results (Doğan and Ehlers, 2021); however, finding one is computationally infeasible (Abdulkadiroğlu and Grigoryan, 2021a). Alternatively, one can look for a Pareto efficient mechanism that always selects a stable matching whenever a Pareto efficient and stable matching exists. The following results shed some light on this approach.

Theorem 10. *(Kesten, 2010) No Pareto efficient and strategy-proof mechanism Pareto dominates the student-optimal matching.*

Kesten (2010) provides an algorithm that Pareto dominates the deferred acceptance algorithm (see also Ehlers and Morrill, 2020). Ergin (2002) identifies an acyclicity condition that is sufficient for the deferred acceptance algorithm to be Pareto efficient.

5.2 The Immediate Acceptance Algorithm

The immediate acceptance algorithm is Pareto efficient with respect to reported preferences (Theorem 2). However, it also induces incentives to strategically misreport preferences. As mentioned in Section 3.3, it is also easy for families to discover

simple preference manipulations. In particular, the algorithm processes applications in the order of student preference rankings. In the first step of the algorithm, the first choices of the students are considered. The highest ranked students are permanently assigned their most preferred schools and the remaining students are rejected. Because popular schools are rapidly filled up, a student risks lowering her chances if she ranks a popular school as her first choice, at which she does not have great odds. The student may therefore find it optimal to not rank her first choice at all in her application. The following example illustrates this point.

Example 6. Consider Example 3. If students report their preferences truthfully, the outcome of the immediate acceptance algorithm is given by the matching:

$$\mu^{\text{immediate}} = ((1, s_2), (2, s_3), (3, s_1)).$$

Note here that students 1 and 3 are assigned their most preferred schools according to their true preferences, so it is impossible for them to get a better match by misrepresenting. Student 2, however, is assigned to her least preferred school. Note that 2's top-ranked school, s_1 , is over-subscribed and she has no chance of being assigned s_1 in the first step. When she is considered at her second choice, s_2 , in the second step, there are no available seats at s_2 . If 2 reports \hat{P}_2 instead of P_2 , where,

$$s_2 \hat{P}_2 s_1 \hat{P}_2 s_3,$$

students 1 and 2 are both considered by school s_2 in Step 1 of the algorithm. Because $2 \succ_{s_2} 1$, 2 is permanently assigned the single seat at s_2 , which is her second most preferred school according to her true preferences P_2 . The student can, consequently, gain by misrepresenting her preferences. This also proves that the immediate acceptance algorithm is not strategy-proof. \square

The fact that the immediate acceptance algorithm is not strategy-proof makes it somewhat more complicated to analyze as its equilibria depend on the information structure (see, e.g., Abdulkadiroğlu et al., 2011b; Ergin and Sönmez, 2006; Pathak and Sönmez, 2008).

Example 3 also demonstrates that the outcome of the immediate acceptance algorithm need not be stable with respect to reported preferences. In particular, student 2 has higher priority at s_2 than 1, but 1 is assigned the single seat at school s_2 by virtue of ranking it as first choice, whereas 2 strictly prefers s_2 to her current matching, s_1 . So even if the immediate acceptance algorithm is intuitive to explain and understand, it is both unstable and easy to manipulate. For further readings, see, for example, Afacan (2013), Chen (2016), Doğan and Klaus (2018), and Kojima and Ünver (2014).

5.3 Top Trading Cycles Algorithm

Top trading cycles was introduced and extensively studied in the context of housing markets (see, e.g., Ma, 1994; Roth, 1982; Roth and Postlewaite, 1977). Abdulkadiroğlu and Sönmez (2003) builds on it to solve the school choice problem in an efficient and strategy-proof manner. Both top trading cycles and deferred acceptance mechanisms are strategy-proof, but they deviate in one important aspect: the top trading cycles algorithm is Pareto efficient and fails to be stable, whereas the deferred acceptance algorithm is stable but not Pareto efficient.

Theorem 11. (*Abdulkadiroğlu and Sönmez, 2003*) *The top trading cycles algorithm is Pareto efficient and strategy-proof.*

This result immediately carries to the serial dictatorship mechanism since it can be implemented via top trading cycles.

Corollary 2. *The serial dictatorship mechanism is Pareto efficient and strategy-proof.*

Note that Pareto efficiency of top trading cycles does not imply that it Pareto dominates the deferred acceptance algorithm even if the latter is not Pareto efficient. Since the top trading cycles algorithm is also strategy-proof, this conclusion follow immediately from Theorem 10, i.e., no Pareto efficient and strategy-proof mechanism Pareto dominates the student-optimal stable matching in the school choice problem. The following example demonstrates that the top trading cycles algorithm is not stable.

Example 7. Consider again Example 3 and recall that the outcome of the top trading cycles algorithm is given by:

$$\mu^{\text{top}} = ((1, s_2), (2, s_1), (3, s_3)),$$

and that student 3's preferences and the ranking of school s_1 are given by

$$\begin{aligned} 3 &: s_1 P_3 s_2 P_3 s_3 \\ s_1 &: 1 \succ_{s_1} 3 \succ_{s_1} 2 \end{aligned}$$

respectively. In this case, $(3, s_1)$ forms a blocking pair since 3 prefers s_1 to her current match $\mu^{\text{top}}(3) = s_3$ and s_1 ranks 3 higher than its current match $\mu^{\text{top}}(s_1) = 2$. Therefore, μ^{top} is cannot be stable. \square

The deferred acceptance algorithm is constrained-optimal since its outcome weakly Pareto dominates any other stable matching. In contrast, there are many Pareto

efficient and strategy-proof mechanisms, e.g., the serial dictatorship and the top trading cycles mechanisms. One can search for a Pareto efficient mechanism that minimizes blocking pairs or priority violations. Formally, a mechanism φ^1 has less priority violations than mechanism φ^2 if, for any given problem, every priority violation instance of φ^1 is also a priority violation instance of φ^2 . The following result then applies in a one-to-one matching setting.

Theorem 12. (*Abdulkadiroğlu et al., 2020b*) *When each school has a single seat, the top trading cycles algorithm has less priority violations than any Pareto efficient and strategy-proof mechanisms.*

As mentioned earlier, minimizing priority violations can be defined in alternative and equally intuitive ways. Doğan and Ehlers (2021) discuss several definitions in the literature. They show that those definitions satisfy certain axioms that guarantee the result above. Therefore, the result in Theorem 12 is robust to various definitions of minimizing priority violations.

Generalizing this result when schools have multiple seats is not trivial. Finding a Pareto efficient mechanism that minimizes priority violations is computationally infeasible (Abdulkadiroğlu and Grigoryan, 2021a). Abdulkadiroğlu et al. (2020b) compare matching mechanisms in terms of “average incidences” in such problems and obtain some domination results in terms of less priority violations for the school choice problem.

Top-trading cycles and deferred acceptance are compared extensively in the literature. The fact that both the deferred acceptance algorithm and the top trading cycles algorithm are strategy-proof, but the former is stable whereas the latter is not and the latter is Pareto efficient whereas the former is not, has been the source of debate in the literature related to choosing the right mechanism. The answer depends on the application and policy objectives. Although the deferred acceptance algorithm seems to dominate the field in school choice programs, there is a sizable literature that investigates and analyses various properties of the top trading cycles algorithm. See, for example, Abdulkadiroğlu et al. (2010, 2020b), Ergin and Sönmez (2006), Haeringer and Klijn (2009), Kesten (2010), Morrill (2013, 2015), Pápai (2000), Pathak and Sönmez (2008), and Pycia and Ünver (2017).

6 Results with Weak Preferences and Priorities

Unlike in earlier applications of matching theory to entry level labor markets, ties and indifferences in school rankings are a common feature in school choice programs.

When schools utilize priorities in admissions, they sort students into thick priority classes. For instance, when a school grants priorities based on whether siblings are enrolled at the school and whether student resides within walking distance of the school, multiple students may have one or both of these priorities.

Recall from Section 5 that when school rankings are strict, the set of stable matchings forms a lattice. Consequently, a unique student-optimal stable matching exists. This result does not extend when school rankings involve indifferences. Consider an example with two students, 1 and 2, and a single school, s , with one seat. Suppose that the school is acceptable for both students and the school ranks both students equally. In this case, there are two stable matchings, $\mu^1 = ((1, s), (2, 2))$ and $\mu^2 = ((1, 1), (2, s))$. Student 1 prefers μ^1 and student 2 prefers matching μ^2 . Therefore, $\{\mu^1, \mu^2\}$ does not form a lattice and each matching is student-optimal in the sense that there is no other stable matching that Pareto dominates either.

Indifferences in school rankings aggravate the trade-offs among stability, efficiency and incentives. As illustrated in Example 1, there may not exist any stable and Pareto efficient matching. However, when school ranking are strict, existence of a stable and Pareto efficient matching can be easily verified by checking whether the deferred acceptance algorithm produces a Pareto efficient matching. In contrast, when school rankings involve indifferences, verifying the existence of a stable and Pareto efficient matching, or even finding one, becomes computationally infeasible.

Theorem 13. *(Abdulkadiroğlu and Grigoryan, 2021a) When school rankings involve indifferences, deciding whether there is a stable and Pareto efficient matching is an NP-complete problem.*

As discussed in Section 4, the algorithms in their simplest forms are defined for strict school rankings. When school rankings involve ties and indifferences, a strict ranking is obtained by breaking ties. Despite added complications created by indifferences in school rankings, tie-breaking preserves certain properties of the algorithms. The following theorem follows from earlier results.

Theorem 14. *If ties in school rankings are broken by any tie-breaker τ , the following results holds:*

- (i) the deferred acceptance algorithm is stable and strategy-proof for the students,*
- (ii) the top trading cycles algorithm is Pareto efficient and strategy-proof for the students,*
- (iii) the immediate acceptance algorithm is Pareto efficient.*

Some properties that holds with strict school rankings, however, do not extend when school rankings involve indifferences. For example, Theorem 12 does not extend into the case of indifference in tie-breaking; hence, the top trading cycles algorithm is not guaranteed to minimize priority violations in the class of Pareto efficient and strategy-proof mechanisms even when each school has a single seat. The problem of tie-breaking in top trading cycles has been studied extensively by Alcalde-Unzua and Molis (2011), Aziz and de Keijzer (2012), Ehlers (2014), Ehlers and Erdil (2010), and Jaramillo and Manjunath (2012).

Indifferences in school rankings imply a fundamental conflict among Pareto efficiency, stability and strategy-proofness. In this case, the deferred acceptance algorithm does not guarantee a stable and Pareto efficient matching (whenever one exists) as demonstrated in the following example.

Example 8. Suppose that there are three students, 1, 2 and 3, and three schools, s_1 , s_2 and s_3 . Each school has one available seat. The strict preferences of the students are given by

$$1 : s_2 P_1 s_1 P_1 s_3$$

$$2 : s_1 P_2 s_2 P_2 s_3$$

$$3 : s_1 P_3 s_2 P_3 s_3$$

All students have the same priority at every school, so there are no inherit strict school rankings. In other words, each school ranks all three students equally. Suppose that each school draws its own lottery to break ties among students and that the school rankings after the lottery tie-breaking are given by

$$s_1 : 1 \succ_{s_1} 3 \succ_{s_1} 2$$

$$s_2 : 2 \succ_{s_2} 1 \succ_{s_2} 3$$

$$s_3 : 1 \succ_{s_3} 2 \succ_{s_3} 3$$

Notice that every non-wasteful matching is stable in this example, since schools do not have strict rankings. The outcome of the deferred acceptance algorithm is:

$$\mu = ((1, s_1), (2, s_2), (3, s_3)).$$

However, it is Pareto dominated by the following stable matching:

$$\mu' = ((1, s_2), (2, s_1), (3, s_3)).$$

Students 1 and 2 are matched to their first choices at matching μ' . Although 3 has a better ranking than 2 at school s_1 due to a better lottery draw, the pair $(3, s_1)$

does not form a blocking pair since 2 and 3 have the same priority at s_1 . Intuitively, random tie-breaking introduces artificial stability constraints. As a result, the deferred acceptance algorithm may not produce a student-optimal matching since it eliminates artificially created blocking pairs such as $(3, s_1)$. Note also that matching μ' is Pareto efficient. In this case, the deferred acceptance algorithm fails to find a stable and efficient matching as well. \square

Since the deferred acceptance algorithm is strategy-proof, the above result is related to a more fundamental problem.

Theorem 15. *(Erdil and Ergin, 2008) No strategy-proof mechanism is student-optimal stable.*

Failure to find a stable and Pareto efficient matching when one exists also relates to a similar fundamental problem.

Theorem 16. *(Abdulkadiroğlu and Grigoryan, 2021a) No strategy-proof mechanism finds a stable and Pareto efficient matching whenever one exists.*

Once ties are broken, the deferred acceptance algorithm becomes a constrained optimal compromise in the following sense:

Theorem 17. *(Abdulkadiroğlu et al., 2009) If ties in school rankings are broken by an arbitrary tie-breaker τ , then no strategy-proof mechanism Pareto dominates the deferred acceptance algorithm.*

This result generalizes Theorem 10 in two ways. First, it considers indifferences in school rankings. Second, it shows impossibility of Pareto improvement instead of impossibility of Pareto efficiency.

If a stable matching is not student-optimal, a Pareto dominant stable matching can be found in polynomial time (Erdil and Ergin, 2008). To see this, consider a stable matching μ . Student i desires school s if $s \succ_i \mu(i)$. Let $D_s(\mu)$ denote the set of students that are ranked highest according to \succ_s among students who desire school s . A stable improvement cycle is a list of students $(i_1, \dots, i_K, i_{K+1} = i_1)$, such that for every $k \in \{1, 2, \dots, K\}$,

- student i_k is assigned a school, i.e., $\mu(i_k) \in S$,
- student i_k desires school $s = \mu(i_{k+1})$,
- student i_k is highest \succ_s -ranked among applicants who desire school s , i.e., $i_k \in D_s(\mu)$.

For every k , if school $\mu(i_{k+1})$ has an empty seat, then $(i_k, \mu(i_{k+1}))$ forms a blocking pair. So stability of matching μ implies that $\mu(i_{k+1})$ is fully assigned. Construct a new matching by assigning i_k to $\mu(i_{k+1})$ for every k without changing the matches of other students. Each student i_k is matched to a more preferred school and remaining matches are unchanged. So the new matching Pareto dominates μ . Note that the new matching is also stable, because student i_k is highest $\succ_{\mu(i_{k+1})}$ -ranked among applicants who desire $\mu(i_{k+1})$. If μ is not student-optimal, then a stable improvement cycle exists, and a Pareto dominant stable matching can be computed in this fashion. This yields the following student-optimal stable algorithm (Erdil and Ergin, 2008):

Each school ranks all applicants first by priority and then by tie-breaker within the priority groups. Find a stable matching μ via the deferred acceptance algorithm.

Given a stable matching μ , if there is a stable improvement cycle, assign each student in the cycle to the school that the student desires in the cycle. Repeat this step until no stable improvement cycle can be found.

The algorithm terminates when there are no more stable improvement cycles. By Theorem 15, the stable improvements mechanism is not strategy-proof.

In the course of designing the New York City High School admissions process, policymakers from the Department of Education and some involved parents argued that, when a single tiebreaker is used, a student with a bad draw would be rejected from every choice, therefore schools should use different lotteries to break ties. Deferred acceptance simulations with preferences submitted by 78,728 8th-grade students in New York City 2006–2007 show that tiebreakers may have non-trivial welfare consequences (Abdulkadiroğlu et al., 2009). With a single tiebreaker, the deferred acceptance algorithm assigns on average 32,105 students to their first choice, while 5,613 remain unassigned. In contrast, with multiple tiebreakers, it assigns 2,256 fewer students to their top choice but leave 5,427 students (about 200 less) unassigned. Families’ concern is justified by slightly higher number of unassigned applicants with single tie breaking. But more applicants are assigned higher in their choice list in return. Abdulkadiroğlu et al. (2009) find that the distributions produced by single and multiple tiebreakers are statistically different, but there is no stochastic dominance ordering between the two rules in the sense of one tiebreaker assigning more applicants to their n -th and higher choices for every n . For example, the number of applicants receiving their seventh choice or better is greater under single tie breaking, while the number of applicants receiving their eighth choice or better is greater under multiple tie breaking. Despite these counteracting empirical observations, there is a clear theoretical comparison between the two types of tie breaking.

Theorem 18. (Ehlers, 2006; Erdil, 2006; Abdulkadiroğlu et al., 2009) *Every student-*

optimal matching can be produced by the deferred acceptance algorithm with some single tie breaker.

This result implies that if a stable matching can be produced by the deferred acceptance algorithm with a set of multiple tiebreakers, but not with any single tie breaker, it cannot be student-optimal and it must necessarily involve inefficiencies.

7 Generalized School Rankings

So far, the models and algorithms introduced in this chapter have assumed that school preferences can be captured by a simple ranking of the students. However, in many real-life school choice programs, the admission policies or school preferences cannot be represented in such a simple way. It will be instructive to illustrate the problem before generalizing the model.

Selective schools in Chicago rank students according to a score that is composed of the CPS High School Admissions Exam (50% of the total) and 7th-grade grades for core subjects (50%). Admissions to these schools also aim to achieve socioeconomic diversity at each school. Chicago Public Schools (CPS) groups residential areas into four socioeconomic tiers according to median family income, percentage of single-parent household, percentage of households where English is not the first language, percentage of homes occupied by the homeowner, the level of educational attainment, and achievement scores of attendance area schools for students living in the census track. Each tier group is supposed to contain about a quarter of students. 30 percent of the seats at selective schools are assigned to highest-performing applicants according to a composite score regardless of their tier group. These students tend to live in wealthier, more educated parts of the city. The remaining 70 percent of the seats are allocated equally among tiers groups and assigned to highest-performing applicants within each tier.²⁵

Although CPS ranks individual students according a composite score, the CPS policy implies that the district prefers a socioeconomically balanced student body to an unbalanced one even if the latter includes the highest ranked students. This type of complex admissions policies can not be captured adequately by a strict ranking of the students.

The canonical school choice model introduced in Section 2 can be extended by generalizing school preferences over sets of students. Such preference rankings are rather complex mathematical objects. Alternatively, school preferences can also be

²⁵See <https://go.cps.edu/high-school/selection>.

modeled via choice functions. Whereas the former asks how a school compares two different sets of students, the choice function approach asks which subset would the school choose from a given set of students. This section will take the choice-function approach. It will also lay a foundation for the next section, in which diversity in school admissions is considered.

Let $\mathcal{C}_s : 2^N \rightarrow 2^N$, be a choice function for school s . Notice that the function takes any subset of the set of students N as input and produces another subset of N . For it to model the selection of school s from any subset of students, assume that $\mathcal{C}_s(I) \subseteq I$ and $|\mathcal{C}_s(I)| \leq q_s$ for all $I \subseteq N$. That is, school s selects a subset of the students in I without exceeding its capacity. The remaining students in $I \setminus \mathcal{C}_s(I)$ are said to be rejected.

In the deferred acceptance algorithm introduced in Section 4, the schools select the highest ranked students at each step. This can be viewed as a choice function that selects students, one at a time, in the rank order up to the school capacity. That choice function is said to be responsive to school rankings (Roth, 1985a).

The deferred acceptance algorithm can be generalized with choice functions as follows:

1. Each student applies to her most preferred acceptable school that has not rejected her.
2. Let $I^s \subseteq N$ denote the set of students applying to school s . Each school s rejects the students in $I^s \setminus \mathcal{C}_s(I^s)$ but not the students in $\mathcal{C}_s(I^s)$.
3. When there are no rejections, each school s is matched with $\mathcal{C}_s(I^s)$ and the algorithm terminates.

The notion of blocking pairs can be generalized to for choice functions as well. A student-school pair (i, s) blocks a matching μ if $s \succ_i \mu(i)$ and $i \in \mathcal{C}_s(\mu(s) \cup \{i\})$. That is, student i prefers school s to her match and s selects i when it considers i along with the set of students $\mu(s)$ it is matched with. Then the definition of stability follows from this new blocking pair notion as before.

Showing existence of a stable matching and the stability of the deferred acceptance algorithm is more challenging in this model. A school may select a group of students over another for many reasons. The theory fails to generalize without further restrictions on choice functions. A well-known cause of such failure is complementarities among students in school preferences. When a school deems two students complementary to each other, it selects one only if it can also select the other. A

choice function is substitutable if for all $i \in I' \subseteq I \subseteq N$:

$$i \in \mathcal{C}_s(I) \Rightarrow i \in \mathcal{C}_s(I').$$

In words, if a school chooses a student from the set I , it chooses the student from every subset I' of I . This implies that there is no student in $I \setminus I'$ that is complementary to the student.

A choice function satisfies the irrelevance of rejected applicants if removing a student that it is not choose from a set does not impact the selection of the remaining students, i.e., for all $I \subseteq N$ and $i \in I$:

$$i \notin \mathcal{C}_s(I) \text{ implies } \mathcal{C}_s(I) = \mathcal{C}_s(I \cup \{i\}).$$

Theorem 19. (*Hatfield and Milgrom, 2005; Aygün and Sönmez, 2013*) *If \mathcal{C}_s is substitutable and satisfy the irrelevance of rejected applicants for all $s \in S$, then the deferred acceptance algorithm is stable and Pareto dominates any other stable algorithm.*

The following restriction on choice sets recovers strategy-proofness of the deferred acceptance algorithm. A choice function satisfies the law of aggregate demand if it chooses weakly more students when it considers additional students, that is, for all $I' \subseteq I \subseteq N$:

$$|\mathcal{C}_s(I')| \leq |\mathcal{C}_s(I)|.$$

Districts prefer to utilize seats at schools when there is demand. If a school selects a smaller set of students from a larger set, then there must be underutilized seats at the school. So, the law of aggregate demand is a plausible and natural restriction in the context of school choice.

Theorem 20. (*Hatfield and Milgrom, 2005*) *If \mathcal{C}_s is substitutable and satisfies the law of aggregate demand for all $s \in S$, then the deferred acceptance mechanism is stable, strategy-proof and Pareto dominates any other stable algorithm.*

8 Extensions

Real-life applications of school choice sometimes involve complications that go beyond the model presented in Section 2. This section considers two major extensions.

The first extension concerns preferences of school districts over socioeconomic distributions of students that are admitted to schools. Section 8.1 generalizes the

model to account for such preferences. It also develops an axiomatic framework to implement such preferences and policies in student assignment.

The second extension concerns multiple school sectors within a district, such as private schools that operate and run their admissions outside of the traditional school system, charter schools that run their own admissions separately from other schools. Whether admissions in different school sectors are administered within a single system or in several parallel or sequential systems have important welfare implications. Section 8.2 discusses this problem in the context of private and public schools.

8.1 Controlled School Choice

Diversity in school admissions is a common goal in the US (Abdulkadiroğlu, 2005; Ehlers, 2010; Ehlers et al., 2014; Hafalir et al., 2013), India (Aygün and Turhan, 2017; Sönmez and Yenmez, 2021), and other countries around the world. The Chicago Public Schools (CPS) case, mentioned above, is an example of aiming for socioeconomic balance at selective schools. New York City, by comparison, aims for academic diversity at certain schools. The city offers educational option, or simply, EdOpt schools. Half of the seats at EdOpt schools are allocated via lottery and the other half according to school preferences. The schools must enroll students with varying skills measured by the 7th grade scores on the state English exams. 16 percent of seats at EdOpt schools are reserved for low-performing students, 68 percent for average performing, and the remaining 16 percent of seats for top-performing students. School choice programs with such diversity controls are usually referred to as controlled school choice.

The generalized model in Section 7 lays the groundwork for formulating preferences over diversity of student groups in addition to individual student rankings. To this end, assume that there is a finite set of types T . In the CPS example, T constitutes of four tiers. Consider a district that aims for academic diversity in enrollment at each school according to average achievement at schools pupils are coming from. If the district grades schools, the type space may be set as $T = \{A, B, C, D \text{ and below}\}$. Notice that T contains mutually exclusive types in each example. If the district, in addition, aims for family income diversity, the type space can be $T = \{\text{Low Income, Middle Income, High Income}\} \times \{A, B, C, D \text{ and below}\}$. Generalization to multidimensional type spaces usually requires restrictive assumptions that are beyond the scope of this chapter. Thus, assume that T contains mutually exclusive types.

The set T partitions all students according to their type. Let $\tau : N \rightarrow T$ be a type

function such that $\tau(i) \in T$ is student i 's type. For a set of students $I \subseteq N$, let

$$I_t = \{i \in I : \tau(i) = t\}$$

be the set of type- t students in I .

In the examples above, certain numbers of seats are reserved for different types of students. It is also common practice to limit the number applicants of a type that can be admitted at schools. Refer to the first type of control as reserves and the second type as quotas. The aforementioned algorithms can be modified to accommodate reserves and quotas in intuitive ways.

Reserves and quotas are two different ways of imposing distributional restrictions. However, they operate in different fashion and have different welfare implications. It is possible that quotas may hurt their intended beneficiaries (Kojima, 2012), while reserves may outperform quotas in terms of student welfare (Hafalir et al., 2013). The following example demonstrates this point.

Example 9. Suppose that there are three students, 1, 2 and 3, and two schools, s_1 and s_2 with $q_{s_1} = 2$ and $q_{s_2} = 1$. Students 1 and 2 are from high income families and student 3 is from a low income family. Let $T = \{L, H\}$ denote student types, where $\tau(1) = \tau(2) = H$ and $\tau(3) = L$. Preferences and priorities are as follows:

Student 1 : $s_1 P_1 s_2$,

Student 2 : $s_1 P_2 s_2$,

Student 3 : $s_2 P_3 s_1$,

School s_1 : $1 \succ_{s_1} 2 \succ_{s_1} 3$,

School s_2 : $2 \succ_{s_2} 3 \succ_{s_2} 1$.

When there are no reserves and quotas, the deferred acceptance algorithm produces the following matching:

$$\mu = ((1, s_1), (2, s_1), (3, s_2)).$$

In this matching, every student is matched with her first choice. Now suppose that school s_1 has a quota of 1 for high income students. The deferred acceptance algorithm can be easily adapted with quotas as follows: at each step of the deferred acceptance algorithm, school s_1 provisionally admits one student at a time until the quota for high income students are met (Abdulkadiroğlu, 2005). Then the school rejects the remaining high income students and admits only low income students. This version of the algorithm produces the following matching:

$$\mu' = ((1, s_1), (2, s_2), (3, s_1)).$$

In this case, imposing quotas for high-income students hurts the low-income student: the low-income student 3 prefers matching μ to matching μ' .

Alternatively, suppose that one seat is reserved for low income students at school s_1 , which guarantees that 3 is assigned to s_1 whenever she applies to it despite of her low priority at the school. Since both 1 and 2's first choice is s_1 and 3's first choice is school s_2 , the deferred acceptance algorithm with reserves at school s_1 produces low-income student 3's preferred matching μ . \square

Earlier work on controlled choice has developed intuitive ad hoc solutions. However, at a more fundamental level, controlled school choice can be viewed as a problem about school preferences. The question can be formulated as how to balance diversity goals with school rankings of individual students. In particular, if a school uses merely its ranking, it is likely to miss its diversity goals. Focusing solely on diversity goals is likely to create many priority violations. So the question reduces to its choice function. Given diversity goals and school rankings of students, what are the desired properties of a choice function? Does any choice function satisfy these properties?

Recently, an axiomatic foundation has been developed for this question (Echenique and Yenmez, 2015). Consider a school with capacity of q and strict ranking \succ . Let \mathcal{C} denote its choice function (note that the school subscript is dropped for notational convenience).

Type-specific reserves guarantee a certain number of seats at a school for students of a certain type. Formally, let r_t be the number of seats reserved for type- t students. Let $R = (r_t)_{t \in T}$ be the vector of reserves such that $\sum_{t \in T} r_t \leq q$. The vector R will be referred to as type-specific reserves. Then, a choice function is reserves-respecting for R if for all $I \subseteq N$ and $t \in T$,

$$|\mathcal{C}(I)_t| \geq \min\{|I_t|, r_t\}.$$

This inequality states that type- t students should be chosen from the set as long as there are enough reserved seats for type t .

Quotas limit the number of students of a certain type at the school. Let q_t is the maximum number of seats that can be assigned to type- t students. Let $Q = (q_t)_{t \in T}$ be the vector of quotas such that $\sum_{t \in T} q_t \geq q$. Q will be referred to as type-specific quotas. Then, a choice function is quotas-respecting for Q if, for all $I \subseteq N$ and $t \in T$, i.e.,

$$|C(I)_t| \leq q_t.$$

A natural restriction on reserves and quotas is that $r_t \leq q_t$ for all $t \in T$.

In the following, fix the school capacity q , a strict school ranking \succ , reserves R and quotas Q .

The first property below concerns utilization of seats. Namely, it states that there is no reason to reject a student unless the quota for the student's type is met. Choice functions that do not satisfy this property may be considered inefficient or wasteful. Formally, a choice function \mathcal{C} is non-wasteful if for all $I \subseteq N$ and $i \in I$, if student i is rejected and there available seats at the school, i.e.,

$$i \in I \setminus \mathcal{C}(I) \text{ and } |\mathcal{C}(I)| < q,$$

then the quota for i 's type must be met, i.e.:

$$|\mathcal{C}(I)_{\tau(i)}| = q_{\tau(i)}.$$

As discussed above, there is a trade-off between meeting reserves and quotas and avoiding priority violations. A choice function \mathcal{C} violates the priority of student $i \in I$ if there is a student $j \in I$ such that i is rejected, j is chosen and i is ranked higher by the school, i.e.,

$$i \in I \setminus \mathcal{C}(I), j \in \mathcal{C}(I) \text{ and } i \succ j.$$

Priority violations can be classified into two groups. The first one concerns violations among same type students. This kind of violation cannot be justified by diversity concerns. Because swapping students between the chosen and rejected sets eliminates the violation without changing type distribution among the selected students. Therefore, any priority violation should happen only among different types of students. This is formalized as follows. A choice function \mathcal{C} is within-type priority compatible if, for all $I \subseteq N$, whenever student i is chosen from I , student j is rejected and they are of the same type, then i must be ranked higher than j by s , i.e.,

$$i \in \mathcal{C}(I), j \in I \setminus \mathcal{C}(I) \text{ and } \tau(i) = \tau(j) \text{ imply } i \succ j$$

for all $i, j \in I$.

A priority violation between two students of different types cannot be justified when there is room for assigning more students of the rejected student's type and the chosen student's type is sufficiently represented. The former means that the quota for rejected student's type is not met. The latter means that enough students of the chosen type is selected, there is no need for the chosen type to occupy a reserved seat for her type. The following property eliminates such violations. A choice function \mathcal{C} is across-type priority compatible if for all $I \subseteq N$, whenever student i is chosen from I , student j is rejected, they are of different types, more students with the same type of student i are chosen than the number of reserved seats for their type but the quota for student j 's type is not met, then student i must be ranked higher than student j by school s . Formally, for all $i, j \in I$ such that $i \in \mathcal{C}(I)$ and $j \in I \setminus \mathcal{C}(I)$,

$$|\mathcal{C}(I)_{\tau(i)}| > r_{\tau(i)} \text{ and } |\mathcal{C}(I)_{\tau(j)}| < q_{\tau(j)} \text{ imply } i \succ j.$$

In general, a stable matching need not be reserves- and quotas-respecting. Therefore, meeting the reserves and quotas may cause priority violations. Minimizing such violations subject to diversity constraints is a reasonable objective. In other words, when diversity is the primary objective, one can look for reserves- and quotas-respecting solutions that minimize priority violations.

This can be formalized as follows. A choice function \mathcal{C} is priority-violations minimal in the class of choice functions Γ , if $\mathcal{C} \in \Gamma$ and for any $\mathcal{C}' \in \Gamma$ and $I \subseteq N$, the number of students whose priority is violated at $\mathcal{C}(I)$ is smaller than that at $\mathcal{C}'(I)$.

The class of reserves- and quotas-respecting choice is rather large (Abdulkadiroğlu and Grigoryan, 2021b). The following reserves-and-quotas choice function selects students for reserves first, then it selects for the remaining seats among remaining students subject to quotas in the order of the school ranking. Formally, the reserves-and-quotas choice function is defined as follows: for each subset of students $I \subseteq N$,

1. For each type $t \in T$, select up to r_t of type- t applicants in I in the order of \succ .
2. Among the remaining students, select one at a time in the order of \succ without exceeding type-specific quotas and school capacity.

This function is non-wasteful and is reserves- and quotas-respecting. Dur et al. (2014) study a situation with an opposite order of processing. To compare the two solutions, suppose that there are no quotas, i.e., that $q_t = q$ for all $t \in T$. The alternative solution first selects $q - \sum_{t \in T} r_t$ of the highest ranked students. Then for each type t , it selects r_t of the highest ranked type t students among the remaining. Finally, it selects the highest ranked applicants from the remaining students for the remaining empty seats. The resulting choice function is reserves-respecting, non-wasteful and within-type priority compatible.

The following example demonstrates the difference between two solutions, which leads to the main results of this section.

Example 10. Suppose that there are three students, 1, 2 and 3. Students 1 and 3 are from low income families, student 2 is from a high income family. Let the type space be $T = \{L, H\}$, where L and H stands for low-income and high-income, respectively. Then $\tau(1) = \tau(3) = L$ and $\tau(2) = H$. Suppose now that the school has two seats, one of them is reserved for low-income students. There are no other distributional requirements. Equivalently, $r_L = 1, r_H = 0$ and $q_L = q_H = q = 2$. The school ranking is $1 \succ 2 \succ 3$.

The reserves-and-quotas choice function first selects student 1 for the reserved seat, then chooses student 2 for the remaining seat. The set of chosen students, $\{1, 2\}$, has

one low-income and one high-income student. The alternative solution first selects student 1 for one of the seats, then it chooses student 3 for the reserved seat. This time, the set of chosen students, $\{1, 3\}$, has only low-income students and a priority violation since $2 \succ 3$. \square

As the above example demonstrates, the processing order of reserves has distributional consequences. Both the constraints imposed by reserves and quotas and the order of processing of reserves determine the type distribution in the final selection. More importantly, the alternative solution causes a priority violation. There is no priority violation with the reserves-and-quotas choice function in this example. This observation is quite general: the reserves-and-quotas choice function is priority violations minimal among all such mechanisms.

Theorem 21. *(Abdulkadiroğlu and Grigoryan, 2021b) A choice function is the reserves-and-quotas choice function if and only if it is priority violations minimal in the class of reserves-respecting, quotas-respecting and non-wasteful choice functions.*

The following characterization provides further justification for the reserves-and-quotas choice function by showing that any priority violation can be justified by diversity constraints. In particular, the reserves-and-quotas choice function does not allow for within-type priority violations. Any remaining violation is justified in the sense of across-type priority compatibility.

Theorem 22. *(Imamura, 2020) A choice function is the reserves-and-quotas choice function if and only if it is a reserves-respecting, quotas-respecting, non-wasteful, within-type priority compatible and across-type priority compatible choice function.*

An alternative but equally plausible objective is to select highest priority students subject to diversity constraints. Formally, a choice function is priority maximal in a class of choice functions Γ , if $\mathcal{C} \in \Gamma$ and for any $\mathcal{C}' \in \Gamma$ and $i, j \in I \subseteq N$,

$$i \in \mathcal{C}(I) \setminus \mathcal{C}'(I) \text{ and } j \in \mathcal{C}'(I) \setminus \mathcal{C}(I) \text{ implies } i \succ j.$$

That is, any student chosen by \mathcal{C} but not by \mathcal{C}' must be ranked higher than any student chosen by \mathcal{C}' but not by \mathcal{C} . This is a plausible but demanding requirement, that is satisfied by the reserves-and-quotas choice function.

Theorem 23. *(Abdulkadiroğlu and Grigoryan, 2021b) A choice function is the reserves-and-quotas choice function if and only if it is priority maximal in the class of reserves-respecting, quotas-respecting, non-wasteful and within-type priority compatible choice functions.*

It turns out that the observation about the welfare implications of reserves and quotas in Example 9 holds generally when there are only two types.

Theorem 24. (*Hafalir et al., 2013*) *Suppose that $T = \{L, H\}$. Let a reserves regime be such that each school reserves r seats for type L students and has no quotas, where $r \leq \min_{s \in S} q_s$. Let a corresponding quotas regime be such that each school s has a quota of $q_s - r$ for type H students and does not reserve any seats. Every school uses the reserves-and-quotas choice function with the associated reserves and quotas. Then every student weakly prefers the reserves regime to the corresponding quotas regime.*

Ehlers et al. (2014) consider the problem with arbitrary many student types. In this generalized setting, they show that the school-proposing deferred acceptance algorithm finds a matching that minimizes violations of controlled choice constraints among the matchings where no student justifiably envies any other student.

Finally, the reserves-and-quotas choice function is substitutable and satisfies the law of aggregate demand. The next result follows from Theorem 20.

Theorem 25. *Consider a controlled school choice problem in which every school has a reserves-and-quotas choice function. Then, the deferred acceptance algorithm is stable, strategy-proof and Pareto dominates any other stable algorithm.*

8.2 Sequential and Parallel Admissions

In Boston, Chicago, and New York City, admissions to elite exam schools are administered separately from admissions to other public schools. In most American school districts, charter schools run their independent admissions. Likewise, private schools also make their admissions decisions outside of the public school admissions. In the OECD countries and their 14 partner economies, 18 percent of all 15-year-old students are admitted to non-public schools (OECD, 2012).²⁶ In many of these countries, private schools have their own separate admission systems. In Sweden, for example, private schools normally take care of their own admissions based on queueing time whereas public schools normally admit students based on walking distance and sibling priority.

Such multi-tiered admissions can be administered via centralized, decentralized, parallel, and sequential systems. For instance, students in New York City (NYC) have to fill in two separate applications when applying to exams schools and non-exam schools. The city first assigns students to exams schools. Then it administers

²⁶The only exceptions are Azerbaijan, Romania and the Russian Federation. They do not have any private schools.

admissions for the general population. Consequently, some students can be assigned both an exam school and a non-exam school.²⁷

Similar systems have been used in France and Turkey for college admissions and in some Swedish municipalities for school choice (see Andersson et al., 2020; Haeringer and Iehle, 2017). In Turkish college admissions, students apply to the private schools in Round 1. Unlike the NYC case, only those who are unassigned in Round 1 are allowed to apply to public schools in Round 2. This is a significant change from the earlier Turkish system, in which all students could apply to public schools in Round 1 and they could also apply to private schools in Round 2 independently of the outcome of Round 1.²⁸ These types of systems will henceforth be referred to as sequential systems, as students are assigned to schools in multiple rounds. In other systems, students are assigned in only one round but in parallel systems. This is the case in several US states as well as in many Swedish municipalities where different school sectors run their own admissions in parallel within the same school district (see, e.g. Andersson et al., 2020; Manjunath and Turhan, 2016).

Such systems are likely to feature inefficiencies and under-utilization of school seats. They also raise strategic issues since students may need to apply to schools in different time periods and have to decide whether or not to accept an offer from a school without knowing the outcome of all their applications. The following example from Andersson et al. (2020) illustrates some of the problems associated with sequential admission systems (a similar example can be constructed for parallel systems).

Example 11. Let $N = \{1, 2, 3, 4, 5\}$ and suppose that the set of schools $S = \{s_1, s_2, s_3, s_4, s_5\}$ is partitioned into private and public schools given by $S^1 = \{s_1, s_2, s_3\}$ and $S^2 = \{s_4, s_5\}$, respectively. Assume further that each school has a single seat. Student preferences are given by:

- Student 1: $s_1 P_1 s_2 P_1 s_5 P_1 1$,
- Student 2: $s_1 P_2 s_4 P_2 s_2 P_2 2$,
- Student 3: $s_4 P_3 s_3 P_3 s_5 P_3 3$,
- Student 4: $s_3 P_4 s_5 P_4 s_4 P_4 4$,
- Student 5: $s_1 P_5 s_2 P_5 s_3 P_5 5$.

For simplicity and without loss of generality, assume that all schools have a common priority order given by $1 \succ 2 \succ 3 \succ 4 \succ 5 \succ s$ for each $s \in S$. Consider now the

²⁷See Abdulkadiroğlu et al. (2005a, 2009, 2017a) for more detailed discussions.

²⁸This is a simplified description of the systems. For more details, see Andersson et al. (2020).

following three matchings:

$$\begin{aligned}\mu^1 &= ((1, s_1), (2, s_4), (3, s_3), (4, s_5), (5, s_2)), \\ \mu^2 &= ((1, s_1), (2, s_2), (3, s_3), (4, s_5), (5, 5)), \\ \mu^3 &= ((1, s_1), (2, s_4), (3, s_3), (4, 4), (5, 5)),\end{aligned}$$

and recall that a student is unassigned if she is matched to herself, like students 4 and 5 in matching μ^3 . In this example, matching μ^1 is generated by the deferred acceptance algorithm for the case when all schools are part of a one round admission system, i.e., when the set of schools S is not partitioned into two subsets. The matchings μ^2 and μ^3 are generated by the new and the old Turkish systems, respectively, assuming that all students report truthfully, which need not be the case.

In the new Turkish system, students first apply to the private schools and are assigned seats according to a serial dictatorship induced by the common \succ . Hence, schools s_1 , s_2 and s_3 are assigned to students 1, 2, and 3, respectively. These students are not allowed to participate in Round 2, where only unmatched students can apply to public schools. Because student 5 is unassigned and does not rank any public school, she will remain unmatched. Consequently, student 4 is the only student that is allowed and willing to participate in Round 2. Therefore, student 4 is assigned her most preferred public school, namely, school s_5 . This gives matching μ^1 .

In the old Turkish system, students are assigned to public schools in Round 1, but all students are allowed to participate in Round 2. If they are assigned to a school in both rounds, they can freely select which of the two schools they would like to get admitted to. As in the new Turkish system, the schools are assigned in each round using a serial dictatorship induced by the common \succ . Using the same arguments as in the above, student 5 will not participate in Round 1, and schools s_4 and s_5 are assigned to students 2 and 1, respectively. In Round 2, the private schools s_1 , s_2 and s_3 are assigned to students 1, 2, and 3, respectively. Students 1 and 2 now select their most preferred schools of their multiple assignments (for example, since $s_4 P_2 s_2$, student 2 selects school s_4). This gives matching μ^2 .

The matchings generated by the old and the new Turkish systems have some obvious problems. For example, in all matchings except μ^1 there is at least one unassigned student. Furthermore, matching μ^3 is Pareto dominated by matching μ^1 , and the matching μ^2 is unstable, since $(2, s_4)$ constitutes a blocking pair. In fact, the only stable and Pareto efficient matching in this example is μ^1 , but this is also the only matching that is not generated in a sequence. \square

In spite of the apparent problems illustrated in the above example, the matching

literature has just recently started to pay serious attention to sequential and parallel systems. A common feature of all sequential admission systems is that some seats at private and public schools may become available between rounds. For instance, if a student rejects his assignment to a private school in the first round, it becomes available in the second round, in which only public school admissions are administered. This has been investigated in the literature by, for example, Doğan and Yenmez (2018a,b) and Dur and Kesten (2018) in the context of school choice, and by Haeringer and Iehle (2017) for college admissions in France, and by Westkamp (2013) in the context of German university assignment.

Sequential systems also create incentives for strategic behavior (Haeringer and Iehle, 2017). First, students that are admitted in an earlier round need to determine if they should participate in a later round. It is not immediately clear that they always gain by participating in later rounds if they are forced to reject their first round offer to participate in the second round, and consequently face the risk of being placed at a lower ranked school. Second, reporting preferences truthfully may no longer be a dominant strategy. The results related to strategy-proofness are mainly negative (Andersson et al., 2020; Doğan and Yenmez, 2018b; Dur and Kesten, 2018; Haeringer and Iehle, 2017). For example, Haeringer and Iehle (2017) show that sequential mechanisms are not strategy-proof even if a strategy-proof mechanism is used at each round. Dur and Kesten (2018) investigates the conflict between efficiency and incentives in sequential admission systems where students have to choose in which round to participate and argue that unified admissions leads to superior welfare and incentive properties as illustrated by matching μ^1 in Example 11. Doğan and Yenmez (2018a,b) analyze developments in the Chicago school system and investigate the welfare effects and incentive properties of the system. They provide a thorough analysis on how students are expected to behave in the sequential admissions system. They also identify a priority structures and school capacities under which an additional period of matching benefits students in equilibrium. Andersson et al. (2020) show that there is a trade-off between a notion of truthfulness, called straightforwardness, and non-wastefulness, so inefficiency of sequential admissions is almost unavoidable.

Inefficiencies, wastefulness, and strategy-proofness are also major concerns in parallel admissions systems. Manjunath and Turhan (2016) investigate a school choice system where separate sets of schools run their admissions processes in parallel using the deferred acceptance algorithm. They demonstrate that the resulting school assignment is often inefficient, and offer a way to Pareto improve upon these assignments by iteratively rematching students. Turhan (2019) further investigates this mechanism. Here, a major question relates to potential gains by merging two parallel systems into one common system.

Dur and Kesten (2018) investigate a sequential system where students are forced to choose the stage of admissions in which they participate in and argue that unified admissions leads to superior welfare and incentive properties. Unlike students, certain schools, such as charters, are granted to run their own admissions. Participation in general admissions remains a decision for such schools. Ekmekci and Yenmez (2019) investigate this issue in a sequential assignments model in which centralized admissions for district schools precedes individual admissions for charter schools. They show that a charter school is better-off by running its independent admissions after the centralized admissions. To overcome this problem, they propose a mechanism and show that an equilibrium can be sustained if all schools participate in the centralized clearinghouse, as in the case of New Orleans.

There is also a distinction between common application and a common admission. For instance, in New Orleans prior to 2003, students could apply to schools through a common applications process, but schools decided who to admit in a decentralized uncoordinated manner, creating congestion problems in the system (Abdulkadiroğlu et al., 2017c). In order to avoid such inefficiencies, the Recovery School District in New Orleans became the first US district in 2012 to unify charter and traditional public school admissions in a single-offer assignment mechanism known as OneApp. Abdulkadiroğlu et al. (2017a) quantify the welfare consequences of coordinating admissions. Using data from New York city prior to 2003, in which admissions to schools were uncoordinated, and data from post 2003 when the city adopted a coordinated admissions system, they identify large welfare gains from switching to the new system. In comparison, further gains from fine-tuning algorithmic details of the coordinated system are much smaller.

9 Research Design in School Choice

Assignment algorithms generate data that can be used to answer empirical questions about school effectiveness and policy interventions. When schools are oversubscribed and priorities are not sufficiently fine to determine the final assignment, tie breakers are used to ration seats, as discussed in Section 6. This generates quasi-experimental variation in student assignment, opening the door to credible research designs for program evaluation in education (see, e.g., Abdulkadiroğlu et al., 2017b, 2022).

When an algorithm uses random tie breakers (such as lottery numbers) and non-random tie breaking (such as entrance exam scores), each applicant’s probability of assignment to a school in her choice list can be readily computed by simulating the algorithm many times by redrawing random and non-random tie-breakers. Given access to the algorithm, one need not know anything about the algorithm, how it

operates, or the theory behind it.

A treatment is not necessarily summarized by assignment to a particular school. However, the probability of assignment to treatment can be computed rather easily. For example, if the treatment is attendance at a charter school, the probability of assignment to treatment would be the sum of probabilities of assignment to charter schools. More complicated treatments can be studied as well. Suppose a policy maker is interested in the impact of attending a school with 50% or more students of applicant's own race. In that case, at a heavily black school, a black student would be treated but white students would be untreated. Regardless, probability of assignment to treatment can be computed for each applicant.

After the probability of assignment to treatment is computed for each applicant, the causal effect of treatment can be estimated via the method of Abadie (2003), or an equivalent 2SLS strategy. Formally, let Z be a binary variable indicating whether an applicant is assigned a treatment or not, $D_z \in \{0, 1\}$ represent potential enrollment in an treatment school given $Z = z$, and Y_{zd} denote potential outcome given $Z = z$ and $D = d$. Let θ denote the preference list and priorities and X denote the set of all potential predetermined student characteristics except θ . Let $P(\theta, X) = \text{Prob}(Z = 1|\theta, X)$ be the probability of assignment to treatment given θ and X . As explained above, this probability is readily computed via simulations.

The following assumptions are standard (see Imbens and Angrist, 1994):

- i. Independence of the instrument: Conditional on (X, θ) , the random vector $((Y_{0d}, Y_{1d})_{d \in \{0,1\}}, D_0, D_1)$ is independent of Z .
- ii. Exclusion restriction: $\text{Prob}(Y_{0d} = Y_{1d}|\theta) = 1$ for $d \in \{0, 1\}$.
- iii. First stage: $0 < \text{Prob}(Z = 1|X, \theta) < 1$ and $\text{Prob}(D_1 = 1|X, \theta) > \text{Prob}(D_0 = 1|X, \theta)$.
- iv. Monotonicity: $\text{Prob}(D_1 \geq D_0|X, \theta) = 1$.

Given the exclusion restriction, let $Y_d = Y_{0d} = Y_{1d}$ for $d \in \{0, 1\}$. Also define the observed treatment status as $D = ZD_1 + (1 - Z)D_0$ and the observed outcome as $Y = DY_1 + (1 - D)Y_0$. Note that an econometrician observes only Z , D , and Y for each student. The main result of Abadie (2003) applies directly to our case:

Theorem 26 (Abadie 2003). *Let $g(Y, D, X, \succ, \vartheta)$ be a measurable real function of (Y, D, X, θ) such that $E[|g(Y, D, X, \theta)|] < \infty$. Define:*

$$\kappa = 1 - \frac{D(1 - Z)}{\text{Prob}(Z = 0|X, \theta)} - \frac{(1 - D)Z}{\text{Prob}(Z = 1|X, \theta)}.$$

If Assumption 1 holds, then $E[g(Y, D, X, \theta)|D_1 > D_0] = \frac{1}{\text{Prob}(D_1 > D_0)} E[\kappa g(Y, D, X, \theta)]$.

In other words, one can estimate $g(Y, D, X, \theta)$ for the sub-population of compliers with $D_1 > D_0$ without explicitly knowing each student's (D_0, D_1) . Consider the following constant-effects causal model:

$$Y = \alpha + \beta D + \gamma X + \varepsilon, \quad (1)$$

where β is the treatment effect. Let:

$$(\alpha, \beta, \gamma) = \arg \min_{a,b,c} E[(Y - (a + bD + cX))^2 | D_1 > D_0].$$

This equation cannot be estimated directly, since one does not observe (D_0, D_1) , but the theorem implies that:

$$(\alpha, \beta, \gamma) = \arg \min_{a,b,c} E[\kappa(Y - (a + bD + cX))^2].$$

Note that this expectation is over the entire population and not just the compliers. The sample counterpart of this gives an unbiased estimator of the parameters. To this end, compute for each student i :

$$\kappa_i = 1 - \frac{D_i(1 - Z_i)}{1 - P(X_i, \theta_i)} - \frac{(1 - D_i)Z_i}{P(X_i, \theta_i)},$$

where $P(X_i, \theta_i)$ is estimated by simulating the algorithm. Then, estimate the parameters by:

$$(\hat{\alpha}, \hat{\beta}, \hat{\gamma}) = \arg \min_{a,b,c} \sum_i (\kappa_i(Y_i - (a + bD_i + cX_i)))^2.$$

Since the assignment probability is determined by preferences and priorities, and these are discrete variables, the probability is a linear function of preferences and priorities. Therefore, an equivalent 2SLS estimator exists (Proposition 5.1, Abadie (2003)).

Abdulkadiroğlu et al. (2017b) develops an asymptotically valid 2SLS strategy without simulations based on large economies for deferred acceptance with random tie breaking. Abdulkadiroğlu et al. (2022) generalizes that approach for deferred acceptance with random and non-random tie breaking and develops a Regression Discontinuity Design. This summary barely scratches the surface and much remains to be done.

10 Conclusions

Parental choice over public schools has become a major part of education reform around the world. In the US alone, the proportion of the largest 100 schools districts with parental choice over public schools doubled between 2000 and 2016 (Whitehurst, 2017). Currently, parents in Belgium, England, Ireland, Italy, the Netherlands, Portugal, Sweden, Wales, and Northern Ireland have the right to or must choose a public or a private school at the primary, lower and upper secondary levels for their children. In fact, for the upper secondary level, only six European countries (Denmark, Greece, France, Cyprus, Malta and Turkey) do not have any type of school choice program and, instead, assign students to schools based on their place of residence (European Commission, 2020). Consequently, the demand for rigorous solutions for student assignment has been growing. Each school-choice program comes with its own institutional and political constraints, which opens the door for further research. In fact, most of the theoretical literature on school choice is motivated by real-life problems identified in the field.

In addition, admissions policies are frequently decided without consulting theory. Ad hoc details and features may also be added at this stage to a yet-to-be designed admissions process. While most of these policies and features can be accommodated by adapting the algorithms introduced in this chapter, certain decisions make it theoretically impossible to meet all requirements. In such instances, approximate market design solutions are developed by the guidance of theory and data analysis.

This chapter has focused on student assignment in school choice by taking preferences and priorities as exogenous. An equally important question concerns the welfare consequences of school choice. In particular, the models in the literature ignore probably one of the most important aspects of school choice: access to a school is determined not only by the ability to list the school in the application form, but also by admissions policies, such as neighborhood priority, and by means for traveling to the school. This makes both preferences and priorities endogenous. While wealthy families can choose affluent neighborhoods with good schools before going through the formal choice process, low income families are shut out of such residential choice. The matching models of school choice regularly ignore housing markets, families' endogenous housing decisions and their impact on welfare. Recent advances in this direction have been made by studying school choice with competitive housing markets and a continuum of students. For example, in a fairly general model, Grigoryan (2021) offers a convincing theoretical argument in favor of school choice by showing that low income families are better-off under deferred acceptance in comparison to solely residence-based school assignment even when applicants are granted neighborhood priorities in deferred acceptance. Much work is still to be done both theoretically and empirically on that frontier.

The welfare consequences of admissions priorities remains an open question. Mechanism design and market design have developed new theories and solutions for the problem of assigning pupils to schools, but have been mostly silent on the design of priorities. Recent advances on that front have been made (see, e.g., Grigoryan, 2021; Kloosterman and Troyan, 2020), but much remains to be done.

More importantly, data generated by assignment algorithms can be used to answer most pressing empirical questions on school effectiveness and policy interventions.²⁹ Seats in a school choice program are rationed by admissions priorities, such as neighborhood priority for pupils living within a certain distance from school, lotteries, student rankings at the school which may be based on an entrance examination, academic records, interviews and other criteria. Such rationing creates quasi-experimental variation in school assignment at unprecedented levels that can be used for credible evaluation of individual schools and of school reform models, such as charters, small schools, and voucher programs. A recent literature focuses on research design with data from centralized admissions and develops econometric techniques.

Finally, there is a growing literature that uses detailed micro data to, e.g., estimate parental preferences over schools. Advances on that front can guide enrollment and school portfolio planning at districts, see, for example, Abdulkadiroğlu et al. (2011a, 2017a), Agarwal and Somaini (2018), Ajayi (2014), Burgess et al. (2015), Calsamiglia et al. (2020), Deming et al. (2014), Hastings et al. (2009), He (2017), Myoung and Hwang (2014), Neilson (2021), and Walters (2018).

References

- Abadie, A. (2003). Semiparametric Instrumental Variables Estimation of Treatment Response Models. *Journal of Econometrics*, 113(2):231–263.
- Abdulkadiroğlu, A. (2005). College admission with affirmative action. *International Journal of Game Theory*, 33:535–549.
- Abdulkadiroğlu, A. (2019). The data from market design solutions for school choice. pages 439–444.
- Abdulkadiroğlu, A., Agarwal, N., and Pathak, P. A. (2017a). The welfare effects of coordinated assignment: Evidence from the New York City high school match. *American Economic Review*, 107:3635–3689.

²⁹See Abdulkadiroğlu (2019) for a brief discussion on empirical possibilities with data generated by centralized admissions systems.

- Abdulkadirođlu, A., Angrist, J. D., Dynarski, S. M., Kane, T. J., and Pathak, P. A. (2011a). Accountability and flexibility in public schools: Evidence from Boston’s charters and pilots. *Quarterly Journal of Economics*, 126:699–748.
- Abdulkadirođlu, A., Angrist, J. D., Narita, Y., and Pathak, P. A. (2017b). Research Design Meets Market Design: Using Centralized Assignment for Impact Evaluation. *Econometrica*, 85(5):1373–1432.
- Abdulkadirođlu, A., Angrist, J. D., Narita, Y., and Pathak, P. A. (2022). Breaking ties: Regression discontinuity design meets market design. *Econometrica*, 90(1):117–151.
- Abdulkadirođlu, A., Che, Y.-K., Pathak, P. A., Roth, A. E., and Tercieux, O. (2017c). Minimizing justified envy in school choice: The design of New Orleans’ Oneapp. *NBER Working Paper 23265*.
- Abdulkadirođlu, A., Che, Y.-K., Pathak, P. A., Roth, A. E., and Tercieux, O. (2020a). Efficiency, justified envy, and incentives in priority-based matching. *American Economic Review: Insights*, 2:425–442.
- Abdulkadirođlu, A., Che, Y.-K., Pathak, P. A., Roth, A. E., and Tercieux, O. (2020b). Efficiency, justified envy, and incentives in priority-based matching. *American Economic Review: Insights*, 2:425–442.
- Abdulkadirođlu, A., Che, Y.-K., and Tercieux, O. (2010). The role of priorities in assigning indivisible objects: A characterization of top trading cycles. *Working Paper*.
- Abdulkadirođlu, A., Che, Y.-K., and Yasuda, Y. (2011b). Resolving conflicting preferences in school choice: the Boston mechanism reconsidered. *American Economic Review*, 101:399–410.
- Abdulkadirođlu, A. and Grigoryan, A. (2021a). Efficient and envy minimal matching. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3646475.
- Abdulkadirođlu, A. and Grigoryan, A. (2021b). Priority-based assignment with reserves and quotas. *NBER Working Paper 28689*.
- Abdulkadirođlu, A., Pathak, P. A., and Roth, A. E. (2005a). The New York City high school match. *American Economic Review*, 95:364–367.
- Abdulkadirođlu, A., Pathak, P. A., and Roth, A. E. (2009). Strategy-proofness versus efficiency in matching with indifferences: Redesigning the NYC high school match. *American Economic Review*, 99:1954–1978.

- Abdulkadirođlu, A., Pathak, P. A., Roth, A. E., and Sönmez, T. (2005b). The Boston public school match. *American Economic Review, Papers and Proceedings*, 95:368–371.
- Abdulkadirođlu, A. and Sönmez, T. (1998). Random serial dictatorship and the core from random endowments in house allocation problems. *Econometrica*, 66:689–701.
- Abdulkadirođlu, A. and Sönmez, T. (2003). School choice: A mechanism design approach. *American Economic Review*, 93:729–747.
- Afacan, M. O. (2013). Alternative characterizations of Boston mechanism. *Mathematical Social Sciences*, 66:176–179.
- Agarwal, N. and Somaini, P. (2018). Demand analysis using strategic reports: An application to a school choice mechanism. *Econometrica*, 86:391–444.
- Ajayi, K. (2014). Does School Quality Improve Student Performance? New Evidence from Ghana. IED Discussion Paper No. 260.
- Alcalde, J. (1996). Implementation of stable solutions to marriage problems. *Journal of Economic Theory*, 69:240–254.
- Alcalde, J. and Barbera, S. (1994). Top dominance and the possibility of strategy-proof stable solutions to matching problems. *Economic Theory*, 4:417–435.
- Alcalde-Unzua, J. and Molis, E. (2011). Exchange of indivisible goods and indifference: The top trading absorbing sets mechanisms. *Games and Economic Behavior*, 73:1–16.
- Andersson, T., Dur, U., Ertemel, S., and Kesten, O. (2020). Sequential school choice with public and private schools. *Working Paper*.
- Andersson, T., Ehlers, L., and Svensson, L.-G. (2013). Budget-balance, fairness and minimal manipulability. *Theoretical Economics*, 9:753–777.
- Aygün, O. and Turhan, B. (2017). Large-scale affirmative action in school choice: Admissions to IITs in India. *American Economic Review*, 107:210–213.
- Aygün, O. and Sönmez, T. (2013). Matching with Contracts: Comment. *American Economic Review*, 103(5):2050–2051.
- Azevedo, E. M. and Budish, E. (2019). Strategy-proofness in the large. *The Review of Economic Studies*, 86:81–116.

- Aziz, H. and de Keijzer, B. (2012). Housing markets with indifferences: a tale of two mechanisms. *Proceedings of the 26th AAAI Conference on Artificial Intelligence*, pages 1249–1255.
- Balinski, M. and Sönmez, T. (1999). A tale of two mechanisms: Student placement. *Journal of Economic Theory*, 84:73–94.
- Bonkougou, S. and Nesterov, A. (2021). Comparing school choice and college admission mechanisms by their immunity to strategic admissions. *Theoretical Economics*, 16:881–909.
- BPS (2004). Introducing the Boston Public Schools: A Guide for Parents and Students. www.bostonpublicschools.org.
- Burgess, S., Greaves, E., Vignoles, A., and Wilson, D. (2015). What parents want: School preferences and school choice. *Economic Journal*, 125:1262–1289.
- Calsamiglia, C., Fu, C., and Güell, M. (2020). Structural estimation of a model of school choices: The Boston mechanism versus its alternatives. *Journal of Political Economy*, 128:642–680.
- Chen, Y. (2016). New axioms for immediate acceptance. *Review of Economic Design*, 20:329–337.
- Chen, Y. and Kesten, O. (2017). Chinese college admissions and school choice reforms: a theoretical analysis. *Journal of Political Economy*, 125:99–139.
- Chen, Y., Pycia, M., and Yenmez, M. B. (2017). Manipulability of stable mechanisms. *American Economic Journal: Microeconomics*, 8:202–214.
- Chen, Y. and Sönmez, T. (2006). School choice: an experimental study. *Journal of Economic Theory*, 127:202–231.
- Decerf, B. (2021). A modification aimed at reducing the manipulability and inefficiency of the boston school choice mechanism. *Social Choice and Welfare*, forthcoming.
- Deming, D. J., Hastings, J. S., Kane, T. J., and Staiger, D. O. (2014). School choice, school quality, and postsecondary attainment. *American Economic Review*, 104:991–1013.
- Doğan, B. and Ehlers, L. (2021). Robust minimal instability of the top trading cycles mechanism. *Theoretical Economics*, 16:1249–1279.
- Doğan, B. and Klaus, B. (2018). Object allocation via immediate-acceptance: Characterizations and an affirmative action application. *Journal of Mathematical Economics*, 79:140–159.

- Doğan, B. and Yenmez, M. B. (2018a). Unified versus divided enrollment in school choice: Improving student welfare in Chicago. *Working Paper*.
- Doğan, B. and Yenmez, M. B. (2018b). When does an additional stage improve welfare in centralized assignment? *Working Paper*.
- Dubins, L. E. and Freedman, D. A. (1981). Machiavelli and the Gale-Shapley algorithm. *American Mathematical Monthly*, 88:485–494.
- Dur, U., Kominers, S., Pathak, P. A., and Sönmez, T. (2014). The Demise of Walk Zones in Boston: Priorities vs. Precedence in School Choice. NBER Working Paper No. 18981.
- Dur, U. M. (2019). The modified boston mechanism. *Mathematical Social Sciences*, 101:31–41.
- Dur, U. M., Hammond, R. G., and Morrill, T. (2018a). Identifying the harm of manipulable school choice mechanisms. *American Economic Journal: Economic Policy*, 10:187–213.
- Dur, U. M., Hammond, R. G., and Morrill, T. (2018b). The secure Boston mechanism: Theory and experiments. *Experimental Economics*, 22:1–36.
- Dur, U. M. and Kesten, O. (2018). Sequential versus simultaneous assignment problems and two applications. *Working Paper*.
- Dur, U. M., Pathak, P. A., Song, F., and Sönmez, T. (2021). Deduction dilemmas: The Taiwan assignment mechanism. *American Economic Journal: Microeconomics*, forthcoming.
- Echenique, F., Immorlica, N., and Vazirani, V., editors (forthcoming 2022). *Online Matching Theory and Market Design*. Cambridge University Press.
- Echenique, F. and Yenmez, M. B. (2015). How to control controlled school choice. *American Economic Review*, 105:2679–2694.
- Ehlers, L. (2006). Respecting Priorities when Assigning Students to Schools. Working Paper.
- Ehlers, L. (2010). School choice with control. *Working Paper CIREQ Cahier 13–2010*.
- Ehlers, L. (2014). Top trading with fixed tie-breaking in markets with indivisible goods. *Journal of Economic Theory*, 151:64–87.
- Ehlers, L. and Erdil, A. (2010). Efficient assignment respecting priorities. *Journal of Economic Theory*, 145:1269–1282.

- Ehlers, L., Hafalir, I. E., Yenmez, M. B., and Yildirim, M. A. (2014). School choice with controlled choice constraints: Hard bounds versus soft bounds. *Journal of Economic Theory*, 153:648–683.
- Ehlers, L. and Morrill, T. (2020). (Il)legal assignments in school choice. *The review of Economic Studies*, 87:1837–1875.
- Ekmekci, M. and Yenmez, B. (2019). Common enrollment in school choice. *Theoretical Economics*, 14:1237–1270.
- Erdil, A. (2006). Two-sided Matching with Ties. PhD dissertation.
- Erdil, A. and Ergin, H. (2008). What’s the matter with tie-breaking? improving efficiency in school choice. *American Economic Review*, 98:669–689.
- Ergin, H. (2002). Efficient resource allocation on the basis of priorities. *Econometrica*, 70:2489–2497.
- Ergin, H. and Sönmez, T. (2006). Games of school choice under the Boston mechanism. *Journal of Public Economics*, 90:215–237.
- European Commission (2020). *Equity in school education in Europe: Structures, policies and student performance*. Luxembourg: Publications Office of the European Union.
- Fack, G., Grenet, J., and He, Y. (2019). Beyond truth-telling: Preference estimation with centralized school choice and college admissions. *American Economic Review*, 1486–1529:109.
- Gale, D. and Shapley, L. S. (1962). College admissions and the stability of marriage. *American Mathematical Monthly*, 69:9–15.
- Gibbard, A. (1973). Manipulation of voting schemes: A general result. *Econometrica*, 41:587–602.
- Grigoryan, A. (2021). School Choice and Housing Market. Working Paper.
- Haeringer, G. and Iehle, V. (2017). Gradual college admission. *Working Paper*.
- Haeringer, G. and Klijn, F. (2009). Constrained school choice. *Journal of Economic Theory*, 144:1921–1947.
- Hafalir, I. E., Yenmez, M. B., and Yildirim, M. A. (2013). Effective affirmative action in school choice. *Theoretical Economics*, 8:325–363.
- Hakimov, R. and Kübler, D. (2021). Experiments on centralized school choice and college admissions: a survey. *Experimental Economics*, 24:434–488.

- Harless, P. (2017). Immediate acceptance with or without skips? comparing school assignment procedures. *Working Paper*.
- Hastings, J., Kane, T. J., and Staiger, D. (2009). Heterogeneous preferences and the efficacy of public school choice. *Brown University Working Paper*.
- Hatfield, J. W. and Milgrom, P. (2005). Matching with Contracts. *American Economic Review*, 95:913–935.
- He, Y. (2017). Gaming the Boston school choice mechanism in Beijing. *Toulouse School of Economics Working Papers 15–551*.
- Hylland, A. and Zeckhauser, R. (1979). The efficient allocation of individuals to positions. *Journal of Political Economy*, 87:293–314.
- Imamura, K. (2020). Meritocracy Versus Diversity. Working Paper.
- Imbens, G. W. and Angrist, J. D. (1994). Identification and Estimation of Local Average Treatment Effects. pages 467–475.
- Jaramillo, P. and Manjunath, V. (2012). The difference indifference makes in strategy-proof allocation of objects. *Journal of Economic Theory*, 147:1913–1946.
- Kelly, J. S. (1988). Minimal manipulability and local strategy-proofness. *Social Choice and Welfare*, 5:81–85.
- Kesten, O. (2010). School choice with consent. *The Quarterly Journal of Economics*, 125:1297–1348.
- Kesten, O. (2012). On two kinds of manipulation for school choice problems. *Economic Theory*, 51:677–693.
- Kloosterman, A. and Troyan, P. (July 2020). School choice with asymmetric information: priority design and the curse of acceptance. *Theoretical Economics*, 15(3).
- Knuth, D. E. (1976). Marriages stables. *Les Presses de l'Université de Montréal*, 10.
- Kojima, F. (2006). Mixed strategies in games of capacity manipulation in hospital-intern markets. *Social Choice and Welfare*, 27:25–28.
- Kojima, F. (2012). School choice: Impossibilities for affirmative action. *Games and Economic Behavior*, 75:685–693.
- Kojima, F. and Ünver, U. (2014). The Boston school choice mechanism: An axiomatic approach. *Economic Theory*, 55:515–544.

- Konishi, H. and Ünver, U. (2006). Games of capacity manipulation in hospital-intern markets. *Social Choice and Welfare*, 27:3–24.
- Ma, J. (1994). Strategy-proofness and the strict core in a market with indivisibilities. *International Journal of Game Theory*, 23:75–83.
- Manjunath, V. and Turhan, B. (2016). Two school systems, one district: What to do when a unified admissions process is impossible. *Games and Economic Behavior*, 95:25–40.
- Maus, S., Peters, H., and Storcken, T. (2007). Anonymous voting and minimal manipulability. *Journal of Economic Theory*, 135:533–544.
- McVitie, D. G. and Wilson, L. B. (1970). Stable marriage assignment for unequal sets. *BIT Numerical Mathematics*, 10:295–309.
- Mennle, T. and Seuken, S. (2015). Trade-offs in school choice: comparing deferred acceptance, the naive and the adaptive Boston mechanism. *Working Paper*.
- Miralles, A. (2008). School choice: The case for the Boston method. *Working Paper*.
- Morrill, T. (2013). An alternative characterization of top trading cycles. *Economic Theory*, 54:181–197.
- Morrill, T. (2015). Two simple variations of top trading cycles. *Economic Theory*, 60:123–140.
- Myerson, R. B. (1979). Incentive compatibility and the bargaining problem. *Econometrica*, 47:61–73.
- Myoung, S. I. and Hwang (2014). A Robust Redesign of High School Match. *Working Paper*.
- Neilson, C. (2021). Targeted vouchers, competition among schools, and the academic achievement of poor students. *Working Paper*.
- OECD (2012). Public and private schools: How management and funding relate to their socio-economic profile. *OECD Publishing*.
- Pápai, S. (2000). Strategyproof assignment by hierarchical exchange. *Econometrica*, 68:1403–1433.
- Pathak, P. A. and Sönmez, T. (2008). Leveling the playing field: Sincere and strategic players in the Boston mechanism. *American Economic Review*, 98:1636–1652.

- Pathak, P. A. and Sönmez, T. (2013). School admissions reform in Chicago and England: Comparing mechanisms by their vulnerability to manipulation. *American Economic Review*, 103:80–106.
- Pycia, M. and Ünver, U. M. (2017). Incentive compatible allocation and exchange of discrete resources. *Theoretical Economics*, 12:287–329.
- Romero-Medina, A. (1998). Implementation of stable solutions in a restricted matching market. *Review of Economic Design*, 3:137–147.
- Roth, A. E. (1982). The economics of matching: Stability and incentives. *Mathematics of Operations Research*, 7:617–628.
- Roth, A. E. (1984a). The evolution of the labor market for medical interns and residents: A case study in game theory. *Journal of Political Economy*, 92:991–1016.
- Roth, A. E. (1984b). Stability and polarization of interests in job matching. *Econometrica*, 52:47–57.
- Roth, A. E. (1985a). The college admissions problem is not equivalent to the marriage problem. *Journal of Economic Theory*, 36:277–288.
- Roth, A. E. (1985b). Common and conflicting interests in two-sided matching markets. *European Economic Review*, 27:75–96.
- Roth, A. E. (1986). On the allocation of residents to rural hospitals: A general property of two-sided matching markets. *Econometrica*, 54:425–427.
- Roth, A. E. (1991). A natural experiment in the organization of entry level labor markets: Regional markets for new physicians and surgeons in the U.K. *American Economic Review*, 81:415–440.
- Roth, A. E. (2002). The economist as engineer: Game theory, experimentation, and computation as tools for design economics. *Econometrica*, 70:1341–1378.
- Roth, A. E. (2008). Deferred acceptance algorithms: History, theory, practice, and open questions. *International Journal of Game Theory*, 537–569:36.
- Roth, A. E. (2018). Marketplaces, markets, and market design. *American Economic Review*, 108:1608–1658.
- Roth, A. E. and Peranson, E. (1997). The effects of the change in the NRMP matching algorithm. *Journal of the American Medical Association*, 278:729–732.
- Roth, A. E. and Postlewaite, A. (1977). Weak versus strong domination in a market with indivisible goods. *Journal of Mathematical Economics*, 4:131–137.

- Roth, A. E. and Sotomayor, M. (1990). *Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis*. Econometric Society Monograph Series, Cambridge University Press.
- Shapley, L. and Scarf, H. (1974). On cores and indivisibility. *Journal of Mathematical Economics*, 1:23–28.
- Sönmez, T. (1997). Manipulation via capacities in two-sided matching markets. *Journal of Economic Theory*, 77:197–204.
- Sönmez, T. and Ünver, M. U. (2011). Matching, allocation, and exchange of discrete resources. In Benhabib, J., Bisin, A., and Jackson, M., editors, *Handbook of Social Economics*, volume 1A, pages 781–852. The Netherlands: North-Holland.
- Sönmez, T. and Yenmez, M. B. (2021). Can economic theory be informative for the judiciary? affirmative action in India via vertical and horizontal reservations. *Working Paper*.
- Svensson, L.-G. (1994). Queue allocation of indivisible goods. *Social Choice and Welfare*, 11:323–330.
- Turhan, B. (2019). Welfare and incentives in partitioned school choice markets. *Games and Economic Behavior*, 113:199–208.
- Walters, C. R. (2018). The demand for effective charter schools. *Journal of Political Economy*, 126:2179–2223.
- Westkamp, A. (2013). An analysis of the German university admissions system. *Economic Theory*, 53:561–589.
- Whitehurst, G. J. R. (2017). Education choice and competition index 2016 summary and commentary. https://www.brookings.edu/wp-content/uploads/2017/03/ccf20170329_cci_full_report.pdf.
- Zhou, L. (1990). On a conjecture by Gale about one-sided matching problems. *Journal of Economic Theory*, 52:123–135.