

Dynamic Matching in School Choice: Efficient Seat Reassignment After Late Cancellations

Itai Feigenbaum,^a Yash Kanoria,^b Irene Lo,^c Jay Sethuraman^d

^aLehman College and the Graduate Center, City University of New York, Bronx, New York 10468; ^bDecision, Risk, and Operations Division, Columbia Business School, New York, New York 10027; ^cDepartment of Management Science and Engineering, Stanford University, Stanford, California 94305; ^dDepartment of Industrial Engineering and Operations Research, Columbia University, New York, New York 10027

Contact: itai.feigenbaum@lehman.cuny.edu (IF); yanoria@columbia.edu,  <http://orcid.org/0000-0002-7221-357X> (YK); ilo@stanford.edu,  <http://orcid.org/0000-0002-0678-3494> (IL); jay@ieor.columbia.edu,  <http://orcid.org/0000-0002-9985-0683> (JS)

Received: May 1, 2017

Revised: July 28, 2018; February 7, 2019; July 15, 2019

Accepted: July 27, 2019

Published Online in Articles in Advance: May 8, 2020

<https://doi.org/10.1287/mnsc.2019.3469>

Copyright: © 2020 INFORMS

Abstract. In the school choice market, where scarce public school seats are assigned to students, a key operational issue is how to reassign seats that are vacated after an initial round of centralized assignment. Practical solutions to the reassignment problem must be simple to implement, truthful, and efficient while also alleviating costly student movement between schools. We propose and axiomatically justify a class of reassignment mechanisms, the permuted lottery deferred acceptance (PLDA) mechanisms. Our mechanisms generalize the commonly used deferred acceptance (DA) school choice mechanism to a two-round setting and retain its desirable incentive and efficiency properties. School choice systems typically run DA with a lottery number assigned to each student to break ties in school priorities. We show that under natural conditions on demand, the second-round tie-breaking lottery can be correlated arbitrarily with that of the first round without affecting allocative welfare and that *reversing* the lottery order between rounds minimizes reassignment among all PLDA mechanisms. Empirical investigations based on data from New York City high school admissions support our theoretical findings.

Funding: Y. Kanoria was supported by the National Science Foundation's Division of Civil, Mechanical, and Manufacturing Innovation [Grant CMMI-1201045]. J. Sethuraman was supported by the National Science Foundation [Grant CMMI-1201045].

History: Accepted by Gad Allon, operations management.

Supplemental Material: The online appendix is available at <https://doi.org/10.1287/mnsc.2019.3469>.

Keywords: dynamic matching • matching markets • school choice • deferred acceptance • tie breaking • cancellations • reassessments

1. Introduction

In public school systems throughout the United States, students submit preferences over the schools for which they are eligible for admission. Because this occurs fairly early in the school year, students typically do not know their options outside the public school system when submitting their preferences. Consequently, a significant fraction of students who are allotted a seat in a public school eventually do not use it, leading to considerable inefficiency. In the New York City (NYC) public high school system, over 80,000 students are assigned to a public school each year in March, and about 10% of these students choose to not attend a public school in September, possibly opting instead to attend a private or charter school.¹ Schools find out about many of the vacated seats only after classes begin, when students do not show up to class; such seats are reassigned in an ad hoc manner by the schools using decentralized procedures that can run months into the school year. A well-designed reassignment process, run after students learn about their outside options, could lead to significant gains in overall welfare.

Yet no systematic way of reassigning students to unused seats has been proposed in the literature. Our goal is to design an explicit reassignment mechanism run at a late stage of the matching process that efficiently reassesses students to vacated seats.

During the past 15 years, insights from matching theory have informed the design of school choice programs in cities around the world. The formal study of this mechanism design approach to school choice originated in a paper by Abdulkadiroglu and Sönmez (2003). They formulated a model in which students have strict preferences over a finite set of schools, each with a given capacity, and each school partitions the set of students into priority groups. There is now a vast and growing literature that explores many aspects of school choice systems and informs how they are designed in practice. However, most models considered in this literature are essentially static. Incorporating dynamic considerations in designing assignment mechanisms, such as students learning new information at an intermediate time, is an important aspect that has only recently started to be addressed.

Our work provides some initial theoretical results in this area and suggests that simple adaptations of one-shot mechanisms can work well in a more general setting.

We consider a two-round model of school assignment with finitely many schools. Students learn their outside option after the first-round assignment and vacate seats that can be reassigned. In the first round, schools have weak priorities over students, and students submit strict ordinal preferences over schools. Students receive a first-round assignment based on these preferences via deferred acceptance with single tie breaking (DA-STB), a variant of the standard deferred acceptance (DA) mechanism, where ties in school preferences are broken via a single lottery ordering across all schools. Afterward, students learn their outside options (such as admission to a private school) and may no longer be interested in the seat allotted to them. In the second round, students are invited to submit new ordinal preferences over schools, reflecting changes in their preferences induced by learning their outside options. The goal is to reassign students so that the resulting assignment is efficient and the two-round mechanism is strategy-proof and does not penalize students for participating in the second round. Because a significant fraction of seats available for reassignment are vacated only after the start of the school year, a key additional goal is to ensure that the reassignment process minimizes the number of students who need to be reassigned.


We introduce a class of reassignment mechanisms with desirable properties: the *permuted lottery deferred acceptance (PLDA) mechanisms*. PLDA mechanisms compute a first-round assignment by running DA-STB and then a second-round assignment by running DA-STB with a permuted lottery. In the second round, each school first prioritizes students who were

assigned to it in the first round, which guarantees each student a second-round assignment that he or she prefers to his or her first-round assignment, then prioritizes students according to their initial priorities at the school, and finally breaks ties at all schools via a permutation of the (first-round) lottery numbers. Our proposed PLDA mechanisms are based on school choice mechanisms currently implemented in the main round of assignments and can be implemented either as *centralized PLDAs*, which run a centralized second round with updated preferences, or as *decentralized PLDAs*, which run a decentralized second round via a waitlist system that closely mirrors current reassignment processes.

Our key insight is that the mechanism designer can design the correlation between tie-breaking lotteries to achieve operational goals. In particular, *reversing* the lottery between rounds minimizes reassignment without sacrificing student welfare. Our main theoretical result is that under an intuitive *order condition*, all PLDAs produce the same distribution over the final assignment, and reversing tie-breaking lotteries between rounds implements the centralized reverse lottery DA (RLDA), which minimizes the number of reassigned students. We axiomatically justify PLDA mechanisms: absent school priorities, PLDAs are equivalent to the class of mechanisms that are two-round strategy-proof while satisfying natural efficiency requirements and symmetry properties.

In a setting where all students agree on a ranking of schools and there are no priorities, our results are very intuitive. By reversing the lottery, we move a few students many schools up their preference list rather than many students a few schools up, thereby eliminating unnecessary cascades of reassignment (see Figure 1). Surprisingly, however, our theoretical result holds in

Figure 1. Running DA with a Reversed Lottery Eliminates the Cascade of Reassignments



Notes. There are six students with identical preferences over schools and six schools each with a single-priority group. All students prefer schools in the order $s_1 > s_2 > \dots > s_6$. The student assigned to school s_1 in the first round leaves after the first round; otherwise, all students find all schools acceptable in both rounds. Running DA with the same tie-breaking lottery reassigned each student to the school that is one better on his or her preference list, whereas reversing the tie-breaking lottery reassigned only the student initially assigned to s_6 .

a general setting with heterogeneous student preferences and arbitrary priorities at schools. The order condition can be interpreted as aggregate student preferences resulting in the same *order* of popularity of schools in the two rounds. Our results show that if student preferences and school priorities produce such agreement in aggregate demand across the two rounds, then reversing the lottery between rounds preserves ex ante allocative efficiency and minimizes reassignment.

We empirically assess the performance of RLDA using data from the NYC public high school system. We first investigate a class of centralized PLDAs that includes RLDA, rerunning DA using the original lottery order (termed forward lottery deferred acceptance (FLDA)), and rerunning DA using an independent random lottery. We find that all these mechanisms provide similar allocative efficiency, but RLDA reduces the number of reassigned students significantly. For instance, in the NYC data set from 2004–2005, we find that FLDA results in about 7,800 reassessments and RLDA results in about 3,400 reassessments out of a total of about 74,000 students who remained in the public school system—that is, fewer than half the number of reassessments under FLDA. The gains become even more marked if we compare with current practice: RLDA results in fewer than 40% of the 8,600 reassessments under decentralized FLDA with waitlists.²

To better evaluate the currently used waitlist systems, we also empirically explore the performance of decentralized FLDA and RLDA as a function of the time available to clear the market. We find that the timing of information revelation can greatly impact both allocative efficiency and congestion. If congestion is caused by students taking time to vacate previously assigned seats (see Figure 1), then reversing the lottery increases allocative efficiency during the early stages of reassignment and decreases congestion. However, if congestion is caused by students taking time to decide on waitlist offers, these findings are reversed. In both cases, for reasonable timescales, the welfare gains from centralizing the system and reducing congestion can be substantial.

The rest of this paper is organized as follows. We outline current practice in school admissions in Section 1.1 and related literature in Section 1.2. In Section 2, we describe our model, our proposed PLDA mechanisms, and their properties. Section 3 presents our main results, and Section 4 provides intuition and a flavor of our analysis via a special case of our model. We provide empirical results in Section 5 and conclude in Section 6.

1.1. Current Practice

School systems in cities across the United States use similar centralized processes for admissions to public

schools. Students seeking admission to a school submit their preferences over schools to a central authority by December through March, for admission starting the subsequent fall. Each school may have priority classes of students, such as priority for students who live in the neighborhood, siblings of students already enrolled at that school, or students from low-income families. An assignment of seats to students is produced using the student-proposing DA algorithm with single tie breaking. Students must register in their assigned school by April or early May.

In March and April, students are also admitted to private and charter schools via processes run concurrently with the public school assignment process. This results in an attrition rate of about 8%–10% of the seats assigned in the main round of public school admissions. Some schools account for this attrition by accepting more students in the first round than they have seats for. However, such oversubscription of students is usually conservative because of hard constraints on space and teacher capacity.³ As a result, most schools have unused seats at the end of the first round that can be reassigned. In addition, most public schools find out about many of these vacant seats only after the start of the school year because they cannot require deposits or other forms of commitment from students before the start of the school year.

Reassignments in most school choice systems are performed using a decentralized waitlist system.⁴ Students are put on waitlists for all schools that they ranked above their first-round assignment and ordered by first-round priorities (after tie breaking). Students who do not register by the deadline are presumed to be uninterested, and their seats are offered to waitlisted students in sequence, with more seats becoming available over time as students receive new offers from outside the system or are reassigned via waitlists to other public school seats. Students offered seats by the waitlist system usually have just under a week to make a decision and are only bound by the final offer they choose to accept.⁵ Overall, this typically results in a drawn-out reassignment process that continues all summer until after classes begin and in some cities (e.g., NYC kindergarten, Boston, and Washington, DC) up to several months after the start of the school year.

Our proposed class of mechanisms generalizes these waitlist systems as follows. Waitlists are PLDA mechanisms where (1) the second round is implemented in a decentralized fashion as information about vacated seats propagates through the system, and (2) the tie-breaking lotteries used in the two rounds are the same. We show that *permuting* the tie-breaking lottery numbers before creating waitlists provides a class of reassignment mechanisms that, given sufficient time, results in similar allocative

efficiency while providing flexibility for optimizing other objectives.

1.2. Related Work

The mechanism design approach to school choice was first formulated by Balinski and Sönmez (1999) and Abdulkadiroglu and Sönmez (2003). Since then, academics have worked closely with school authorities to redesign school choice systems to increase student welfare.⁶ A significant portion of the literature has focused on providing solutions for a single round of centralized school assignment—see, for example, Pathak (2011) and Abdulkadiroglu and Sönmez (2011) for recent surveys. Many of these works provide axiomatic justifications for two canonical mechanisms—DA (Gale and Shapley 1962) and top trading cycles (TTCs)—and their variants, in terms of their desirable properties. We provide a similar framework for the reassignment problem by proposing and characterizing PLDA mechanisms by their incentive and efficiency properties.

There is a growing operations literature on designing the school choice process to optimize quantitative objectives. Ashlagi and Shi (2014) consider how to improve community cohesion in school choice by correlating the lotteries of students in the same community, and Ashlagi and Shi (2015) show how to maximize welfare given busing cost constraints. Several papers also explore how school districts can use rules for breaking ties in school priorities as policy levers. Arnosti (2015), Ashlagi and Nikzad (2016), and Ashlagi et al. (2015) show that DA-STB assigns more students to one of their top k schools (for small k) compared with DA using independent lotteries at different schools, and Abdulkadiroglu et al. (2009) empirically compare these tie-breaking rules. Erdil and Ergin (2008) also exploit indifferences to improve allocative efficiency. We explore the design of tie-breaking rules in the reassignment setting and correlate tie breaking across rounds.

There is also a vast literature on dynamic matching and reassignments. The reassignment of donated organs has been extensively studied in work on kidney exchange (see, e.g., Roth et al. 2004; Anderson et al. 2015, 2017; Ashlagi et al. 2019). Reassignments resulting from cancellations also arise in online assignment settings such as kidney transplantation (see, e.g., Zenios 1999, Su and Zenios 2006) and public housing allocation (see, e.g., Kaplan 1987, Arnosti and Shi 2017). An important difference is that these are *online* settings where agents and objects arrive over time and are matched on an ongoing basis. In such settings, matches are typically irrevocable, so optimal assignment policies account for typical cancellation and arrival statistics and optimize for agents arriving in the future (see, e.g., Dickerson and Sandholm

2015). In our setting, the matching for the entire system is coordinated in time, and we improve welfare by controlling both the initial assignment and subsequent reassignment of objects among the same set of agents.

Another relevant strand in the reassignment literature is the work of Abdulkadiroglu and Sönmez (1999) on house allocation models with housing endowments. Our second round can be thought of as school seat allocation where some agents already own a seat and we wish to reassign seats to reach an efficient assignment. There are also a growing number of papers that consider a dynamic model for school admissions (see e.g., Compte and Jehiel 2008, Combe et al. 2016). A critical distinction between these works and ours is that in our model, the initial endowment is determined endogenously by preferences, so we propose reassignment mechanisms that are impervious to students manipulating their first-round endowment to improve their final assignment.

A number of recent papers, such as those by Dur (2012), Pereyra (2013), and Kadam and Kotowski (2014), focus on the strategic issues in dynamic reassignment. These works develop solution concepts in finite markets with specific cross-period constraints and propose DA-like mechanisms that implement them. In recent complementary work, Narita (2016) analyzes preference data from NYC school choice, observes that a significant fraction of preferences is permuted after the initial match, and proposes a modified version of DA with desirable properties in this setting. We similarly propose PLDA mechanisms for their desirable incentive and efficiency properties. In addition, our large market and consistency assumptions allow us to uncover considerable structure in the problem and provide conditions under which we can optimize over the entire class of PLDA mechanisms.

Our work also has some connections to the queuing literature. The class of mechanisms that emerges in our setting involves choosing a permutation of the initial lottery order, and we find that the reverse lottery minimizes reassignment within this class. This is similar to choosing a service policy in a queuing system (e.g., first in first out, last in first out, shortest remaining time first, etc.) in order to minimize cost functions (see, e.g., Lee and Srinivasan 1989). “Work-conserving” service policies can result in identical throughput but different expected waiting times, and we similarly find that PLDA mechanisms may have identical allocative efficiency but different numbers of reassignments. Our continuum model parallels fluid limits and deterministic models employed in queuing (Whitt 2002), revenue management (Talluri and Van Ryzin 2006), and other contexts in operations management.

2. Model

2.1. Definitions and Notation

We consider the problem of assigning a set of students Λ to seats in a finite set of schools $S = \{s_1, \dots, s_N\}$. Each student can attend at most one school. There is a *continuum*⁷ of students with an associated measure η : for any (measurable) subset $A \subseteq \Lambda$, we use $\eta(A)$ to denote the mass of students in A . The outside option is $s_{N+1} \notin S$. The capacities of the schools are $q_1, \dots, q_N \in \mathbb{R}_+$ and $q_{N+1} = \infty$. A set of students of η -measure at most q_i can be assigned to school s_i .

Each student submits a strict preference ordering over his or her acceptable schools, and each school partitions eligible students into priority groups. Each student has a *type* $\theta = (\succ^\theta, \hat{\succ}^\theta, p^\theta)$ that encapsulates both his or her preferences and school priorities. The student's first- and second-round preferences, respectively, \succ^θ and $\hat{\succ}^\theta$ are strict ordinal preferences over $S \cup \{s_{N+1}\}$, and schools before (after) s_{N+1} in the ordering are acceptable (unacceptable). We think of s_{N+1} as the best guaranteed outside option available to the student, with the understanding that it can "improve" from the first to the second round—for example, because a new private school offer comes in. The student's *priority class* p^θ encodes his or her priority p_i^θ at each school s_i . Each school s_i has n_i priority groups. We assume that schools prefer higher priority groups, that students ineligible for school s_i have priority $p_i = -1$, and that $p_i \in \{-1, 0, 1, \dots, n_i - 1\}$. Eligibility and priority groups are exogenously determined and publicly known. Each student $\lambda = (\theta^\lambda, L(\lambda)) \in \Lambda$ also has a first-round lottery number $L(\lambda) \in [0, 1]$. We sometimes use the notation $(\succ^\lambda, \hat{\succ}^\lambda, p^\lambda)$ as a less cumbersome alternative to $(\succ^{\theta^\lambda}, \hat{\succ}^{\theta^\lambda}, p^{\theta^\lambda})$. Let Θ be the set of all student types so that $\Lambda = \Theta \times [0, 1]$ denotes the set of students. For each $\theta \in \Theta$, let $\zeta(\theta) = \eta(\{\theta\} \times [0, 1])$ be the measure of all students with type θ .

We assume that all students have consistent preferences, defined as follows.

Definition 1. Preferences $(\succ, \hat{\succ})$ are *consistent* if the second-round preferences $\hat{\succ}$ are obtained from \succ via truncation—that is, (a) schools do not become acceptable only in the second round, $\forall s_i \in S \ s_i \hat{\succ} s_{N+1}$ implies $s_i \succ s_{N+1}$, and (b) the relative ranking of schools is unchanged across rounds, $\forall s_i, s_j \in S \text{ if } s_i \hat{\succ} s_{N+1} \text{ and } s_i \hat{\succ} s_j, \text{ then } s_i \succ s_j$. Type θ is *consistent* if $(\succ^\theta, \hat{\succ}^\theta)$ are consistent.

Assumption 1 (Consistent Preferences). *If $\zeta(\theta) > 0$, then the type θ is consistent.*

In other words, we assume that the only change that a student makes to his or her preferences between the two rounds is truncating his or her preference list. The consistency assumption allows us to study the effects

of students leaving the public school system but precludes the possibility of students revising their relative ranking of schools between rounds—for example, as a result of making mistakes in the first round—or obtaining more information about the schools. It also allows us to propose strategy-proof mechanisms⁸ and to optimize for allocative efficiency and reassignment. In settings where the consistency assumption does not hold, we note that the mechanisms we will propose are still well defined, and we will show that they retain many of their desirable properties.

We also assume that consistent types have full support. We use this only to characterize our proposed mechanisms (Theorem 3) and do not need it for our positive results (Theorems 1 and 2).


Assumption 2 (Full Support). *For all consistent types $\theta \in \Theta$, it holds that $\zeta(\theta) > 0$.*

Finally, we assume that the first-round lottery numbers are drawn independently and uniformly from $[0, 1]$ and do not depend on preferences: $\eta(\{\theta\} \times (a, b)) = (b - a)\zeta(\theta) \ \forall \theta \in \Theta, 0 \leq a \leq b \leq 1$.⁹

An *assignment* $\mu : \Lambda \rightarrow S$ specifies the school to which each student is assigned. For any assignment μ , we let $\mu(\lambda)$ denote the school to which student λ is assigned, and in a slight abuse of notation, we let $\mu(s_i)$ denote the set of students assigned to school s_i . We assume that $\mu(s_i)$ is η -measurable and that the assignment is *feasible*—that is, $\eta(\mu(s_i)) \leq q_i$ for all $s_i \in S$, and if $\mu(\lambda) = s_i$, then $p_i^\lambda \geq 0$. We let μ and $\hat{\mu}$ denote the first- and second-round assignments, respectively.

2.1.1. Timeline. Students report first-round preferences \succ . The mechanism designer obtains a first-round assignment μ by running DA-STB with lottery L and announces μ and L . Students then observe their outside options and update their preferences to $\hat{\succ}$. Finally, students report their second-round preference $\hat{\succ}$, and the mechanism designer obtains a second-round assignment $\hat{\mu}$ by running a reassignment mechanism M and announces $\hat{\mu}$. We illustrate the timeline in Figure 2.

2.1.2. Informational Assumptions. Eligibility and priorities are exogenously determined and publicly known. The mechanism is publicly announced before preferences are submitted. Before first-round reporting, each student knows his or her first-round preferences and that his or her second-round preferences will be obtained from these preferences via truncation. Each student has imperfect information regarding his or her own second-round preferences (i.e., the point of truncation) at that stage and believes with positive probability that his or her preferences in both rounds will be identical.¹⁰ We assume that students know the distribution η over student types and lotteries (an

Figure 2. Timeline of the Two-Round Mechanism Design Problem


assumption we need only for our characterization result, Theorem 3). Each student is assumed to learn his or her lottery number after the first round because, in practice, students are often permitted to inquire about their position on each school's waitlist; our results hold even if students do not learn their lottery numbers.

Definition 2. A student $\lambda \in \Lambda$ is a *reassigned student* if he or she is assigned to a different school in S in the second round than in the first round. That is, λ is a reassigned student under reassignment $\hat{\mu}$ if $\mu(\lambda) \neq \hat{\mu}(\lambda)$ and¹¹ $\mu(\lambda) \neq s_{N+1}$, $\hat{\mu}(\lambda) \neq s_{N+1}$.

Most reassessments happen around the start of the school year, a time when they are costly for schools and students alike. Hence, in addition to providing an efficient final assignment, we also want to reduce congestion by minimizing the number of reassigned students.

2.2. Mechanisms

A *mechanism* is a function that maps the realization of first-round lotteries L , school priorities p , and students' first-round preference reports $>$ into an assignment μ . A *reassignment mechanism* is a function that maps the realization of first-round lotteries L , first-round assignment μ , school priorities p , and students' second-round reports $\hat{>}$ into a second-round assignment¹² $\hat{\mu}$. A *two-round mechanism* obtained from a reassignment mechanism M is a two-round mechanism where the first-round mechanism is DA-STB (see Definition 3) and the second-round mechanism is M .

In the first round, seats are assigned according to the student-optimal DA algorithm with single tie breaking (STB) as follows. A *single* lottery ordering of the students L is used to resolve ties in the priority groups at all schools, resulting in an instance of the two-sided matching problem with strict preferences and priorities. In each step of DA, unassigned students apply to their most-preferred school that has not yet rejected them. A school with a capacity of q *tentatively accepts* the q highest-ranked eligible applicants (according to its priority ranking of students after breaking all ties) and rejects any remaining applicants. The algorithm runs until there are no new

student applications, at which point it terminates and assigns each student to his or her tentatively assigned school seat. The strict student preferences, weak school priorities, and use of DA-STB mirror current practice in many school choice systems, such as those in NYC, Chicago, and Denver (see, e.g., Abdulkadiroglu and Sönmez 2003).

DA can also be formally defined in terms of admissions scores and cutoffs.

Definition 3 (Deferred Acceptance; Azevedo and Leshno 2016). The DA mechanism with single tie breaking (DA-STB) is a function $DA_\eta((>^\lambda, p^\lambda)_{\lambda \in \Lambda}, L)$ mapping student preferences, priorities, and lottery numbers into an assignment μ , defined by a vector of cutoffs $\mathbf{C} \in \mathbb{R}_+^N$ as follows. Each student λ is given a score $r_i^\lambda = p_i^\lambda + L(\lambda)$ at school s_i and is assigned to his or her most-preferred school as per his or her preferences, among those where his or her score exceeds the cutoff:

$$\mu(\lambda) = \max_{>^\lambda} \left(\left\{ s_i \in S : r_i^\lambda \geq C_i \right\} \cup \{s_{N+1}\} \right). \quad (1)$$

Moreover, \mathbf{C} is market clearing—namely,

$$\eta(\mu(s_i)) \leq q_i, \text{ for all } s_i \in S, \text{ with equality if } C_i > 0. \quad (2)$$

Azevedo and Leshno (2016) showed that the set of assignments satisfying Equations (1) and (2) forms a nonempty complete lattice and typically consists of a single uniquely determined assignment.¹³ This unique assignment in the continuum further corresponds to the scaling limit of the set of stable matches obtained in finite markets as the number of students grows (with school capacities growing proportionally). Throughout this paper, in the (knife-edge) case where there are multiple assignments satisfying Definition 3, we pick the student-optimal matching.

Given cutoffs $\{C_i\}_{i=1}^N$, we will also find it helpful to define for each priority class π the *cutoffs within the priority class* at each school $C_{\pi,i} \in [0, 1]$ by $C_{\pi,i} = 0$ if $C_i \leq \pi_i$, $C_{\pi,i} = 1$ if $C_i \geq \pi_i + 1$, and $C_{\pi,i} = C_i - \pi_i$ otherwise. Thus, $C_{\pi,i}$ is the lowest lottery number a student θ with priority $p^\theta = \pi$ can have and still be able to attend school s_i .

We now turn to the mechanism design problem. We emphasize that we consider only two-round mechanisms whose first-round mechanism is the currently used DA-STB—that is, the only freedom afforded the planner is the design of the reassignment mechanism. We propose the following class of two-round mechanisms. Intuitively, these mechanisms run DA-STB twice, once in each round. They explicitly correlate the lotteries used in the two rounds via a permutation P and in the second round give each student top priority in the school to which he or she was assigned in the first round to guarantee that each student receives a (weakly) better assignment.

Definition 4 (Permuted Lottery Deferred Acceptance (PLDA) Mechanisms). Let $P : [0, 1] \rightarrow [0, 1]$ be a measure-preserving bijection, let L be the realization of first-round lottery numbers, and let μ be the first-round assignment obtained by running DA with lottery L . Define a new economy $\hat{\eta}$, where to each student $\lambda \in \Lambda$ with priority vector p^λ and first-round lottery and assignment $L(\lambda) = l$, $\mu(\lambda) = s_i$, we (a) assign a lottery number $P(l)$ and (b) give top second-round priority $\hat{p}_i^\lambda = n_i$ at their first-round assignment s_i and unchanged priority $\hat{p}_j^\lambda = p_j^\lambda$ at all other schools $s_j \neq s_i$. $PLDA(P)$ is the two-round mechanism obtained using the reassignment mechanism $DA_{\hat{\eta}}((\hat{\lesssim}^\lambda, \hat{p}^\lambda)_{\lambda \in \Lambda}, P \circ L)$.

We use $\hat{C}_{\pi, i}^P$ to denote the second-round cutoff for priority class π in school i under $PLDA(P)$.

We highlight two particular PLDA mechanisms. The RLDA (reverse lottery) mechanism uses the reverse permutation $R(x) = 1 - x$, and the FLDA (forward lottery) mechanism, which preserves the original lottery order, uses the identity permutation $F(x) = x$. By default, school districts often use a decentralized version of the FLDA mechanism, implemented via waitlists. In this paper, we provide evidence that supports using the *centralized RLDA mechanism* in a school system such as that in NYC, where a large proportion of vacated seats is revealed close to or after the start of the school year and where reassignments are costly for both students and the school administration.

The PLDA mechanisms are an attractive class of two-round assignment mechanisms for a number of reasons. They are intuitive to understand and simple to implement in systems already using DA. (A decentralized implementation would be even simpler to integrate with current practice; the currently used waitlist mechanism for reassignments can be retained with the simple modification of permuting the lottery numbers just before waitlists are constructed.) In addition, we will show that the PLDA mechanisms have desirable incentive and efficiency properties, which we now describe.

Any reassignment mechanism that takes away a student's initial assignment against his or her will is

impractical. Thus, we require our mechanism to respect first-round guarantees:

Definition 5. A two-round mechanism (or a second-round assignment $\hat{\mu}$) *respects guarantees* if every student (weakly) prefers his or her second-round assignment to his or her first-round assignment—that is, $\hat{\mu}(\lambda) \hat{\lesssim}^\lambda \mu(\lambda)$ for every $\lambda \in \Lambda$.

One of the reasons for the success of DA in practice is that it respects priorities: if a student is not assigned to a school he or she wants, it is because that school is filled with students with higher priority at that school. This leads to the following natural requirement in our two-round context.

Definition 6. A two-round mechanism (or a second-round assignment $\hat{\mu}$) *respects priorities (subject to guarantees)* if $\forall s_i \in S$, every eligible student $\lambda \in \Lambda$ such that $s_i \hat{\lesssim}^\lambda \hat{\mu}(\lambda)$, and every student λ' such that $\hat{\mu}(\lambda') = s_i \neq \mu(\lambda')$ it holds that λ' is eligible for s_i and $p_i^{\lambda'} \geq p_i^\lambda$.

Thus, our definition of respecting priorities (subject to guarantees) requires that every student who was upgraded to a school s in the second-round must have a (weakly) higher priority at that school than every eligible student λ who prefers s to his or her second-round assignment.

We now turn to incentive properties. In the school choice problem, it is reasonable to assume that students will be strategic in how they interact with the mechanism at each stage. Hence, it is desirable that whenever a student (with consistent preferences) reports preferences, conditional on everything that has happened up to that point, it is a dominant strategy for him or her to report truthfully. To describe the properties formally, we start by fixing an arbitrary profile of first- and second-round preferences $(\succ^{-\lambda}, \lesssim^{-\lambda})$ for all the students other than student λ . For any preference report of student λ in the first round, he or she will receive an assignment that is probabilistic because of the tie-breaking lottery; then, after observing his or her first-round assignment and lottery number and his or her updated outside option, he or she can submit a second-round preference report, based on which his or her final assignment is computed. This leads to two natural notions of strategy-proofness.

Definition 7. A two-round mechanism is *strongly strategy-proof* if for each student λ (with consistent preferences) truthful reporting is a dominant strategy—that is, for each realization of lottery numbers (including his or her own lottery number) and profile of first- and second-round reported preferences of the students other than λ , reporting his or her preferences truthfully in each of the two rounds is a best response for student λ .

Our definition of *strong strategy-proofness* is rather demanding: it requires that no student be able to

manipulate the mechanism even if he or she has full knowledge of the first- and second-round preferences of *all* other students and the lottery numbers. We shall also consider a weaker version of strategy-proofness that applies when a manipulating student does not know the lottery number realizations when he or she submits his or her first-round preference report and learns all lottery numbers only after the end of the first round. In that case, each student views his or her first-round assignment as a probability vector; his or her second-round assignment is also random but is a deterministic function of the first-round outcome, the second-round reports, and the first-round lottery numbers. We make precise the notion of a successful manipulation in this setting as follows.

Definition 8. A two-round mechanism is *weakly strategy-proof* if the following conditions hold:

- Knowing the specific realization of first-round assignments (and lottery numbers) and the second-round preferences of the students other than λ , it is optimal for student λ to submit his or her second-round preference truthfully, given what the other students do.
- For each student λ (with consistent preferences), and for each profile of first- and second-round preferences of the students other than λ , the probability that student λ is assigned to one of his or her top k schools in the second round is maximized when he or she reports truthfully in the first round (assuming truthful reporting in the second round), for each $k = 1, 2, \dots, N$.

In other words, in each stage of the dynamic game, the outcome from truthful reporting stochastically dominates the outcomes of all other strategies. We emphasize that the uncertainty in the first-round assignment is solely due to the lottery numbers, which students initially do not know.

Note that a two-round mechanism that uses the first-round assignment as the initial endowment for a mechanism such as top trading cycles in the second round will *not* be two-round strategy-proof because students can benefit from manipulating their first-round reports to obtain a more popular initial assignment that they could use to their advantage in the second round.

Finally, we discuss some efficiency properties. To be efficient, clearly a mechanism should not leave unused any seats that are desired by students.

Definition 9. A two-round mechanism is *nonwasteful* if no student is assigned to a school for which he or she is eligible that he or she prefers less than a school not at capacity; that is, for each student $\lambda \in \Lambda$ and schools s_i, s_j , if $\hat{\mu}(\lambda) = s_i$ and $s_j \succ^\lambda s_i$ and $p_j^\lambda \geq 0$, then $\eta(\hat{\mu}(s_j)) = q_j$.

It is also desirable for a two-round mechanism to be Pareto efficient. We do not want any students to be

able to improve their utility by swapping probability shares in second-round assignments. However, we also require that our reassignment mechanism respect guarantees and priorities (see Definitions 5 and 6), which is incompatible with Pareto efficiency even in a static, one-round setting.¹⁴ This motivates the following definitions. Consider a second-round assignment $\hat{\mu}$. A *Pareto-improving cycle* is an ordered set of types $(\theta_1, \theta_2, \dots, \theta_m) \in \Theta^m$, sets of students $(\Lambda_1, \Lambda_2, \dots, \Lambda_m) \in \Lambda^m$, and schools $(\tilde{s}_1, \tilde{s}_2, \dots, \tilde{s}_m) \in S^m$ such that $\eta(\Lambda_i) > 0$ and $\tilde{s}_{i+1} \succ^{\theta_i} \tilde{s}_i$ (where we define $\tilde{s}_{m+1} = \tilde{s}_1$), for all i , and such that for each i , $\theta^{\lambda} = \theta_i$, $\hat{\mu}(\lambda) = \tilde{s}_i$ for all $\lambda \in \Lambda_i$.

Let \hat{p} be the second-round priorities obtained by giving each student λ a top second-round priority $\hat{p}_i^\lambda = n_i$ at their first-round assignment $\mu(\lambda) = s_i$ (if $s_i \in S$) and unchanged priority $\hat{p}_j^\lambda = p_j^\lambda$ at all other schools $s_j \neq s_i$. We say that a Pareto-improving cycle (in a second-round assignment) *respects (second-round) priorities* if $\hat{p}_{\tilde{s}_{i+1}}^{\theta_i} \geq \hat{p}_{\tilde{s}_{i+1}}^{\theta_{i+1}}$ for all i (where we define $\theta_{m+1} = \theta_1$).

Definition 10. A two-round mechanism is *constrained Pareto efficient* if the second-round assignment has no Pareto-improving cycles that respect second-round priorities.

We remark that this is the same notion of efficiency that is satisfied by static, single-round DA-STB (Definition 3)—the resulting assignment has no Pareto-improving cycles that respect priorities. In other words, the constrained Pareto-efficiency requirement is informally to be “as efficient as static DA.” We also note here that as a result of the requirement to respect second-round priorities, Pareto-improving cycles considered must include only reassigned students.

Finally, for equity purposes, it is desirable that a mechanism be anonymous.

Definition 11. A two-round mechanism is *anonymous* if students with the same first-round assignment and the same first- and second-round preference reports have the same distribution over second-round assignments.

We show that PLDA mechanisms satisfy all the aforementioned properties.

Proposition 1. PLDA mechanisms respect guarantees and priorities and are nonwasteful, constrained Pareto efficient, and anonymous. Moreover, if student preferences are consistent, PLDA mechanisms are strongly two-round strategy-proof.

We will show in Section 3.1 that in a setting without priorities, the PLDA mechanisms are the only mechanisms that satisfy all these properties (and some additional technical requirements), even if we only require weak strategy-proofness (Theorem 3).

Finally, it is simple to show that the natural counterparts to PLDA mechanisms in a discrete setting

(with a finite number of students) respect guarantees and priorities and are nonwasteful, constrained Pareto efficient, and anonymous. We make these claims formal in the online appendix and also provide an informal argument that the discrete PLDA mechanisms are also approximately strategy-proof when the number of students is large.

3. Main Results

In this section, we show that the defining characteristic of a PLDA mechanism—the permutation of lotteries between the two rounds—can be chosen to achieve desired operational goals. We first provide an intuitive *order condition* and show that, under this condition, all PLDA mechanisms give the same *ex ante* allocative efficiency. Thus, when the primitives of the market satisfy the order condition, it is possible to pursue secondary operational goals without sacrificing allocative efficiency. Next, in the context of reassigning school seats at the start of the school year, we consider the specific problem of minimizing reassignment and show that when the order condition is satisfied, *reversing* the lottery minimizes reassignment among all centralized PLDA mechanisms. In Section 5, we empirically demonstrate using data from NYC public high schools that reversing the lottery minimizes reassignment (among a subclass of centralized PLDA mechanisms) and does not significantly affect allocative efficiency, even when the order condition does not hold exactly. Our results suggest that centralized RLDA is a good choice of mechanism when the primary goal is to minimize reassignments while providing a second-round assignment with high allocative efficiency. In Section 3.1, we provide an axiomatic justification for PLDA mechanisms, and later in Section 6, we discuss how the choice of lottery permutation can be used to achieve other operational goals, such as maximizing the number of students with improved assignments.

We begin by defining the order condition, which we will need to state our main results.

Definition 12. The *order condition* holds on a set of primitives (S, q, Λ, η) if for every priority class π , the first- and second-round school cutoffs under RLDA within that priority class are in the same order—that is, for all $s_i, s_j \in S$,

$$C_{\pi,i} > C_{\pi,j} \Rightarrow \hat{C}_{\pi,i}^R \geq \hat{C}_{\pi,j}^R.$$

We emphasize that the order condition is a condition on the market primitives—namely, school capacities and priorities and student preferences (though checking whether it holds involves investigating the output of RLDA). We may interpret the order condition as an indication that the relative demand for the schools is consistent between the two rounds. Informally

speaking, this means that the revelation of the outside options does not change the *order* in which schools are overdemanded. One important setting where the order condition holds is the case of uniform dropouts and a single priority type. In this setting, each student independently with probability ρ either remains in the system and retains his or her first-round preferences in the second round or drops out of the system entirely; student first-round preferences and school capacities are arbitrary. In Section 4, we use the uniform dropouts setting to provide some intuition for our general results.

To compare the allocative efficiency of different mechanisms, we define *type equivalence* of assignments. In words, two second-round assignments are type equivalent if the masses of different student types θ assigned to each school are the same across the two assignments.

Definition 13. Two second-round assignments $\hat{\mu}$ and $\hat{\mu}'$ are said to be *type equivalent* if

$$\begin{aligned} \eta\left(\left\{\lambda \in \Lambda : \theta^\lambda = \theta, \hat{\mu}(\lambda) = s_i\right\}\right) \\ = \eta\left(\left\{\lambda \in \Lambda : \theta^\lambda = \theta, \hat{\mu}'(\lambda) = s_i\right\}\right) \forall \theta \in \Theta \text{ and } s_i \in S. \end{aligned}$$

In our continuum model, if two two-round mechanisms produce type-equivalent second-round assignments, we may equivalently interpret them as providing each individual student of type θ with the same *ex ante* distribution (before lottery numbers are assigned) over assignments.

Our first main result is the surprising finding that all PLDAs are allocatively equivalent.

Theorem 1 (Order Condition Implies Type Equivalence). *If the order condition (Definition 12) holds, all PLDA mechanisms produce type-equivalent second-round assignments.*

Thus, if the order condition holds, the measure of students of type $\theta \in \Theta$ assigned to each school in the second round is independent of the permutation P . We remark that type equivalence does not imply an equal (or similar) amount of reassignment (e.g., see Figure 1) because type equivalence depends only on the second-round assignment, whereas reassignment (Definition 2) measures the difference between the first- and second-round assignments. This brings us to our second result.

Theorem 2 (Reverse Lottery Minimizes Reassignment). *If all PLDA mechanisms produce type-equivalent second-round assignments, then RLDA minimizes the measure of re-signed students among PLDA mechanisms.*

The intuition for this result is that reversing the lottery shortcuts improvement chains and moves a few students many schools up their preference list instead of a large number of students a few schools up their preference list. In particular, we show that

RLDA never reassigned both a student of type θ into a school s_i and another student of type θ out of the same school s_i .

Proof of Theorem 2. Fix $\theta = (\succ^\theta, \lesssim^\theta, p^\theta) \in \Theta$ and $s_i \in S$. We show that among all type-equivalent mechanisms, RLDA minimizes the measure of students with type θ who were reassigned to s_i because it never reassigned both a student of type θ into a school s_i and another student of type θ out of s_i .

For every permutation P , let the measures of students with type θ leaving and entering school s_i in the second round under $\text{PLDA}(P)$ be denoted by $\ell_P = \eta(\{\lambda \in \Lambda : \theta^\lambda = \theta, \mu(\lambda) = s_i, \hat{\mu}^P(\lambda) \neq s_i\})$ and $e_P = \eta(\{\lambda \in \Lambda : \theta^\lambda = \theta, \mu(\lambda) \neq s_i, \hat{\mu}^P(\lambda) = s_i\})$, respectively. As a result of type equivalence, there is a constant c , independent of P , such that $\ell_P = e_P - c$. We will show that $e_R \leq e_P$ for all permutations P , specifically by showing that either $\ell_R = 0$ and $e_R = c$ or $e_R = 0$.

If both $e_R > 0$ and $\ell_R > 0$, then students of type θ who entered s_i in the second round of RLDA had worse first- and second-round lottery numbers than students of type θ who left s_i in the second round of RLDA, which contradicts the reversal of the lottery. Because $e_P = \ell_P + c \geq c$ and $e_P \geq 0$, this completes the proof. \square

Our results present a compelling case for using the centralized RLDA mechanism when the main goals are to achieve allocative efficiency and minimize the number of reassigned students. Theorems 1 and 2 show that when the order condition holds, centralized RLDA is unequivocally optimal in the class of PLDA mechanisms because all PLDA mechanisms give type-equivalent assignments and centralized RLDA minimizes the number of reassigned students. In addition, we remark that the order condition can be checked easily by running RLDA (e.g., on historical data).¹⁵ Finally, even if the order condition does not hold, RLDA moves as few students as possible to reach the RLDA assignment.

Next, we give examples of when the order condition holds and does not hold and illustrate the resulting implications for type equivalence. We illustrate these in Figure 3.

Example 1. There are $N = 2$ schools, each with a single priority group. School s_1 has lower capacity and is initially more overdemanded. Student preferences are such that when all students who want only s_2 drop out, the order condition holds, and when all students who want only s_1 drop out, then s_2 becomes more overdemanded under RLDA, and the order condition does not hold.

School capacities are given by $q_1 = 2, q_2 = 5$. There is measure 4 of each of the four types of first-round student preferences. Let θ_i denote the student type

that finds only school s_i acceptable, and let $\theta_{i,j}$ denote the type that finds both schools acceptable and prefers s_i to s_j . (We will define the second-round preferences of each student type below; each type will either leave the system completely or keep the same preferences.) If we run DA-STB, the first-round cutoffs are $(C_1, C_2) = (\frac{3}{4}, \frac{1}{2})$.

Suppose that all type θ_2 students leave the system and all students of other types stay in the system and keep the same preferences as in the first round. This frees up 2 units at s_2 . Under RLDA, the second-round cutoffs are $(\hat{C}_1^R, \hat{C}_2^R) = (1, \frac{3}{4})$. In this case, the order condition holds, and FLDA and RLDA are type equivalent. It is simple to verify that both FLDA and RLDA assign measure $\hat{\mu}(s)$ of students of type $(\theta_1, \theta_{1,2}, \theta_{2,1})$ to school s , where

$$\hat{\mu}^F = \hat{\mu}^R = (\hat{\mu}(s_1), \hat{\mu}(s_2)) = ((1, 1, 0), (0, 2, 3)).$$

Suppose that all type θ_1 students leave the system and all students of other types stay in the system and keep the same preferences as in the first round. This frees up 1 unit at s_1 . Under RLDA, no new students are assigned to s_2 , and the previously bottom-ranked (but now top-ranked) measure 1 of students who find s_1 acceptable are assigned to s_1 . Hence, the second-round cutoffs are $(\hat{C}_1^R, \hat{C}_2^R) = (\frac{7}{8}, 1)$. In this case, the order condition does not hold. Type equivalence also does not hold because the FLDA and RLDA assignments are

$$\begin{aligned} \hat{\mu}^F &= ((2, 0, 0), (1/3, 7/3, 7/3)), \\ \hat{\mu}^R &= ((1.5, 0.5, 0), (1, 2, 2)). \end{aligned}$$


3.1. Axiomatic Justification of PLDA Mechanisms

We have shown that PLDA mechanisms satisfy a number of desirable properties. Namely, PLDA mechanisms respect guarantees and priorities and are two-round strategy-proof (in a strong sense), non-wasteful, constrained Pareto efficient, and anonymous. In this section, we show that in a setting with a single priority class, the PLDA mechanisms are the *only* mechanisms that satisfy all these properties as well as two mild technical conditions on the symmetry of the mechanism, even when we require only the weaker version of two-round strategy-proofness.

Definition 14. A two-round mechanism satisfies the *averaging* axiom if for every type θ and pair of schools (s, s') , the randomization of the mechanism does not affect the measure of students with type θ assigned to (s, s') in the first and second rounds, respectively. That is, for all θ, s, s' , there exists a constant $c_{\theta, s, s'}$ such that $\eta(\{\lambda \in \Lambda : \theta^\lambda = \theta, \mu(\lambda) = s, \hat{\mu}(\lambda) = s'\}) = c_{\theta, s, s'}$ w.p. 1.

In other words, the averaging axiom requires that despite potential randomness in how the mechanism produces both first- and second-round assignments

Figure 3. (Color online) In Example 1, FLDA and RLDA Are Type Equivalent When the Order Condition Holds and Give Different Assignments to Students of Every Type When the Order Condition Does Not Hold



Notes. The initial economy and first-round assignment are depicted on the top left. On the right, we show the second-round assignments under FLDA and RLDA when type θ_2 students (who want only s_2) drop out and when type θ_1 students (who want only s_1) drop out. Students toward the left have larger first-round lottery numbers. The patterned boxes above each column of students indicate the affordable sets for students in that column. When students who want only s_2 drop out, the order condition holds, and FLDA and RLDA are type equivalent. When students who want only s_1 drop out, s_2 becomes more overdemanded in RLDA, and FLDA and RLDA give different ex ante assignments to students of every remaining type.

for individual students, the distribution over (first-round assignment, second-round assignment) for students of a given type is deterministic. The averaging axiom serves the purpose of excluding mechanisms with aggregate randomness, which adds unnecessary complexity to the assignment process. For example, a mechanism that with probability 1/2 runs FLDA and with probability 1/2 runs RLDA satisfies all our other axioms besides averaging and violates averaging by adding unnecessary uncertainty. (Instead of relying on a coin-flip decision between FLDA and RLDA, the designer could choose the better option between the two based on their goals.) The averaging axiom precludes such mechanisms, and its particular technical form facilitates a sharp characterization of PLDA mechanisms.

Definition 15. A two-round mechanism is *nonatomic* if any single student changing his or her preferences

has no effect on the assignment probabilities of other students.

Our characterization result is the following.

Theorem 3. Suppose that student preferences are consistent and student types have full support (Assumptions 1 and 2). A nonatomic two-round assignment mechanism with first round DA-STB respects guarantees and is nonwasteful, (weakly) two-round strategy-proof, constrained Pareto efficient, anonymous, and averaging if and only if the second-round assignment is given by PLDA.

We remark that we require two-round strategy-proofness only for students whose true preference type is consistent. This is because preference inconsistencies across rounds can lead to conflicts between the desired first-round assignment with respect to first-round preferences and the desired first-round guarantee with respect to second-round preferences, making it unclear how to even define a best response.

Moreover, it is reasonable to assume that students who are sophisticated enough to strategize about misreporting in the first round in order to affect the guarantee structure in the second round will also know their second-round preferences over schools in S (i.e., everything except where they rank their outside option) at the beginning of the first round and hence will have consistent preferences.¹⁶ We remark also that the “only if” direction of this result is the only place where we require the full support assumption (Assumption 2).

The main focus of our result is the effect of cross-round constraints. By assumption, the first-round mechanism is DA-STB. It is relatively straightforward to deduce that the second-round mechanism also has to be DA-STB. Strategy-proofness in the second round, together with nonwastefulness, respecting priorities and guarantees, and anonymity, constrains the second round to be DA, with each student given a guarantee at the school to which he or she was assigned in the first round, and constrained Pareto efficiency forces the tie breaking to be in the same order at all schools. The cross-round constraints are more complicated but can be understood using *affordable sets*. A student’s affordable set is the set of schools that he or she can choose to attend—that is, the first-round affordable set is the set of schools for which he or she meets the first-round cutoff, and the affordable set is the set of schools for which he or she meets the first- or second-round cutoff. The set of possible affordable sets is uniquely determined by the order of cutoffs. By carefully using two-round strategy-proofness and anonymity, we show that a student’s preference type does not affect the joint distribution over his or her first-round affordable set and affordable set, and hence his or her second-round lottery is a permutation of his or her first-round lottery that does not depend on his or her preference type.

Our result mirrors similar large market cutoff characterizations for single-round mechanisms by Liu and Pycia (2016) and Ashlagi and Shi (2014), which show, in settings with single and multiple priority types, respectively, that a mechanism is non-atomic, strategy-proof, symmetric, and efficient (in each priority class) if and only if it can be implemented by lottery-plus-cutoff mechanisms, which provide random lottery numbers to each student and admit them to their favorite school for which they meet the admission cutoff. We obtain such a characterization in a two-round setting using the fact that the mechanism respects guarantees and introducing an *affordable set* argument to isolate the second round from the first. This simplification allows us to employ arguments similar to those used in Liu and Pycia 2016 and Ashlagi and Shi 2014 to show that the first- and

second-round mechanisms can be individually characterized using lottery-plus-cutoff mechanisms.

4. Intuition for Main Results

In this section, we provide some intuition for our main results. This section may be skipped at a first reading without loss of continuity.

A key insight is that we simplify analysis by shifting away from assignments, which depend on preferences, to considering the schools that a student can attend, which are independent of his or her preferences. Specifically, if we define the *affordable set* for each student as the set of schools for which he or she meets either the first- or second-round cutoffs, then each student is assigned to his or her favorite school in his or her affordable set at the end of the second round, and changing the student’s preferences does not change his or her affordable set in our continuum model. Moreover, affordable sets and preferences uniquely determine demand.

The main technical idea that we use in establishing our main results is that the order condition is equivalent to the following seemingly much more powerful “global” order condition.

Definition 16. We say that PLDA(P) satisfies the *local order condition* on a set of primitives (S, q, Λ, η) if, for every priority class π , the first- and second-round school cutoffs within that priority class are in the same order under PLDA(P). That is, for all $s_i, s_j \in S$,

$$C_{\pi,i} > C_{\pi,j} \Rightarrow \hat{C}_{\pi,i}^P \geq \hat{C}_{\pi,j}^P.$$

We say that the *global order condition* holds on a set of primitives (S, q, Λ, η) if

a. (Consistency across rounds) PLDA(P) satisfies the local order condition on $(S, q, \Lambda, \eta) \forall P$.

b. (Consistency across permutations) For every priority class π , for all pairs of permutations P, P' and schools $s_i, s_j \in S \cup \{s_{N+1}\}$, it holds that $\hat{C}_{\pi,i}^P > \hat{C}_{\pi,j}^P \Rightarrow \hat{C}_{\pi,i}^{P'} \geq \hat{C}_{\pi,j}^{P'}$.

In other words, the global order condition requires that all PLDA mechanisms result in the same order of school cutoffs in both rounds. Surprisingly, if the cutoffs are in the same order in both rounds under RLDA, then they are in the same order in both rounds under any PLDA.

Theorem 4. *The order condition (Definition 12) holds for a set of primitives (S, q, Λ, η) if and only if the global order condition holds for (S, q, Λ, η) .*

The affordable set framework can provide some intuition as to why Theorem 4 holds. Under the reverse permutation, the sets of schools that enter a student’s affordable set in the first and second rounds, respectively, are maximally misaligned. Hence, if as required by the local order condition the cutoff

order is consistent across both rounds under the reverse permutation, then the cutoff order should also be consistent across both rounds under any other permutation.

4.1. Global Order Condition Implies Type Equivalence

The affordable set framework can also be used to show that when the global order condition holds, all PLDA mechanisms produce type-equivalent assignments. Fix a mechanism and suppose that the first- and second-round cutoffs are in the same order. Then each student λ 's affordable set is of the form $X_i = \{s_i, s_{i+1}, \dots, s_N\}$ for some $i = i(\lambda)$, where schools are indexed in decreasing order of their cutoffs for the relevant priority group p^{θ^A} , and the probability that a student receives some affordable set is independent of his or her preferences. Moreover, because affordable sets are nested, $X_1 \supseteq X_2 \supseteq \dots \supseteq X_N$, and because the lottery order is independent of student types, the demand for schools is uniquely identified by the proportion of students whose affordable set contains s_i for each i . When the global order condition holds, this is true for every PLDA mechanism individually, which provides enough structure to induce type equivalence. We formalize this argument in the online appendix (Lemmas 1 and 2).

4.2. Uniform Dropouts

We now introduce a special case of our model. For this special case, we prove that the order condition holds. It follows that all PLDA mechanisms give type-equivalent assignments.

Definition 17 (Informal). A market satisfies *uniform dropouts* if there is exactly one priority group at each school, students leave the system independently with some fixed probability ρ , and the students who remain in the system retain their preferences.

Intuitively, in the uniform dropouts model, each student drops out of the system with probability ρ —for example, because he or she leaves the city after the first round for reasons that are independent of the school choice system. The second-round problem can thus be viewed as a rescaled version of the first-round problem. This suggests that schools should fill in the same order regardless of the choice of permutation because the measure of remaining students who were assigned to each school s_i in the first round is $(1 - \rho)q_i$, the measure of students of each type θ assigned to each school is scaled down by $1 - \rho$, the capacity of each school is still q_i , and the measure of students of each type θ who are still in the system is scaled down by $1 - \rho$.

Formally, let student types be defined by $\theta = (>^\theta, \hat{>}^\theta, \mathbf{1})$ (which we will write as $\theta = (>^\theta, \hat{>}^\theta)$ because there

are no priorities). We define *uniform dropouts with probability ρ* by

$$\begin{aligned} & \zeta\left(\left\{\theta = \left(>^\theta, \hat{>}^\theta\right) \in \Theta : >^\theta = >, \hat{>}^\theta = s_{N+1} > \dots\right\}\right) \\ &= \rho \zeta\left(\left\{\theta = \left(>^\theta, \hat{>}^\theta\right) \in \Theta : >^\theta = >\right\}\right), \\ & \zeta\left(\left\{\theta = \left(>^\theta, \hat{>}^\theta\right) \in \Theta : >^\theta = >, \hat{>}^\theta = >\right\}\right) \\ &= (1 - \rho) \zeta\left(\left\{\theta = \left(>^\theta, \hat{>}^\theta\right) \in \Theta : >^\theta = >\right\}\right). \end{aligned} \quad (3)$$

That is, all students with probability ρ find the outside option s_{N+1} the most attractive in the second round and otherwise retain the same preferences in the second round.¹⁷

Theorem 5. *In any market with uniform dropouts (Definition 17), the global order condition (Definition 16) holds.*

We prove Theorem 5 in the online appendix. The main steps of the proof mirror those used to prove Theorem 4, albeit in a simpler and more transparent setting, and can provide the interested reader with a taste of our more general proof techniques.

Remark. We can also use the uniform dropouts setting to explore what happens when students' assignments affect their preferences. We say that a market satisfies *uniform dropouts with inertia* if there is exactly one priority group at each school, students leave the system independently with some fixed probability ρ , students remain and wish to stay at their first round assignment with some fixed probability ρ' (have “inertia”), and students otherwise remain and retain their first-round preferences.¹⁸ It can be shown that in such a market, the global order condition always holds, and RLDA minimizes reassignment among all type-equivalent allocations. Moreover, if all students are assigned in the first round, it can also be shown that PLDA mechanisms produce type-equivalent allocations.

5. Empirical Analysis of PLDA Mechanisms

In this section, we use data from the NYC high school choice system to simulate and evaluate the performance of centralized PLDA mechanisms under different permutations P . The simulations indicate that although our theoretical results are for markets satisfying the order condition, they are real-world relevant. Different choices of P yield similar allocative efficiency: the number of students assigned to their k th choice for each rank k and the number of students remaining unassigned are similar for different permutations P . At the same time, the difference in the number of reassigned students is significant and is minimized under RLDA.

Motivated by current practice, we also simulate decentralized versions of FLDA and RLDA. In a version where students take time to vacate previously assigned

seats, reversing the lottery increases allocative efficiency during the early stages of reassignment and decreases the number of reassessments at every stage. However, in a version where students take time to decide on offers from the waitlist, the efficiency comparisons are reversed.¹⁹ In both versions, both FLDA and RLDA took tens of stages to converge. Our simulations suggest that decentralized waitlist mechanisms can achieve some of the efficiency gains of a centralized mechanism but incur significant congestion costs, and the effects of reversing the tie-breaking order before constructing waitlists will depend on the specific time and informational constraints of the market.

5.1. Data

We use data from the high school admissions process in NYC for the academic year 2004–2005 as follows.

5.1.1. First-Round Preferences. In our simulation, we take the first-round preferences \succ of every student to be the preferences they submitted in the main round of admissions. The algorithm used in practice is essentially strategy-proof (see Abdulkadiroglu et al. 2005a), justifying our assumption that reported preferences are true preferences.²⁰

5.1.2. Second-Round Preferences. In our simulation, students either drop out of the system entirely in the second round or maintain the same preferences. Students are considered to drop out if the data do not record them as attending any public high school in NYC the following year (this was the case for about 9% of the students each year).²¹

5.1.3. School Capacities and Priorities. Each school's capacity is set to the number of students assigned to it in the first-round assignment in the data. This is a lower bound on the true capacity but lets us compute the final assignment under PLDA with the true capacities because the occupancy of each school with vacant seats decreases across rounds in our setting. School priorities over students are obtained directly from the data. (We obtain similar results in simulations with no school priorities.)

5.2. Simulations

In a setting with a finite number of students, DA-STB uses an iterative process of student application and school tentative acceptance to assign students according to student preferences and school preference rankings after tie breaking, as described in Section 2.2. PLDA mechanisms are reassignment mechanisms that run DA-STB with modified school preferences \hat{p} in the second round: for each school s_i , students $\lambda \in \Lambda$ for whom $\mu(\lambda) = s_i$ are given additional priority n_i at school s_i to produce updated priorities \hat{p} , and ties

within the updated priority groups \hat{p} are broken according to the permuted lottery $P \circ L$ (in favor of the student with the larger permuted lottery number).

5.2.1. Centralized PLDA. We first consider the following family of centralized PLDA mechanisms, parameterized by a single parameter α that smoothly interpolates between RLDA and FLDA. Each student λ receives a uniform *independent and identically distributed* (i.i.d.) first-round lottery number $L(\lambda)$ (a normal variable with mean 0 and variance 1) that generates a uniformly random lottery order.²² The second-round “permuted lottery” of λ is given by $\alpha L(\lambda) + \tilde{L}(\lambda)$, where $\tilde{L}(\lambda)$ is a new i.i.d. normal variable with mean 0 and variance 1, and α is identical for all students. RLDA corresponds to $\alpha = -\infty$, and FLDA corresponds to $\alpha = \infty$. For a fixed real α , every realization of second-round scores corresponds to some permutation of first-round lottery numbers, with α roughly capturing the correlation of the second-round order with that of the first round. We quote averages across simulations.

5.2.2. Decentralized PLDA. In order to evaluate the performance of waitlist systems, we experiment with two different versions of managing the waitlists in a decentralized fashion. The starting point for both versions is the first-round assignment of students to schools. Each school maintains an ordered waitlist of students constructed as follows: any student who is not assigned to his or her most-preferred school is automatically included in the waitlist at every school he or she strictly prefers to his or her first-round assignment, and each school orders its waitlist in priority order, breaking ties using the chosen permutation of the first-round lottery numbers. We emphasize that students could appear on multiple waitlists and that individual schools do not know the preferences of the students, the priority of a student at another school, or the number of available seats at other schools.

Version 1. The algorithm proceeds iteratively in stages. The residual capacity of a school at the beginning of a stage is its original capacity minus the number of students in the system who are assigned to that school at that time. (Thus, if a school was filled to capacity in the first round, its residual capacity at the beginning of the first stage of the second round would simply be the number of its assigned students who dropped out of the system.) At the beginning of any stage ℓ , each school s_i makes an offer of admission to the top \tilde{q}_i^ℓ students on its waitlist (or to all students on its waitlist if there are fewer students), where \tilde{q}_i^ℓ is the residual capacity of school i at the beginning of stage ℓ . Note that these include offers to students who have dropped out of the system because the school is not aware of this

fact. A student who drops out of the system rejects all offers; students still in the system accept their most-preferred school among their first-round assignment and all offers received at this stage of the second round and reject every other offer. The schools update their waitlists and residual capacity, and the algorithm proceeds to the next stage.

There can be two types of “wasted” offers at each stage of version 1: a school may make an offer to a student who has already dropped out of the system because the school does not know that the student has dropped out (recall that when a student drops out, only his or her assigned school in the first round is notified), and multiple schools may make offers to a student at the same stage of the second round, and the “losing” schools come to know only at the end of the stage. Such wasted offers are avoided by version 2 of the algorithm described next.

Version 2. The algorithm proceeds iteratively in stages. The residual capacity of a school at the beginning of a stage is its original capacity minus the number of students in the system who are assigned to that school at that time. At each stage ℓ , the school-proposing DA algorithm (Definition 18) is run, with the school capacities being the residual capacities at that stage, student preferences restricted to schools that the student strictly prefers to his or her current assignment, and ties broken using the chosen permutation of first-round lottery numbers.

In contrast to version 1, here the underlying premise is that students who have dropped out or those who receive multiple offers *immediately* reject all but their most-preferred offer, thus allowing schools to make more offers within the stage than their residual

capacity permits if some offers are rejected. However, as in version 1, schools adjust their residual capacity only at the end of the stage. That is, a student who accepts an offer from a school at some stage ℓ in the second round does not notify his or her previously assigned school until the end of the stage (equivalently, that school is allowed to fill this student’s spot only in stage $\ell + 1$).

Version 1 of the decentralized PLDA mechanisms mirrors a decentralized process where students take time to make decisions. However, it does so in a rather naive fashion by assuming that students take the same unit of time (one stage) both to directly respond to offers (i.e., accept or reject an offer) and to indirectly respond (i.e., notify a school that they were previously assigned to that they have been assigned to a different school). Version 2 captures a decentralized process where direct responses to offers are instantaneous, so rejected offers may result in many offers in a single stage, whereas updates to schools that are indirectly affected happen only at the end of each stage. Accordingly, the efficiency outcomes at a given stage of version 2 dominate those of version 1 at the same stage because more information is communicated during each stage.

In practice, we expect that the dynamics of waitlist systems would lie somewhere on the spectrum between these two extreme versions of decentralized PLDA.


5.3. Results

The results of our centralized PLDA computational experiments based on 2004–2005 NYC high school admissions data appear in Table 1 and Figure 4. Figure 4 shows that the mean number of reassignments is minimized at $\alpha = -\infty$ (RLDA) and increases

Table 1. Centralized PLDA Simulation Results: 2004–2005 NYC High School Admissions

α	Reassignments N	Unassigned %	$k = 1$ %	$k \leq 2$ %	$k \leq 3$ %
Round 1 (no reassignment)	0	9.31	50.14	64.14	72.44
Round 2					
FLDA: ∞	7,797	5.89	55.41	69.85	78.03
8.00	7,606	5.90	55.40	69.85	78.02
6.00	7,512	5.90	55.40	69.85	78.03
4.00	7,325	5.89	55.38	69.84	78.02
2.00	6,863	5.89	55.33	69.81	78.02
0.00	5,220	5.87	54.96	69.65	77.97
-2.00	3,686	5.81	54.52	69.37	77.82
-4.00	3,480	5.79	54.47	69.33	77.78
-6.00	3,433	5.79	54.46	69.32	77.77
-8.00	3,416	5.79	54.45	69.31	77.77
RLDA: $-\infty$	3,391	5.79	54.45	69.30	77.75

Notes. We show the mean percentage of students remaining unassigned or getting at least their k th choice averaged across 100 realizations for each value of α . All percentages are out of the total number of students remaining in the second round. The data contained 81,884 students, 74,366 students remaining in the second round, and 652 schools. The percentage of students who dropped out was 9.18%. The variation in the number of reassignments across realizations was ~ 100 students.

Figure 4. Number of Reassigned Students vs. α


Note. The number of reassigned students under the extreme values of α —namely, $\alpha = \infty$ (FLDA) and $\alpha = -\infty$ (RLDA)—are shown by dotted lines.

with α , which is consistent with our theoretical result in Theorem 2. The mean number of reassessments is as large as 7,800 under FLDA compared with just 3,400 under RLDA.

Allocative efficiency appears not to vary much across values of α : the number of students receiving at least their k th choice for each $1 \leq k \leq 12$ and the number of unassigned students vary by less than 1% of the total number of students. There is a slight trade-off between allocative efficiency resulting from reassignment and allocative efficiency from assigning previously unassigned students, with the percentage of unassigned students and percentage of students obtaining their top choice both decreasing in α by about 0.1% and 1% of students, respectively.²³ We further find that for most students, the likelihoods of getting one of their top k choices under FLDA and under RLDA are very close to each other. (For instance, for 87% of students, these likelihoods differ by less than 3% for all k .) This is consistent with what we would expect based on our theoretical finding of type equivalence (Theorem 1) of the final assignment under different PLDA mechanisms.

The results of our decentralized PLDA computational experiments appear in Table 2. When implementing PLDAs in a decentralized fashion, our measures of congestion can be more nuanced. We let a *reassignment* be a movement of a student from a school in S to a different school in S , possibly during an interim stage of the second round, and let a *temporary reassignment* be a movement of a student from a school in $S \cup \{s_{N+1}\}$ to a different school in S that is not his or her final assignment. We will also be interested in the number of stages it takes to clear the market.

In the first version of decentralized PLDAs, FLDA reassesses more students than RLDA but far outperforms RLDA in terms of minimizing congestion and maximizing efficiency. FLDA takes, on average, 17 stages to converge, whereas RLDA requires 33 stages. FLDA performs 780 temporary transfers, whereas RLDA performs 2,420, creating much more unnecessary congestion. FLDA takes two and five stages to achieve 50% and 90%, respectively, of the total increase in the number of students assigned to their top school, whereas RLDA takes three and nine stages, respectively. FLDA also dominates RLDA in terms of the number of students assigned to one of their top k choices in the first ℓ stages, for all k and all ℓ , and the percentage of unassigned students in the first ℓ stages for almost all small ℓ .²⁴

In the second version of decentralized PLDAs, FLDA still reassesses more students and now achieves less allocative efficiency than RLDA during the initial stages of reassignment. RLDA has fewer unassigned students by stage ℓ than FLDA for all ℓ . RLDA also dominates FLDA in terms of the number of students assigned to one of their top k choices in the first two stages and achieves most of its allocative efficiency by the second stage, improving the allocative efficiency by fewer than 100 students from that point onward. In the limit, FLDA is still slightly more efficient than RLDA, so for large ℓ , FLDA achieves higher allocative welfare than RLDA after ℓ stages. However, FLDA also requires more stages to converge, taking, on average, 12 stages compared with 9 stages for RLDA.

Our empirical findings have mixed implications for implementing decentralized waitlists. Our clearest finding is the benefit of centralization in reducing

Table 2. Decentralized PLDA Simulation Results: 2004–2005 NYC High School Admissions

α	Reassignments Total (temporary)	Unassigned %	$k = 1$ %	$k \leq 2$ %	$k \leq 3$ %
Round 1 (no reassessments)	0	9.31	50.14	64.17	72.45
Round 2 FLDA, version 1					
Stage 1	3,461 (447)	7.89	52.68	66.62	74.47
Stage 2	2,126 (206)	7.04	53.93	68.03	76.14
Stage 3	1,258 (80)	6.55	54.60	68.83	76.96
Stage 4	727 (30)	6.27	54.97	69.28	77.42
Stage 5	425 (11)	6.11	55.18	69.53	77.68
Total (stage ≈ 17)	8,590 (780)	5.87	55.46	69.87	78.05
Round 2 RLDA, version 1					
Stage 1	1,004 (835)	7.85	51.38	65.70	74.09
Stage 2	1,077 (577)	7.18	52.24	66.72	75.15
Stage 3	838 (369)	6.78	52.82	67.39	75.83
Stage 4	640 (234)	6.52	53.23	67.86	76.30
Stage 9	180 (24)	5.97	54.22	69.02	77.45
Total (stage ≈ 33)	5,818(2419)	5.79	54.51	69.37	77.80
Round 2 FLDA, version 2					
Stage 1	4,139 (457)	7.62	53.21	67.14	75.21
Stage 2	2,333 (166)	6.69	54.50	68.66	76.75
Stage 3	1137 (42)	6.24	55.06	69.35	77.48
Stage 4	511 (9)	6.03	55.30	69.65	77.80
Total (stage ≈ 12)	8,503 (677)	5.89	55.47	69.87	78.04
Round 2 RLDA, version 2					
Stage 1	2,863 (199)	6.15	54.14	68.85	77.24
Stage 2	489 (17)	5.88	54.38	69.16	77.58
Stage 3	165 (2)	5.82	54.46	69.26	77.69
Stage 4	63 (0)	5.79	54.49	69.30	77.73
Total (stage ≈ 9)	3,624 (220)	5.79	54.51	69.33	77.76

Notes. We show the mean number of reassessments (number of movements of a student from a school in S to a different school in S) and the mean number of temporary reassessments (number of movements of a student from a school in $S \cup \{s_{N+1}\}$ to a different school in S that is not their final assignment) in parentheses. We also show mean percentage of students remaining unassigned or getting at least their k th choice. All figures are averaged across 100 realizations for each value of α , and all percentages are out of the total number of students remaining in the second round. The data contained 81,884 students, 74,366 students remaining in the second round, and 652 schools.

congestion. In most school districts, students are given up to a week to make decisions. If students take this long both to reject undesirable offers and to vacate previously assigned seats, our simulations on NYC data suggest that in the best case the market could take at least 4 months to clear. Even if students make quick decisions, if it takes them a week to vacate their previously assigned seats, our simulations suggest that the market would take at least 2 months to clear. In both cases, the congestion costs are prohibitive. If, despite these congestion costs, a school district wishes to implement decentralized waitlists, our results suggest that the optimal permutation for the second-round lottery for constructing waitlists will depend on the informational constraints in the market.

5.4. Strategy-Proofness of PLDA

One of the aspects of the DA mechanism that makes it successful in school choice in practice is that it is strategy-proof. While we have shown that PLDA mechanisms are two-round strategy-proof in a continuum setting, it is natural to ask to what extent

PLDA mechanisms are two-round strategy-proof in practice. We provide a numerical upper bound on the incentives to deviate from truthful reporting using computational experiments based on 2004–2005 NYC high school data and find that, on average, a negligible proportion of students ($< 0.01\%$) could benefit from misreporting within their consideration set of programs. Specifically, 0.8% of sampled students could misreport in a potentially beneficial manner in at least 1 of 100 sampled lotteries, and no students could benefit in more than 3 of 100 sampled lotteries from misreporting. Moreover, for 99.8% of lotteries, the proportion of students who could successfully manipulate their report is at most 1%.²⁵

6. Proposals and Discussion

6.1. Summary of Findings

We have proposed the PLDA mechanisms as a class of reassignment mechanisms with desirable incentive and efficiency properties. These mechanisms can be implemented with a centralized second round at the start of the school year or with a decentralized second

round via waitlists, and a suitable implementation can be chosen depending on the timing of information arrival and subsequent congestion in the market. Moreover, the key defining characteristic of the mechanisms in this class, the permutation used to correlate the tie-breaking lotteries between rounds, can be used to optimize various objectives. We propose implementing centralized RLDA at the start of the school year because both in our theory and in simulations on data this allows us to maintain efficiency while eliminating the congestion caused by sequentially reassigning students and minimizes the number of reassessments required to reach an efficient assignment.

6.2. RLDA Is Practical

Reversing the lottery between rounds is simple to understand and implement. It also has the nice property of being equitable in an intuitive manner because students who receive a poor draw of the lottery in the first round are prioritized in the second round. This may make RLDA more palatable to students than other PLDA mechanisms. Indeed, Random Hall, a Massachusetts Institute of Technology (MIT) dorm, uses a mechanism for assigning rooms that resembles the reverse-lottery mechanism we have proposed. Freshmen rooms are assigned using serial dictatorship. At the end of the year (after seniors leave), students can claim the rooms vacated by the seniors using serial dictatorship, where the initial lottery numbers (from their first match) are reversed.²⁶

6.3. Optimizing Other Objectives

Our results suggest that PLDA mechanisms are an attractive class of mechanisms in more general settings, and the choice of PLDA mechanism will vary with the policy goal. If, for instance, it were viewed as more equitable to allow more students to receive (possibly small) improvements to their first-round assignment, implementing FLDA optimizes this. Our type-equivalence result (Theorem 1) shows that when students have consistent preferences and the order condition holds, this choice can be made without sacrificing allocative efficiency.

6.4. Discussion of Axiomatic Characterization

Our characterization for PLDA mechanisms (Theorem 3) does not incorporate priorities. In a model with priorities, we find that natural extensions of our axioms continue to describe PLDA mechanisms but also include undesirable generalizations of PLDA mechanisms. Specifically, suppose that we add an axiom requiring that for each school s , the probability that a student who reports a top choice of s then receives it in the first or second round is independent of that student's priority at other schools. This new set of axioms

describes a class of mechanisms that strictly includes the PLDA mechanisms, as well as mechanisms where the permutation of the lottery can depend on students' priorities. Characterizing the class of mechanisms satisfying these axioms in the richer setting with school priorities remains an open question. It may also be possible to characterize PLDA mechanisms in a setting with priorities using a different set of axioms. We leave both questions for future research.

6.5. Finite Markets

It is natural to ask what implications our results have for finite markets. Azevedo and Leshno (2016) have shown that if a sequence of (large) discrete economies converges to some limiting continuum economy with a unique stable matching (defined via cutoffs), then the stable matchings of the discrete economies converge to the stable matching of the continuum. This suggests that our theoretical results should approximately hold for large, discrete economies. As an example, we provide a heuristic argument for why PLDA mechanisms satisfy the strategy-proofness in the large condition defined by Azevedo and Budish (2013). By definition, PLDA mechanisms satisfy the efficiency and anonymity requirements in finite markets as well. In the second round, it is clearly a dominant strategy to be truthful, and intuitively, for a student to benefit from a first-round manipulation, his or her report should affect the second-round cutoffs in a manner that gives him or her a second-round assignment that he or she would not have received otherwise. If the market is large enough, the cutoffs will converge to their limiting values, and the probability that the student could benefit from such a manipulation would be negligible. (Indeed, in simulations on NYC high school data, we find that the average proportion of students who can successfully manipulate their report is $< 0.01\%$; see Section 5.4.) A similar argument suggests that an approximate version of our characterization result (Theorem 3) should hold for finite markets with no priorities. Our type-equivalence result (Theorem 1) and result showing that RLDA minimizes transfers (Theorem 2) should also be approximately valid in the large market limit.²⁷

6.6. Inconsistent Preferences

Another natural question is how to deal with inconsistent student preferences. Narita (2016) observed that in the current reapplication process in the NYC public school system, about 5% of students reapplied with second-round preferences that were inconsistent with their first-round reported preferences.²⁸ PLDAs retain most of their desirable properties even when students report inconsistent preferences in the second round. We believe that some of our insights on

optimizing allocative efficiency and reassignment remain valid if a small fraction of students has an idiosyncratic change in preferences or if a small number of new students enter in the second round. We also believe that our qualitative insights that reversing the lottery reduces reassignment and that the choice of PLDA has smaller effects on allocative efficiency in settings where aggregate demand in the two rounds are well aligned and will continue to hold quite generally. However, new effects may emerge if students have arbitrarily different preferences in the two rounds. In such settings, strategy-proofness is no longer well defined. It can also be shown that the order condition is no longer sufficient to guarantee type equivalence and optimality of RLDA, and the relative efficiency of the PLDA mechanisms will depend on the details of school supply and student demand. Our simulations further suggest that in such settings there may be a more significant trade-off between decreasing the number of unassigned and reassigned students and increasing the efficiency of the final assignment. Mechanisms such as RLDA that prioritize students with poor first-round assignments are likely to perform better on the former, whereas mechanisms such as FLDA that minimize the constraints imposed by first-round guarantees will likely perform better on the latter. We leave a theoretical study of this trade-off for future research.

6.7. More Than Two Rounds

Finally, what insights do our results provide for when assignment is done in three or more rounds? For instance, one could consider mechanisms under which the lottery is reversed (or permuted) after a certain number of rounds and thereafter remains fixed. At what stage should the lottery be reversed? Clearly, there are many other mechanisms that are reasonable for this problem, and we leave a more comprehensive study of this question for future work.

Acknowledgments

The authors are grateful to Atila Abdulkadiroglu for introducing them to this problem, and Yusuke Narita and seminar participants at Caltech, Columbia, Cornell, Duke, Facebook Research, Georgia Tech, Google Research, Johns Hopkins, Hong Kong University of Science and Technology, Microsoft Research, Massachusetts Institute of Technology, Northwestern, National University of Singapore, Princeton, Purdue, Stanford, and Yale for their helpful comments. They would like to thank research assistants Paul Laliberte, Alexej Proskynitopoulos, Pavithra Srinath, Avani Tanna, and Cyndia Yu for their help with data analysis and simulations.

Endnotes

¹ In the 2004–2005 school year, 9.22% of a total of 81,884 students dropped out of the public school system after the first round. Numbers for 2005–2006 and 2006–2007 are similar.

² A decentralized version of FLDA is used in most cities and in NYC kindergarten admissions.

³ Capacity constraints are binding in most schools. Most states impose maximum class sizes and fund schools based on enrollment after the first 2–3 weeks of classes, which incentivizes schools to enroll as many students as permissible.

⁴ We describe the decentralized reassignment processes currently used in NYC kindergarten; Boston; Washington, DC; Denver; Seattle; New Orleans; and Chicago. A similar process was also used in NYC high school admissions until a few years ago, when the system abandoned reassignments entirely, anecdotally because of the excessive logistical difficulties created by market congestion.

⁵ Students who have accepted an offer off the waitlist of one school are allowed to accept offers off the waitlists of other schools. Because registration for one school automatically cancels the student's previous registrations, this would automatically release the seat the student accepted from the first school to other students on the waitlist.

⁶ See, for example, Abdulkadiroglu et al. (2005a, b) for an overview of the redesigns in New York City (2003) and Boston (2005), respectively. These were followed by New Orleans (2012), Denver (2012), and Washington, DC (2013)—among others. See Abdulkadiroglu et al. (2017) for welfare analysis of the changes in NYC.

⁷ Our continuum model can be viewed as a two-round version of the model introduced by Azevedo and Leshno (2016). Continuum models have been used in a number of papers on school choice; see Agarwal and Somaini (2018), Ashlagi and Shi (2014), and Azevedo and Leshno (2016). Intuitively, one can think of the continuum model as a reasonable approximation of the discrete model in Online Appendix B when the number of students is large, although establishing a formal relationship between the discrete and continuum models is beyond the scope of our paper.

⁸ Consistency is required to meaningfully define strategy-proofness in our two-round setting because we require truthful reporting in the first round to be optimal for both the student's first- and second-round assignments.

⁹ This can be justified via an axiomatization of the kind obtained by Al-Najjar (2004).

¹⁰ This ensures that students will report their full first-round preferences in the first round instead of truncating in the first round based on their beliefs about their second-round preferences.

¹¹ Several alternative definitions of reassigned students—such as counting students who are initially unassigned and end up at a school in S and/or counting initially assigned students who end up unassigned—could also be considered. We note that our results continue to hold for all these alternative definitions.

¹² Here we make the restriction that the second-round assignment depends on the first-round reports only indirectly, through the first-round assignment μ . We believe that this is a reasonable restriction, given that the second round occurs a significant period of time after the first round, and the mechanism should come across as fair to the students.

¹³ If the demand function is continuously differentiable in the cutoffs, the assignment is unique. For an arbitrary demand function, the resulting assignment is unique for all but a measure zero set of capacity vectors.

¹⁴ When schools have strict preferences, an assignment respects priorities if and only if it is stable, and it is well known that in two-sided matching markets with strict preferences, there exist preference structures for which every stable assignment can be Pareto improved (Erdil and Ergin 2008).

¹⁵ We are not suggesting that the mechanism should involve checking the order condition and then using centralized RLDA only if this condition is satisfied (based on the guarantee in Theorems 1 and 2). However, one could check whether the order condition holds on historical data and accordingly decide whether to use the centralized RLDA mechanism or not.

¹⁶ One obvious objection is that students may also obtain extra utility from staying at a school between rounds, or equivalently, they may have a disutility for moving, creating inconsistent preferences where the school they are assigned to in the first round becomes preferred to previously more desirable schools. We remark that Theorem 3 extends to the case of students whose preferences incorporate additional utility if they stay put, provided that the utility is the same at every school for a given student or satisfies a similar noncrossing property.

¹⁷ We remark that there is a well-known technical measurability issue with regard to a continuum of random variables and that this issue can be handled; see, for example, Al-Najjar (2004).

¹⁸ This market is slightly beyond the scope of our general model because the type of the student now also has to encode second-round preferences that depend on the first-round assignment—namely, whether they have inertia.

¹⁹ This is due to a phenomenon that occurs when the second round is decentralized (not captured by our theoretical model), where under the reverse lottery the students with the worst lottery in the first round increase the waiting time for other students in the second round by increasing the waiting time for other students in the second round by considering multiple offers off the waitlist that they eventually decline.

²⁰ The algorithm is not completely strategy-proof because students may rank no more than 12 schools. However, only a very small percentage of students rank 12 schools. Another issue is that there is some empirical evidence that students do not report their true preferences even in school choice systems with strategy-proof mechanisms; see, for example, Hassidim et al. (2015) and Narita (2016).

²¹ For a minority of the students (9.2%–10.45%), attendance in the following year could not be determined by our data, and hence we assume that they drop out randomly at a rate equal to the dropout rate for the rest of the students (9.2%).

²² School preferences are then generated by considering students in the lexicographic ordering first in terms of priority and then by lottery number. We may equivalently renormalize the set of realized lottery numbers to lie in the interval $[0, 1]$ before computing scores.

²³ Intuitively, prioritizing students with lower lotteries has the desirable effect of decreasing the number of unassigned students because it prioritizes more students who were not assigned in the first round. However, it also decreases allocative efficiency by artificially increasing the constraints imposed by first-round guarantees.

²⁴ The mechanism driving the inefficiency of decentralized RLDA is that all offers in the first few stages of round 2 are made to a small number of students: those with the worst first-round lottery numbers. This inefficiency should be mitigated when students respond quickly to most of their offers, as we see in the second version of decentralized PLDA.

²⁵ These upper bounds were computed as follows: approximately 2,700 students were sampled, and RLDA was run for each of these students using 100 different sampled lotteries. For a given student, let S be the set of schools that were a part of the student's first-round preferences in the data. We allowed the student to unilaterally misreport in the first round, reporting at most one school from S in the first round instead of his or her true preferences. We then counted the number of such students who by doing so could either (a) change their first-round assignment (for the worse) but second-round assignment for the better or (b) create a rejection cycle. This provides a provable upper bound on the number of students who can benefit from misreporting (and possibly reordering) a subset of S in the first round. We omit the formal details in the interest of space.

²⁶ The MIT Random Hall matching is more complicated because sophomores and juniors can also claim the vacated rooms, but the lottery only gets reversed at the end of freshman year. Afterward, if a

sophomore switches rooms, his or her priority drops to the last place of the queue.

²⁷ Specifically, consider a sequence of markets of increasing size. If the global order condition holds in the continuum limit, this should lead to approximate type-equivalence under all PLDAs and to RLDA approximately minimizing transfers among PLDAs in the finite markets as market size grows. Moreover, if the order condition holds, then in large, finite economies and for every permutation P , the set of students who violate a local order condition on PLDA(P) will be small relative to the size of the market.

²⁸ About 7% of students reapplied, and about 70% of these reapplicants reported second-round preferences that were inconsistent with their first-round reported preferences.

References

Abdulkadiroglu A, Sönmez T (1999) House allocation with existing tenants. *J. Econom. Theory* 88(2):233–260.

Abdulkadiroglu A, Sönmez T (2003) School choice: A mechanism design approach. *Amer. Econom. Rev.* 93(3):729–747.

Abdulkadiroglu A, Sönmez T (2011) *Matching Markets: Theory and Practice* (Cambridge University Press, Cambridge, UK).

Abdulkadiroglu A, Agarwal N, Pathak PA (2017) The welfare effects of coordinated assignment: Evidence from the New York City high school match. *Amer. Econom. Rev.* 107(12):3635–3689.

Abdulkadiroglu A, Pathak PA, Roth AE (2005a) The New York City high school match. *Amer. Econom. Rev.* 95(2):364–367.

Abdulkadiroglu A, Pathak PA, Roth AE (2009) Strategy-proofness vs. efficiency in matching with indifferences: Redesigning the NYC high school match. *Amer. Econom. Rev.* 99(5):1954–1978.

Abdulkadiroglu A, Pathak PA, Roth AE, Sönmez T (2005b) The Boston public school match. *Amer. Econom. Rev.* 95(2):368–371.

Agarwal N, Somaini P (2018) Demand analysis using strategic reports: An application to a school choice mechanism. *Econometrica* 86(2):391–444.

Al-Najjar NI (2004) Aggregation and the law of large numbers in large economies. *Games Econom. Behav.* 47(1):1–35.

Anderson R, Ashlagi I, Gamarnik D, Kanoria Y (2017) Efficient dynamic barter exchange. *Oper. Res.* 65(6):1446–1459.

Anderson R, Ashlagi I, Gamarnik D, Rees M, Roth AE, Sönmez T, Ünver MU (2015) Kidney exchange and the alliance for paired donation: Operations research changes the way kidneys are transplanted. *Interfaces* 45(1):26–42.

Arnold N (2015) Short lists in centralized clearinghouses. *Proc. 16th ACM Conf. Econom. Comput.* (Association for Computing Machinery, New York), 751.

Arnold N, Shi P (2017) How (not) to allocate affordable housing. *Proc. 2017 ACM Conf. Econom. Comput.* (Association for Computing Machinery, New York), 679.

Ashlagi I, Nikzad A (2016) What matters in school choice tie-breaking? How competition guides design. *Proc. 17th ACM Conf. Econom. Comput.* (Association for Computing Machinery, New York).

Ashlagi I, Shi P (2014) Improving community cohesion in school choice via correlated-lottery implementation. *Oper. Res.* 62(6): 1247–1264.

Ashlagi I, Shi P (2015) Optimal allocation without money: An engineering approach. *Management Sci.* 62(4):1078–1097.

Ashlagi I, Nikzad A, Romm A (2015) Assigning more students to their top choices: A tiebreaking rule comparison. *Proc. 16th ACM Conf. Econom. Comput.* (Association for Computing Machinery, New York), 755–756.

Ashlagi I, Burq M, Jaillet P, Manshadi V (2019) On matching and thickness in heterogeneous dynamic markets. *Oper. Res.* 67(4): 927–949.

Azevedo EM, Budish EB (2013) Strategy-proofness in the large. Working paper, University of Chicago, Chicago.

Azevedo EM, Leshno JD (2016) A supply and demand framework for two-sided matching markets. *J. Political Econom.* 124(5): 1235–1268.

Balinski M, Sönmez T (1999) A tale of two mechanisms: Student placement. *J. Econom. Theory* 84(1):73–94.

Combe J, Tercieux O, Terrier C (2016) The design of teacher assignment: Theory and evidence. *ACM Conf. Econom. Comput.* (Association for Computing Machinery, New York).

Compte O, Jehiel P (2008) Voluntary participation and reassignment in two-sided matching. Working paper, Paris School of Economics, Paris.

Dickerson JP, Sandholm T (2015) Futurematch: Combining human value judgments and machine learning to match in dynamic environments. *Proc. 29th AAAI Conf. Artificial Intelligence* (Association for the Advancement of Artificial Intelligence, Palo Alto, CA), 622–628.

Dur U (2012) Dynamic school choice problem. Working paper, North Carolina State University, Raleigh.

Erdil A, Ergin H (2008) What's the matter with tie-breaking? Improving efficiency in school choice. *Amer. Econom. Rev.* 98(3):669–689.

Gale D, Shapley LS (1962) College admissions and the stability of marriage. *Amer. Math. Monthly* 69(1):9–15.

Hassidim A, Marciano-Romm D, Romm A, Shorrer RI (2015) “Strategic” behavior in a strategy-proof environment. Working paper, Bar-Ilan University, Ramat Gan, Israel.

Kadam SV, Kotowski MH (2014) Multi-period matching. Working paper, Harvard University, Cambridge, MA.

Kaplan EH (1987) Analyzing tenant assignment policies. *Management Sci.* 33(3):395–408.

Lee H-S, Srinivasan MM (1989) Control policies for the MX/G/1 queueing system. *Management Sci.* 35(6):708–721.

Liu Q, Pycia M (2016) Ordinal efficiency, fairness, and incentives in large markets. Working paper, Columbia University, New York.

Narita Y (2016) Match or mismatch: Learning and inertia in school choice. Working paper, Yale University, New Haven, CT.

Pathak PA (2011) The mechanism design approach to student assignment. *Annual Rev. Econom.* 3(1):513–536.

Pereyra JS (2013) A dynamic school choice model. *Games Econom. Behav.* 80:100–114.

Roth AE, Sönmez T, Ünver MU (2004) Kidney exchange. *Quart. J. Econom.* 119(2):457–488.

Su X, Zenios SA (2006) Recipient choice can address the efficiency-equity trade-off in kidney transplantation: A mechanism design model. *Management Sci.* 52(11):1647–1660.

Talluri KT, Van Ryzin GJ (2006) *The Theory and Practice of Revenue Management* (Springer Science and Business Media, New York).

Whitt W (2002) *Stochastic-Process Limits: An Introduction to Stochastic-Process Limits and Their Application to Queues* (Springer Science and Business Media, New York).

Zenios SA (1999) Modeling the transplant waiting list: A queueing model with reneging. *Queueing Systems* 31(3–4):239–251.