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

Itai Feigenbaum* Yash Kanoria† Irene Lo‡ Jay Sethuraman§

July 17, 2018

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 *reversing* the lottery order between rounds minimizes reassignment among all PLDA mechanisms. Empirical investigations based on data from NYC high school admissions support our theoretical findings.

Keywords: dynamic matching, matching markets, school choice, deferred acceptance, tie-breaking, cancellations, reassignments.

1 Introduction

In public school systems throughout the United States, students submit preferences over the schools for which they are eligible for admission. As this occurs fairly early in the school year, students typically do not know their options outside of 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 NYC public high school

*Lehman College and the Graduate Center, CUNY, Bronx, NY; itai.feigenbaum@lehman.cuny.edu

†Decision, Risk and Operations Division, Columbia Business School, New York, NY; ykanoria@columbia.edu

‡IEOR Department, Columbia University, New York, NY; iyl2104@columbia.edu

§IEOR Department, Columbia University, New York, NY; jay@ieor.columbia.edu

system, over 80,000 students are assigned to 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 reassigns students to vacated seats.

During the past fifteen 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 of 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 which can be reassigned. In the first round, schools have weak priorities over students, and students submit strict ordinal

[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.

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 mechanism (DA) where ties in school preferences are broken via a single lottery ordering across all schools. Afterwards, 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. As 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 she prefers to 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 assignment, 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

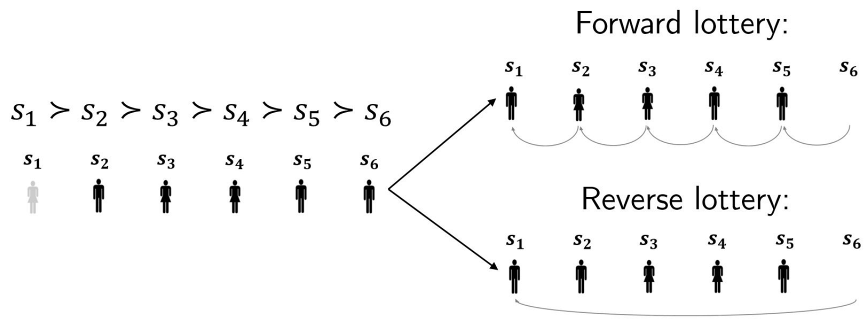

Figure 1: Running DA with a reversed lottery eliminates the cascade of reassignments.

There are 6 students with identical preferences over schools, and 6 schools each with a single priority group. All students prefer schools in the order $s_1 \succ s_2 \succ \cdots \succ s_6$. 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 reassigns each student to the school one better on her preference list, whereas reversing the tie-breaking lottery reassigns only the student initially assigned to s_6.

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). Suprisingly, however, our theoretical result holds in 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 New York City 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 or FLDA), and rerunning DA using an independent random lottery. We find 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 reassignments and RLDA results in about 3,400 reassignments out of a total of about 74,000 students who remained in the public school system, i.e. fewer than half the number of reassignments under FLDA. The gains become even more marked if we compare with current practice: RLDA results in fewer than 40% of the 8,600 reassignments 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.

Organization of the paper. 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.

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

1.1 Current Practice

Schools systems in cities across the US 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 Deferred Acceptance algorithm with single-tiebreaking. Students must register in their assigned school by April or early May.

In March and April students are also admitted to private and charter school 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 making "over offers" in the first round and accepting more students than they have seats for (Szuflita, b). However, such oversubscription of students is usually conservative (Szuflita, b), due to 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, as 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

[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 New York 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 due to the excessive logistical difficulties created by market congestion.

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 "huge slow round robin" (Szuflita, a) of reassignment 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 generalize 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 rouns 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, result 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 e.g. Pathak (2011) and Abdulkadiroglu and Sönmez (2011) for recent surveys. Many of these works provide axiomatic justifications for two canonical mechanisms, Deferred Acceptance Gale and Shapley (1962) and Top Trading Cycles (TTC), and their variants, in terms of their desirable properties. We provide a similar framework for the reassignment problem by

[5]Students who have accepted an offer off the waitlist of one school are allowed to accept offers off the waitlists of other schools. Since 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 e.g. Abdulkadiroglu et al. (2005a) and Abdulkadiroglu et al. (2005b) 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. (2015) for welfare analysis of the changes in NYC.

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 to 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., 2017). Reassignments due to cancellations also arise in online assignment settings such as kidney transplantation (e.g., see Zenios, 1999; Su and Zenios, 2006) and public housing allocation (e.g., see 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, and so optimal assignment policies account for typical cancellation and arrival statistics and optimize for agents arriving in the future (e.g., see 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 (e.g., see 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, and 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), Kadam and Kotowski (2014) and Pereyra (2013), 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 are 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 queueing 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 queueing system (e.g. FIFO, LIFO, SRPT 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 queueing (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 $S = \{s_1, \ldots, s_N\}$. 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 $q_1, \ldots, q_N \in \mathbb{R}_+$, 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 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 her preferences and school priorities. The student's first- and second-round preferences, respectively $\succ^\theta$ 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, e.g., because a new private school offer comes in. The student's *priority class* p^θ encodes 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, students ineligible for school s_i have priority $p_i = -1$, and that $p_i \in \{-1, 0, 1, \ldots, n_i - 1\}$. 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.

[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 (2014), 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 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.

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, i.e.: (1) 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 (2) the relative ranking of schools is unchanged across rounds,* $\forall s_i, s_j \in S$ *if* $s_i \hat{\succ} s_{N+1}$ *and* $s_i \hat{\succ} s_j$ *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.*

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

Consistency is required to meaningfully define strategy-proofness in our two-round setting, as we require truthful reporting in the first round to be optimal for both the student's first- and second-round assignments. We use the full support assumption only to characterize our proposed mechanisms (Theorem 3) and do not need it for our positive results (Theorems 1 and 2). We also assume that the first-round lottery numbers are drawn independently and uniformly from $[0, 1]$ and do not depend on preferences, i.e. $\forall \theta \in \Theta$ and $0 \leq a \leq b \leq 1$, $\eta(\{\theta\} \times (a, b)) = (b - a)\zeta(\theta)$.[8]

An *assignment* $\mu : \Lambda \to S$ specifies the school that each student is assigned to. 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*, i.e., $\eta(\mu(s_i)) \leq q_i$ 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.

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.

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

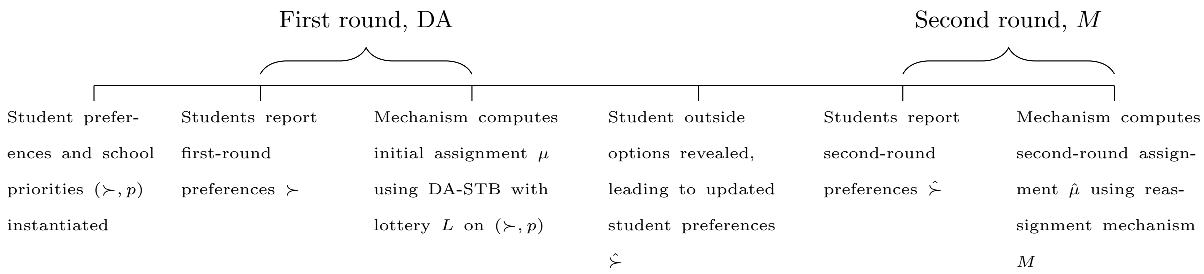

Figure 2: Timeline of the two-round mechanism design problem

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 her first-round preferences, and that her second-round preferences will be obtained from these preferences via truncation. Each student has imperfect information regarding her own second-round preferences (i.e., the point of truncation) at that stage, and believes with positive probability her preferences in both rounds will be identical.[9] We assume students know the distribution η over student types and lotteries (an assumption we need only for our characterization result, Theorem 3). Each student is assumed to learn her lottery number after the first round, as in practice students are often permitted to inquire about their position in 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 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 $\mu(\lambda) \neq \hat{\mu}(\lambda)$ and*[10] $\mu(\lambda) \neq s_{N+1}, \hat{\mu}(\lambda) \neq s_{N+1}$.

The majority of reassignments 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.

[9]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.

[10]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.

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 $\succ$ 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 $\hat{\succ}$ into a second-round assignment[11] $\hat{\mu}$. 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 Deferred Acceptance (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 the 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 her tentatively assigned school seat. The strict student preferences, weak school priorities, and the use of DA-STB mirror current practice in many school choice systems, such as those in New York City, Chicago, and Denver (see, e.g., Abdulkadiroglu and Sönmez, 2003).

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

Definition 3 (Deferred Acceptance (Azevedo and Leshno, 2016))**.** *The Deferred Acceptance mechanism with single tie-breaking (DA-STB) is a function* $\mathrm{DA}_\eta\big((\succ^\lambda, p^\lambda)_{\lambda\in\Lambda}, L\big)$ *mapping student preferences, priorities and lottery numbers into an assignment* μ*, defined by a vector of cutoffs* $\mathbf{C} \in \mathbb{R}_+^{\mathbf{N}}$

[11] 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.

as follows. Each student λ is given a score $r_i^\lambda = p_i^\lambda + L(\lambda)$ at school s_i and assigned to her most-preferred school as per her preferences, among those where her score exceeds the cutoff:

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

Moreover, $\mathbf{C}$ *is market-clearing, namely*

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

Azevedo and Leshno (2016) showed that the set of assignments satisfying equations (1) and (2) forms a non-empty complete lattice, and typically consists of a single uniquely determined assignment.[12] 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 in the priority class π 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, i.e., 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 she was assigned to in the first round to guarantee that each student receives a (weakly) better assignment.

[12] 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.

Definition 4 (Permuted Lottery Deferred Acceptance (PLDA) mechanisms)**.** *Let $P : [0,1] \to [0,1]$ 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: (1) assign a lottery number $P(l)$; and (2) 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* $\mathrm{DA}_{\hat{\eta}}\big((\hat{\succ}^\lambda, \hat{p}^\lambda)_{\lambda\in\Lambda}, P \circ L\big)$.

We use $\hat{C}^P_{\pi,i}$ 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 like that in NYC, where a large proportion of vacated seats are 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 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 her second-round assignment to her first-round assignment, that is, $\hat{\mu}(\lambda)\hat{\succeq}^{\lambda}\mu(\lambda)$ 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 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 $s_i \hat{\succ}^{\lambda} \hat{\mu}(\lambda)$, 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 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 her to report truthfully. To describe the properties formally, we start by fixing an arbitrary profile of first and second round preferences $(\succ^{-\lambda}, \hat{\succ}^{-\lambda})$ for all the students other than student λ. For any preference report of student λ in the first round she will receive an assignment that is probabilistic because of the tie-breaking lottery; then, after observing her first-round assignment and lottery number and her updated outside option, she can submit a second-round preference report, based on which 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, i.e., for each realization of lottery numbers (including her own lottery number) and profile of first- and second-round reported preferences of the students other than λ, reporting 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 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 she submits her first-round preference report and learns all lottery numbers only after the end of the first round. In that case, each student views her first-round assignment as a probability vector; 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 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 her top k schools in the second round is maximized when she reports truthfully in the first round (assuming truthful reporting in the second round), for each $k = 1, 2, \ldots, N$.*

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 like 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 **non-wasteful** if no student is assigned to a school she is eligible for that 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 $s_j \hat{\succ}^{\lambda} s_i$ 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.[13] This motivates the following definitions. Consider a second-round assignment $\hat{\mu}$. A *Pareto-improving cycle* is an ordered set of types $(\theta_1, \theta_2, \ldots, \theta_m) \in \Theta^m$, sets of students $(\Lambda_1, \Lambda_2, \ldots, \Lambda_m) \in \Lambda^m$, and schools $(\tilde{s}_1, \tilde{s}_2, \ldots, \tilde{s}_m) \in S^m$, such that $\eta(\Lambda_i) > 0$ and $\tilde{s}_{i+1} \hat{\succ}^{\theta_i} \tilde{s}_i$ (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 $\theta_{m+1} = \theta_1$).

[13]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).

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. *Suppose student preferences are consistent. Then PLDA mechanisms respect guarantees and priorities, and are strongly two-round strategy-proof, non-wasteful, constrained Pareto efficient, and anonymous.*

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 non-wasteful, constrained Pareto efficient, and anonymous. We make these claims formal in Appendix B 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 will 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 (amongst 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, i.e., for all* $s_i, s_j \in S$,

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

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, it 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 her first-round preferences in the second round, or drops out of the system entirely; student first-round preferences and school capacities are arbitrary. We establish the order condition and provide direct proofs of several of our theoretical results for this setting in Section 4, in order to give a flavor of the arguments employed to establish our results in the general setting.

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*

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

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 the permutation P. We remark that type equivalence does not imply an equal (or similar) amount of reassignment (e.g., see Figure 1), as type-equivalence depends only on the second-round assignment, while 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 reassigned students among PLDA mechanisms.*

Proof of Theorem 2. Fix $\theta = (\succ^\theta, \hat{\succ}^\theta, p^\theta) \in \Theta$ 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, as it never reassigns both a student of type θ into a school s_i and another student of type θ out of s_i.

Formally, for every permutation P, let the measures of students with type θ leaving and entering school s_i in the second round under 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. Due to 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 who left s_i in the second round of RLDA, which contradicts the reversal of the lottery. Formally, suppose $e_R > 0$ and $\ell_R > 0$. If $s_{N+1} \hat{\succ}^\theta s_i$ then $e_R = 0$, so we may assume $s_i \hat{\succ}^\theta s_{N+1}$. Since $e_R > 0$, there exists some student $\lambda \in \Lambda$ with type $\theta^\lambda = \theta$ for whom $s_i = \hat{\mu}^R(\lambda) \hat{\succ}^\theta \mu(\lambda)$. Since $\ell_R > 0$, there exists $\lambda' \in \Lambda$ with type $\theta^{\lambda'} = \theta$ for whom $s_j = \hat{\mu}^R(\lambda') \hat{\succ}^\theta \mu(\lambda') = s_i$. By consistency, we have $s_i \succ^\theta \mu(\lambda)$, and therefore λ wished to be assigned to $\mu(\lambda')$ in the first round and hence $L(\lambda') > L(\lambda)$. Note that since $s_i \hat{\succ}^\theta s_{N+1}$, it follows that $s_j \hat{\succ}^\theta s_i \hat{\succ}^\theta s_{N+1}$. Now, since λ' received a better second-round assignment under RLDA than λ and both λ and λ' were reassigned under RLDA, it follows that $R(L(\lambda')) > R(L(\lambda))$, which is a contradiction. Since $e_P = \ell_P + c \geq c$ and $e_P \geq 0$ this completes the proof. $\square$

Our results present a strong 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, since all PLDA mechanisms give type-equivalent assignmentsand 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).[14]

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 $(C_1, C_2) = \left(\frac{3}{4}, \frac{1}{2}\right)$.

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 $\left(\hat{C}_1^R, \hat{C}_2^R\right) = \left(1, \frac{3}{4}\right)$. In this case, the order condition holds

[14] 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.

and FLDA and RLDA are type-equivalent. It is simple to verify that both FLDA and RLDA assigns 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 $\left(\hat{C}_1^R, \hat{C}_2^R\right) = \left(\frac{7}{8}, 1\right)$. In this case, the order condition does not hold. Type equivalence also does not hold, since the FLDA and RLDA assignments are

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

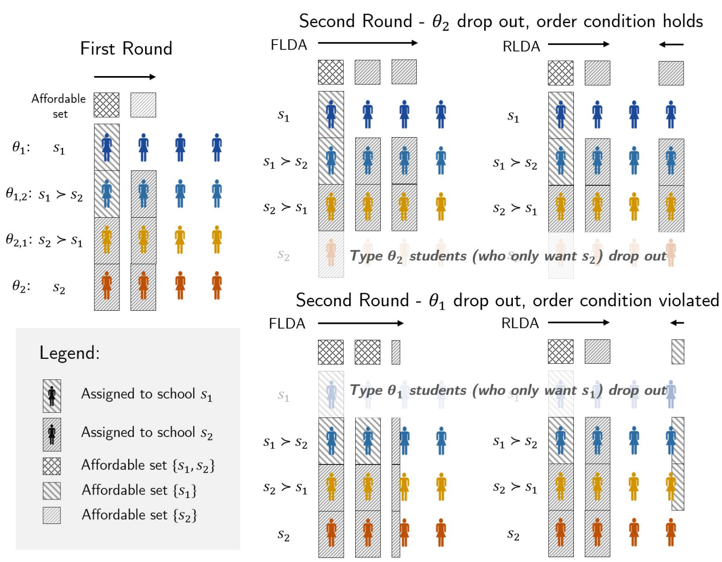

Figure 3: In Example 1, FLDA and RLDA are type-equivalent when the order condition holds, and give different assignments to students of *every* type when the order condition does not hold.

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.

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.*

Definition 15. *A two-round mechanism is* ***non-atomic*** *if any single student changing 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 & 2). A non-atomic two-round assignment mechanism with first round DA-STB respects guarantees and is: non-wasteful, (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.[15] 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 non-wastefulness, respecting priorities and guarantees, and anonymity, constrain the second round to be DA, with each student given a guarantee at the school she was assigned to in the first round, and constrained Pareto efficiency forces the tiebreaking 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 she can choose to attend, i.e., the first-round affordable set is the set of schools for which she meets the first-round cutoff, and the affordable set is the set of schools for which 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 her first-round affordable set and affordable set, and hence her second-round lottery is a permutation of her first-round lottery that does not depend on her preference type.

Our result mirrors similar large market cutoff characterizations for single-round mechanisms by

[15] 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 non-crossing property.

Liu and Pycia (2016) and Ashlagi and Shi (2014), which show, in settings with a 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, and furnish full proofs for a special case of our model to give the interested reader a taste of the general proof techniques in a more transparent setting. This section may be skipped at a first reading without loss of continuity.

We begin with some definitions and intuition for our general results. 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 her preferences. Specifically, if we define the *affordable set* for each student as the set of schools for which she meets either the first- or second-round cutoffs, then each student is assigned to her favorite school in her affordable set at the end of the second round, and changing the student's preferences does not change 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* $s_i, s_j \in S$,

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

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

1. *(Consistency aross rounds) PLDA(P) satisfies the local order condition on* (S, q, Λ, η) $\forall P$*;*
2. *(Consistency aross 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* $\hat{C}^P_{\pi,i} > \hat{C}^P_{\pi,j} \;\Rightarrow\; \hat{C}^{P'}_{\pi,i} \geq \hat{C}^{P'}_{\pi,j}$.

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)*.*

We provide some intuition as to why Theorem 4 holds by using the affordable set framework. 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 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.

The affordable set framework also sheds some light on the power of the global order condition. 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 $X_i = \{s_i, s_{i+1}, \ldots, s_N\}$ for some $i = i(\lambda)$, where schools are indexed in decreasing order of their cutoffs for the relevant priority group p^{θ^λ}, and the probability that a student receives some affordable set is independent of her preferences. Moreover, since affordable sets are nested $X_1 \supseteq X_2 \supseteq \cdots \supseteq X_N$, and since 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 now introduce a special case of our model. For this special case, we will prove that the order condition holds, and show 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.*

Before formalizing the definition and results for this setting, we provide some intuition for why the global order condition always holds under uniform dropouts. In the uniform dropouts model, each student drops out of the system with probability ρ, e.g. due to leaving 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; in particular, 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 $1-\rho$. Thus schools fill in the same order regardless of the choice of permutation.

Let us now formalize our definitions and results. Throughout the rest of this section, since there are no priorities, we will let student types be defined either by $\theta = (\succ^\theta, \hat{\succ}^\theta, \mathbf{1})$ or simply by $\theta = (\succ^\theta, \hat{\succ}^\theta)$. We define uniform dropouts with probability ρ by

$$\zeta(\{\theta = (\succ^\theta, \hat{\succ}^\theta) \in \Theta : \succ^\theta = \succ, \hat{\succ}^\theta = s_{N+1} \succ \ldots\}) = \rho\zeta(\{\theta = (\succ^\theta, \hat{\succ}^\theta) \in \Theta : \succ^\theta = \succ\}),$$

$$\zeta(\{\theta = (\succ^\theta, \hat{\succ}^\theta) \in \Theta : \succ^\theta = \succ, \hat{\succ}^\theta = \succ\}) = (1-\rho)\zeta(\{\theta = (\succ^\theta, \hat{\succ}^\theta) \in \Theta : \succ^\theta = \succ\}), \qquad (3)$$

i.e. 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.[16]

We show first that the global order condition (Definition 16) holds in the setting with uniform dropouts. The high level steps and algebraic tools used in this proof are similar to those used to show that the order condition is equivalent to the global order condition in our general framework (Theorem 4), although the analysis in each step is greatly simplified. We provide some intuition as to the differences in this section, and furnish the full proof of Theorem 4 in the Appendix.

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

Proof of Theorem 5. The main steps in the proof are as follows: (1) *Assuming that every student's affordable set is X_i for some i,* for every school s_j, guess the proportion of students who should receive an affordable set that contains s_j. (2) Calculate the corresponding second-round cutoffs $\tilde{C}_j$ for school s_j. (3) Show that these cutoffs are in the same order as the first-round cutoffs. (4) Use the fact that the cutoffs are in the same order to verify that the cutoffs are market-clearing, and deduce that the constructed cutoffs are precisely the PLDA(P) cutoffs.

Throughout this proof, we amend the second-round score of a student λ under PLDA(P) to be $\hat{r}_i^\lambda = P(L(\lambda)) + \mathbb{1}_{\{L(\lambda) \geq C_i\}}$, meaning that we give each student a guarantee at any school for which she met the cutoff in the first round. By consistency of preferences, it is easily seen that this has no effect on the resulting assignment or cutoffs. Let the first-round cutoffs be $C_1, C_2, \ldots, C_N$, where without loss of generality we index the schools such that $C_1 \geq C_2 \geq \cdots \geq C_N$.

(1) In the setting with uniform dropouts, since the second-round problem is a rescaled version of the first-round problem (with a $(1-\rho)$ fraction of the original students remaining), we guess that

[16] We remark that there is a well-known technical measurability issue w.r.t. a continuum of random variables, and that this issue can be handled; see, for example, Al-Najjar (2004).

we want the proportion of students with an affordable set containing s_j to be $\frac{1}{1-\rho}$ times the original proportion. (In the general setting, we no longer have a rescaled problem and so we instead guess that the proportion of students with each affordable set is the same as that under RLDA.)

(2) We translate this into cutoffs. Let $f_i^P(x) = |\{l : l \geq C_i \text{ or } P(l) \geq x\}|$ be the proportion of students who receive school s_i in their (second-round) affordable set with the amended second-round scores under permutation P if the first- and second-round cutoffs are C_i and x respectively. Notice that $f_i(x)$ is non-increasing for all i, $f_i(0) = 1$, $f_i(1) = 1 - C_i$, and if $i < j$ then $f_i(x) \leq f_j(x)$ for all $x \in [0,1]$. Let the cutoff $\tilde{C}_i^P \in [0,1]$ be the minimal cutoff satisfying the equation $f_i(\tilde{C}_i^P) = \frac{1}{1-\rho}(1 - C_i)$, and let $\tilde{C}_i^P = 0$ if $C_i < \rho$. (In the general setting the cutoffs are defined using the same functions $f_i^P(\cdot)$ with the proportions being equal to those that arise under RLDA, as mentioned in step (1) above.)

(3) We now show that the cutoffs $\tilde{\mathbf{C}}$ are in the right order. Suppose that $i < j$. If $\tilde{C}_i^P = 0$ then $C_j \leq C_i \leq \rho$ and so $\tilde{C}_j^P = 0 \leq \tilde{C}_i^P$ as required. Hence we may assume that $\tilde{C}_i^P, \tilde{C}_j^P > 0$. In this case, since $f_j(\cdot)$ is non-increasing and $\tilde{C}_j^P$ is minimal, we can deduce that $\tilde{C}_j^P \leq \tilde{C}_i^P$ if $f_j\left(\tilde{C}_j^P\right) \geq f_j\left(\tilde{C}_i^P\right)$. It remains to establish the latter. Using the definition of f_j and f_i, we have

$$\begin{aligned} f_j\left(\tilde{C}_i^P\right) &= f_i\left(\tilde{C}_i^P\right) + \left|\left\{l : l \in [C_j, C_i),\ P(l) < \tilde{C}_i^P\right\}\right| \\ &\leq \frac{1}{1-\rho}(1 - C_i) + (C_i - C_j) \leq \frac{1}{1-\rho}(1 - C_j) = f_j\left(\tilde{C}_j^P\right), \end{aligned}$$

where both inequalities hold since $C_j \leq C_i$. It follows that $\tilde{C}_i^P \geq \tilde{C}_j^P$, as required. (In the general setting, since we cannot give closed form expressions for the proportions $f_i\left(\tilde{C}_i^P\right)$ in terms of the cutoffs C_i, this step requires using the intermediate value theorem and an inductive argument.)

(4) We now show that $\tilde{\mathbf{C}}^P$ is the set of market-clearing DA cutoffs for the second round of PLDA(P). Note that $\gamma_i = C_{i-1} - C_i$ is the proportion of students whose first-round affordable set is X_i (where $C_0 = 1$). Since dropouts are uniform at random, this is the proportion of such students out of the

total number of remaining students both before and after dropouts.

Consider first the case $\tilde{C}_i^P > 0$. Now $f_i\left(\tilde{C}_i^P\right)$ is the proportion of students whose second-round affordable set contains s_i, and since $C_1 \geq C_2 \geq \cdots \geq C_N$ and $\tilde{C}_1^P \geq \tilde{C}_2^P \geq \cdots \geq \tilde{C}_N^P$, it follows that the affordable sets are nested. Hence the proportion of students (of those remaining after students drop out) whose second-round affordable set is X_i is given by (where $f_0(\cdot) \equiv 1$)

$$\gamma_i^P = f_i\left(\tilde{C}_i^P\right) - f_{i-1}\left(\tilde{C}_{i-1}^P\right) = \frac{C_{i-1} - C_i}{1-\rho} = \frac{\gamma_i}{1-\rho}.$$

For each $\theta = (\succ, \succ)$ and set of schools S, let $D^\theta(S)$ be the maximal school in S under $\succ$, and let $\theta' = (\succ, \hat{\succ})$ be the student type consistent with θ that finds all schools unacceptable in the second round. Then a set of students of measure

$$\sum_{j \leq i} \sum_{\theta \in \Theta : D^\theta(X_j) = i} \gamma_j^P \zeta(\theta) = \sum_{j \leq i} \gamma_j \sum_{\theta \in \Theta : D^\theta(X_j) = i} \frac{\zeta(\theta)}{1-\rho} = \sum_{j \leq i} \gamma_j \sum_{\theta \in \Theta : D^\theta(X_j) = i} \zeta(\theta) + \zeta(\theta')$$

choose to go to school s_i in the second round under the second-round cutoffs $\tilde{\mathbf{C}}^P$. We observe that the expression on the right gives the measure of the set of students who choose to go to school s_i in the first round under first-round cutoffs $\mathbf{C}$.

In the case where $\tilde{C}_i^P = 0$ the above expressions give upper bounds on the measure of the set of students who choose to go to school s_i in the second round under the second-round cutoffs $\tilde{\mathbf{C}}^P$. Since $\mathbf{C}$ are market-clearing cutoffs, and $\tilde{C}_i^P > 0 \Rightarrow C_i^P > 0$, it follows that $\tilde{\mathbf{C}}^P$ are market-clearing cutoffs too. We have shown that in PLDA(P), the second-round cutoffs are exactly the constructed cutoffs $\tilde{\mathbf{C}}^P$ and they satisfy $\tilde{C}_1^P \geq \cdots \geq \tilde{C}_N^P$, and so the global order condition holds. $\square$

The general proof of Theorem 4 uses the cutoffs for RLDA in steps (1) and (2) above to guess the proportion of students who receive an affordable set that contains s_j, and requires that each student priority type be carefully accounted for. However, the general structure of the proof is similar, and the tools used are straightforward generalizations of those used in the proof above.

Under uniform dropouts all PLDA mechanisms give type-equivalent assignments.

Proposition 2. *In any market with uniform dropouts (Definition 17), all PLDA mechanisms produce type-equivalent assignments.*

Proof. The proposition follows immedately from the fact that the proportion γ_i^P of students whose second-round affordable set is X_i does not depend on P.

Specifically, consider first the case when all schools reach capacity in the second round of PLDA. We showed in the proof of Theorem 5 that for all i and all student types θ, the proportion of students of type θ with affordable set X_i in the second round under PLDA(P) is given by $\gamma_i^P = \frac{\gamma_i}{1-\rho}$, where γ_i is the proportion of students of type θ with affordable set X_i in the first round. It follows that all PLDAs are "type-equivalent" to each other because they are type-equivalent to the first-round assignment in the following sense. For each preference order $\succ$, let $\tilde{\succ}$ be the preferences obtained from $\succ$ by making the outside option the most desirable, i.e., $s_{N+1} \tilde{\succ} \cdots$. Then

$$\begin{aligned}
\eta(\{\lambda \in \Lambda : \theta^\lambda = (\succ, \succ), \hat{\mu}^P(\lambda) = s_i\}) &= \frac{1}{1-\rho}\eta(\{\lambda \in \Lambda : \theta^\lambda = (\succ, \succ), \mu(\lambda) = s_i\}) \\
&= \eta(\{\lambda \in \Lambda : \theta^\lambda \in \{(\succ, \succ), (\succ, \tilde{\succ})\}, \mu(\lambda) = s_i\}),
\end{aligned}$$

where the second equality holds since students stay in the system uniformly-at-random with probability $1-\rho$. Under uniform dropouts this holds for all student types that remain in the system, and so it follows that $\hat{\mu}^P$ is type-equivalent to $\hat{\mu}^{P'}$ for all permutations P, P'.

When some school does not reach capacity in the second round, we can show by induction on the number of such schools that all PLDAs are type-equivalent to RLDA. □

Remark. Most of the results of this section extend to the following generalization of the uniform dropouts setting. 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, remain and wish to stay at their first round assignment with some fixed probability ρ' (have 'inertia'), and otherwise remain and retain their first round preferences.[17] It can be shown that in such a market, the global order condition always holds, and RLDA minimizes reassignment amongst 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 New York City (NYC) high school choice system to simulate and evaluate the performance of centralized PLDA mechanisms under different permutations P. The simulations indicate that our theoretical results are real-world relevant. Different choices of P are found to yield similar allocative efficiency: the number of students assigned to their k-th choice for each rank k, as well as 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 reassignments at every stage. However, in a version where students take time to decide on offers from the waitlist, the efficiency comparisons are reversed.[18] 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

[17] This market is slightly beyond the scope of our general model, as 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.

[18] 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 increase the waiting time for other students in the second round by considering multiple offers off the waitlist that they eventually decline.

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 years 2004–2005, 2005–2006, and 2006–2007, as follows.

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.[19]

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 does 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).[20]

School capacities and priorites. 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, since 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.)

[19]The algorithm is not completely strategy-proof, since 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, e.g., Hassidim et al. (2015) and Narita (2016).

[20]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 they drop out randomly at a rate equal to the dropout rate for the rest of the students ($8.9\% - 9.2\%$).

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).

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 i.i.d. first-round lottery number $L(\lambda)$ (a normal variable with mean 0 and variance 1), which generates a uniformly random lottery order.[21] 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 the 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.

Decentralized PLDA. In order to evaluate the performance of waitlist systems, we also ran simulations using two versions of decentralized PLDA with second rounds run in multiple "stages":

Version 1. At stage ℓ, school s_i has residual capacity $\tilde{q}_i^\ell$ equal to the number of students previously assigned to school s_i who rejected school s_i in the previous stage (and $\tilde{q}_i^1$ is the number of students assigned to school s_i in the first round who dropped out of the system). Each school s_i proposes to

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

the top $\tilde{q}_i^\ell$ students on their waitlist (including students who dropped out) and removes them from the waitlist, students who dropped out reject all offers, and all remaining students are (tentatively) assigned to their favorite school that offered them a seat in the first round or in the second round thus far and reject the rest. The stages of reassignment continue until there are no new proposals.

Version 2. At stage ℓ, school s_i has residual capacity $\tilde{q}_i^\ell$ equal to the number of students previously assigned to school s_i who rejected school s_i in the previous stage (and $\tilde{q}_1^\ell$ is the number of students assigned to school s_i in the first round who dropped out of the system). We run DA-STB on the residual economy where each school s_i has capacity $\tilde{q}_i^\ell$ and each student only finds schools strictly better than their current assignment acceptable.[22] This results in some students being reassigned and new residual capacities for stage $\ell + 1$, equal to the sum of the number of unfilled seats at the end of stage ℓ and the number of students who left the school due to an upgrade in stage ℓ. The stages of reassignment continue until there are no new proposals.

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 amount of time to accept an offer, to reject an offer, or to inform a school that they were previously assigned to that they have been assigned to a different school. Version 2 captures a decentralized process where students also take time to both make and communicate decisions, but take much longer to tell schools that they were previously assigned to that they have been assigned to a different school. Accordingly the efficiency outcomes at a given stage of version 2 dominate those of version 1 at the same stage, as more information is communicated during each stage.

Version 2 simulations a setting where the main driver behind congestion is chains of student reassignment. Version 2 is more realistic in settings where schools are the primary drivers behind updated information, since a school is much more likely to ask for decisions from students who are

[22]We provide results using school-proposing DA, as this more closely mirrors the structure of waitlist systems. Results using student-proposing DA were similar.

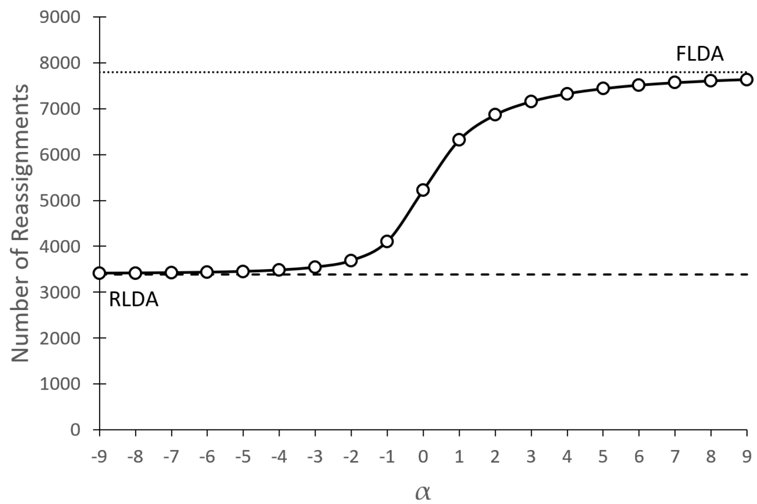

Figure 4: Number of reassigned students versus α. The number of reassigned students under the extreme values of α, namely, $\alpha = \infty$ (FLDA) and $\alpha = -\infty$ (RLDA), are shown via dotted lines.

undecided about an offer from the school rather than from students who have already accepted an offer from the school. In many school districts information about previously assigned students being reassigned to other schools is processed centrally, and it is also reasonable to assume that this would occur on a slower timescale than rejections of offers. 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. Results for 2005–2006 and 2006–2007 were similar. Figure 4 shows that the mean number of reassignments is minimized at $\alpha = -\infty$ (RLDA) and increases with α, which is consistent with our theoretical result in Theorem 2. The mean number of reassignments is as large as 7,800 under FLDA compared to 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 $1 \leq k \leq 12$, as well as the number of unassigned students,

α	Reassignments #	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: ∞	7797	5.89	55.41	69.85	78.03
8.00	7606	5.90	55.40	69.85	78.02
6.00	7512	5.90	55.40	69.85	78.03
4.00	7325	5.89	55.38	69.84	78.02
2.00	6863	5.89	55.33	69.81	78.02
0.00	5220	5.87	54.96	69.65	77.97
-2.00	3686	5.81	54.52	69.37	77.82
-4.00	3480	5.79	54.47	69.33	77.78
-6.00	3433	5.79	54.46	69.32	77.77
-8.00	3416	5.79	54.45	69.31	77.77
RLDA: $-\infty$	3391	5.79	54.45	69.30	77.75

Table 1: Centralized PLDA simulation results: 2004–2005 NYC high school admissions.

We show the mean percentage of students remaining unassigned or getting at least their kth 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 only about 100 students.

vary by less than 1% of the total number of students. There is a slight trade-off between allocative efficiency due to 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 computation 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

[23] Intuitively, prioritizing students with lower lotteries both decreases the number of unassigned students and decreases allocative efficiency by artificially increasing the constraints from providing first-round guarantees.

α	Reassignments # total (# temporary)	Unassigned %	$k = 1$ %	$k \leq 2$ %	$k \leq 3$ %
Round 1 (No Reassignments)	0	9.31	50.14	64.17	72.45
Round 2 FLDA, Version 1					
Stage 1	3461 (447)	7.89	52.68	66.62	74.47
Stage 2	2126 (206)	7.04	53.93	68.03	76.14
Stage 3	1258 (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 $\approx$ 17)	8590 (780)	5.87	55.46	69.87	78.05
Round 2 RLDA, Version 1					
Stage 1	1004 (835)	7.85	51.38	65.70	74.09
Stage 2	1077 (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 $\approx$ 33)	5818(2419)	5.79	54.51	69.37	77.80
Round 2 FLDA, Version 2					
Stage 1	4139 (457)	7.62	53.21	67.14	75.21
Stage 2	2333 (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 $\approx$ 12)	8503 (677)	5.89	55.47	69.87	78.04
Round 2 RLDA, Version 2					
Stage 1	2863 (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 $\approx$ 9)	3624 (220)	5.79	54.51	69.33	77.76

Table 2: Decentralized PLDA simulation results: 2004–2005 NYC high school admissions.

We show the mean number of reassignments (number of movements of a student from a school in S to a different school in S) as well as the mean number of temporary reassignments (number of movements of a student from a school in $S \cup \{s_{N+1}\}$ to a 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 kth 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.

movement of a student from a school in $S \cup \{s_{N+1}\}$ to a different school in S that is not their 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 reassigns more students than RLDA but far outperforms RLDA in terms of minimizing congestion and maximizing efficiency. FLDA takes on average 17 stages to converges, while RLDA requires 33. FLDA performs 780 temporary transfers while RLDA performs 2420, creating much more unnecessary congestion. FLDA takes 2 and 5 stages to achieve 50% and 90% respectively of the total increase in number of students assigned to their top school, whereas RLDA takes 3 and 9 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.

In the second version of decentralized PLDAs, FLDA still reassigns 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 2 stages, and achieves most of its allocative efficiency by the second stage, improving the allocative efficiency by fewer than 100 students from that point onwards. In the limit FLDA is still slightly more efficient than RLDA, and 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 to 9 for RLDA.

Our empirical findings have mixed implications for implementing decentralized waitlists. Our clearest finding is the benefit of centralization in reducing 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 one 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%.[24]

6 Proposals & Discussion

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

[24] These upper bounds were computed as follows. Approximately 2700 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 students 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 their true preferences. We then counted the number of such students who by doing so could either (1) change their first-round assignment (for the worse) but second-round assignment for the better, or (2) 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.

RLDA at the start of the school year, as 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 reassignments required to reach an efficient assignment.

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, as 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, an 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.[25]

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 the relative overdemand for schools stays the same this choice can be made without sacrificing allocative efficiency.

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

[25] 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 room, her priority drops to the last place of the queue.

round is independent of their priority at other schools. This new set of axioms describes a class of mechanisms that strictly includes the PLDA mechanisms. 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.

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, her report should affect the second-round cutoffs in a manner that gives her a second-round assignment she would not have received otherwise. If the market is large enough, the cutoffs will converge to their limiting values, and the probability that she 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.[26]

[26]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

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, although only about 7% of students reapplied, about 70% of these reapplicants reported second-round preferences that were inconsistent with their first-round reported preferences. Note that PLDAs allow students to report inconsistent preferences in the second round. We believe that some of our insights remain valid if a small fraction of students have an idiosyncratic change in preferences, or if a small number of new students enter in the second round. 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 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.

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.

References

A. Abdulkadiroglu and T. Sönmez. Random serial dictatorship and the core from random endowments in house allocation problems. *Econometrica*, 66(3):689–701, 1998.

A. Abdulkadiroglu and T. Sönmez. House allocation with existing tenants. *Journal of Economic Theory*, 88(2):233–260, 1999.

A. Abdulkadiroglu and T. Sönmez. School choice: A mechanism design approach. *American Economic Review*, 93(3):729–747, 2003.

A. Abdulkadiroglu and T. Sönmez. *Matching Markets: Theory and Practice*. Cambridge University Press, 2011. ISBN 9781139060011.

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.

A. Abdulkadiroglu, P. A. Pathak, and A. E. Roth. The New York City high school match. *American Economic Review*, 95(2):364–367, 2005a.

A. Abdulkadiroglu, P. A. Pathak, A. E. Roth, and T. Sönmez. The Boston public school match. *American Economic Review*, 95(2):368–371, 2005b.

A. Abdulkadiroglu, P. A. Pathak, and A. E. Roth. Strategy-proofness versus efficiency in matching with indifferences: Redesigning the NYC high school match. *American Economic Review*, 99(5): 1954–78, December 2009.

A. Abdulkadiroglu, N. Agarwal, and P. A. Pathak. The welfare effects of coordinated assignment: Evidence from the NYC HS match. Technical report, National Bureau of Economic Research, 2015.

N. Agarwal and P. Somaini. Demand analysis using strategic reports: An application to a school choice mechanism. Technical report, National Bureau of Economic Research, 2014.

N. I. Al-Najjar. Aggregation and the law of large numbers in large economies. *Games and Economic Behavior*, 47(1):1–35, 2004.

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

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

N. Arnosti. Short lists in centralized clearinghouses. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation, EC '15, Portland, OR, USA, June 15–19, 2015*, page 751, 2015.

N. Arnosti and P. Shi. How (not) to allocate affordable housing. 2017.

I. Ashlagi and A. Nikzad. What matters in school choice tie-breaking? How competition guides design. In *Proceedings of the Seventeenth ACM Conference on Economics and Computation, EC '16*, 2016.

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

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

I. Ashlagi, A. Nikzad, and A. Romm. Assigning more students to their top choices: A tiebreaking rule comparison. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation, EC '15, Portland, OR, USA, June 15–19, 2015*, pages 755–756, 2015.

I. Ashlagi, M. Burq, P. Jaillet, and V. Manshadi. On matching and thickness in heterogeneous dynamic markets. 2017.

E. M. Azevedo and E. B. Budish. Strategy-proofness in the large. *Chicago Booth Research Paper*, 2013.

E. M. Azevedo and J. D. Leshno. A supply and demand framework for two-sided matching markets. *Journal of Political Economy*, 124(5):1235–1268, 2016.

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

J. Combe, O. Tercieux, and C. Terrier. The design of teacher assignment: Theory and evidence. In *ACM Conference on Economics and Computation*, 2016.

O. Compte and P. Jehiel. Voluntary participation and reassignment in two-sided matching. Working paper, 2008.

J. P. Dickerson and T. Sandholm. Futurematch: Combining human value judgments and machine learning to match in dynamic environments. In *AAAI*, pages 622–628, 2015.

U. Dur. Dynamic school choice problem. Working paper, 2012.

A. Erdil and H. Ergin. What's the matter with tie-breaking? Improving efficiency in school choice. *The American Economic Review*, 98(3):669–689, 2008.

D. Gale and L. S. Shapley. College admissions and the stability of marriage. *The American Mathematical Monthly*, 69(1):9–15, 1962.

A. Hassidim, D. Marciano-Romm, A. Romm, and R. I. Shorrer. Strategic behavior in a strategy-proof environment. Working paper, 2015.

S. V. Kadam and M. H. Kotowski. Multi-period matching. Working paper, 2014.

E. H. Kaplan. Analyzing tenant assignment policies. *Management science*, 33(3):395–408, 1987.

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

Q. Liu and M. Pycia. Ordinal efficiency, fairness, and incentives in large markets. Working paper, 2016.

Y. Narita. Match or mismatch: Learning and inertia in school choice. Working paper, 2016.

P. A. Pathak. The mechanism design approach to student assignment. *Annual Review of Economics*, 3(1):513–536, 2011.

J. S. Pereyra. A dynamic school choice model. *Games and Economic Behavior*, 80:100–114, 2013.

A. E. Roth, T. Sönmez, and M. U. Ünver. Kidney exchange. *The Quarterly Journal of Economics*, 119(2):457–488, 2004.

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

J. Szuflita. Tale of two waitlists: there are too many people at your beloved zoned school. http://www.nycschoolhelp.com/blog/2016/3/16/tale-of-two-waitlists, a. March 16, 2016.

J. Szuflita. Tale of two waitlists 2017. http://www.nycschoolhelp.com/blog/2017/3/14/tale-of-two-wait-lists-2017, b. March 14, 2017.

K. T. Talluri and G. J. Van Ryzin. *The Theory and Practice of Revenue Management*, volume 68. Springer Science & Business Media, 2006.

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

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

Online Appendix

A List of Notation

A.1 Model

- $S = \{s_1, \dots, s_N\}$: schools
- s_{N+1}: outside option
- q_i: capacity of school i
- $\Lambda = \Theta \times [0,1]$: set (continuum) of students
- η: measure over Λ
- $\theta = (\succ^\theta, \hat{\succ}^\theta, p^\theta)$: student types
- Θ: space of student types θ
- $\zeta(\theta)$: measure of students with type θ
- L: student lottery numbers
- n_i: the number of priority groups at school s_i

A.2 Mechanisms

- P: permutation
- μ: first-round assignment
- $\hat{\mu}$: second-round assignment
- $\hat{\mu}_P$: second-round assignment from PLDA with permutation P
- $\mathbf{C}$: first-round cutoffs
- $\hat{\mathbf{C}}^P$: second-round cutoffs from PLDA with permutation P
- $\mathbf{C}_\pi$: first-round cutoffs restricted to priority class π
- $\hat{\mathbf{C}}_\pi^P$: second-round cutoffs from PLDA with permutation P restricted to priority class π

A.3 Proof of Theorem 2

- $\ell_P = \eta(\{\lambda \in \Lambda : \theta^\lambda = \theta, \mu(\lambda) = s_i, \hat{\mu}_P(\lambda) \neq s_i\})$: the measure of students with type θ leaving school s_i in the second round under PLDA with permutation P
- $e_P = \eta(\{\lambda \in \Lambda : \theta^\lambda = \theta, \mu(\lambda) \neq s_i, \hat{\mu}_P(\lambda) = s_i\})$: the measure of students with type θ entering school s_i in the second round under PLDA with permutation P

A.4 Proofs for Uniform Dropouts (Section 4)

- ρ: probability that a student drops out
- $\tilde{\mathbf{C}}$: constructed second-round cutoffs
- $f_i^P(x)$: proportion of students with s_i in their affordable set with permutation P and first- and second-round cutoffs (C_i, x)
- γ_i^P: the fraction of students whose affordable set in the second round of PLDA with permutation P is X_i

A.5 Notation in the Appendix

- $\hat{r}_i^\lambda = P(L(\lambda)) + n_i \mathbb{1}_{\{L(\lambda) \geq C_i\}} + p_i^\lambda \mathbb{1}_{\{L(\lambda) < C_i\}}$: the amended second-round score of student λ under PLDA
- $X_i = \{s_i, \ldots, s_{N+1}\}$: schools (weakly) after s_i in the cutoff ordering
- γ_i: the proportion of students whose first-round affordable set is X_i

A.6 Proof of Theorems 1 and 4

- $\beta_{i,j} = \eta(\{\lambda \in \Lambda : \operatorname{argmax}_{\succ^\lambda} X_j = s_i\})$: the measure of students who, when their set of affordable schools is X_j, will choose s_i
- $E^\lambda(\mathbf{C})$: the set of schools affordable for type λ in the first round under PLDA with permutation P
- $\hat{E}^\lambda(\hat{\mathbf{C}}^P)$: the set of schools affordable for type λ in the second round under PLDA with permutation P

- $\gamma_i^P = \eta(\{\lambda \in \Lambda : \hat{E}_P^\lambda(\hat{\mathbf{C}}^P) = X_i\})$: the fraction of students whose affordable set in the second round of PLDA with permutation P is X_i
- q_π: restricted capacity vector for priority class π
- Λ_π: set of students with priority class π
- η_π: restriction of η to students with priority class π
- $\mathcal{E}_\pi = (S, q_\pi, \Lambda_\pi, \eta_\pi)$: restricted primitives for priority class π
- $s_{\sigma_\pi(i)}$: i-th school under second-round overdemand ordering for $\mathcal{E}_\pi$
- $\tilde{\mathbf{C}}^P$: second-round cutoffs defined for PLDA with the amended second-round scores from the RLDA cutoffs $\hat{\mathbf{C}}^R$
- $\tilde{\mathbf{C}}_\pi^P$: second-round cutoffs defined for PLDA on $\mathcal{E}_\pi$ with the amended second-round scores from the RLDA cutoffs $\hat{\mathbf{C}}^R$
- $\hat{n}$: smallest index of a school affordable to everyone

A.7 Proof of Theorem 3

- $s_{\sigma(i)}$: i-th school under second-round overdemand ordering in a non-atomic mechanism M satisfying axioms (1)–(5)
- $\tilde{X}_i = \{s_{\sigma(i)}, s_{\sigma(i+1)}, \dots, s_{\sigma(N+1)}\}$: schools (weakly) after $s_{\sigma(i)}$ in the second-round overdemand ordering
- $\gamma_{i,j}$: proportion of students under the constructed PLSM whose first-round affordable set was X_i and whose second-round affordable set was $\tilde{X}_j$
- $i(S') = \max\{j \ : \ s_j \in S'\}$: the maximum index of a school in S'
- $I_i^j = [C_i, C_j]$: the first-round scores that give students first-round affordable sets $\{X_{j+1}, X_{j+2}, \dots, X_i\}$
- $\rho^\theta(I, S')$: the proportion of students with type θ who, under the mechanism M, have a first-round score in the interval I and are assigned to a school in S' in the second round

B PLDA for a Discrete Set of Students

In this section we formally define and show how to implement PLDA mechanisms in a discrete setting with a finite number of students, and prove that they retain almost all the desired incentive and efficiency properties discussed in Section 2.2.

B.1 Discrete Model

A finite set $\Lambda = \{1, 2, \dots, n\}$ of students are to be assigned to a set $S = \{s_1, \dots, s_N\}$ of schools. Each student can attend at most one school. As in the continuum model, for every school $s_i \in S$, let $q_i \in \mathbb{N}_+$ be the *capacity* of school s_i, i.e., the number of students the school can accommodate. Let $s_{N+1} \notin S$ denote the outside option, and assume $q_{N+1} = \infty$. For each set of students $A \subseteq \Lambda$ we let $\eta(A) = |A|$ be the number of students in the set. As in the continuum model, each student $\lambda = (\theta^\lambda, L(\lambda)) \in \Lambda$ has a type $\theta^\lambda = (\succ^\lambda, \hat{\succ}^\lambda, p^\lambda)$ and a first-round lottery number $L(\lambda) \in [0, 1]$, which encode both student preferences and school priorities. The first-round lottery numbers $L(\lambda)$ are i.i.d. random variables drawn uniformly from $[0, 1]$ and do not depend on preferences. These random lottery numbers L generate a uniformly random permutation of the students based on the order of their lottery numbers.

An *assignment* $\mu : \Lambda \to S$ specifies the school that each student is assigned to. For an 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. As in the continuum model, we say that a student $\lambda \in \Lambda$ is a *reassigned student* if she is assigned to a school in S in the second round that is different to her first-round assignment.

B.2 PLDA Mechanisms & Their Properties

We now formally define PLDA mechanisms in a setting with a finite number of students. In order to do so, we use the algorithmic description of DA and extend it to a two-round setting. This also provides a clear way to implement PLDA mechanisms in practice.

We first reproduce the celebrated and widely deployed DA algorithm, and then proceed to define PLDAs.

Definition 18. *The Deferred Acceptance algorithm with single tie-breaking is a function* $DA\big((\succ^\lambda, p^\lambda)_{\lambda\in\Lambda}, L\big)$ *mapping the student preferences in the first round, priorities and lottery numbers into* *an assignment μ constructed as follows. In each step, unassigned students apply to their most-preferred school that has not yet rejected them. A school with a capacity of q* tentatively accepts *its q highest-ranked applicants, ranked according to its priority ranking of the students with ties broken by giving preference to higher lottery numbers L (or tentatively accepts all applicants, if fewer than q have applied), and rejects any remaining applicants, and the algorithm moves on to the next step. The algorithm runs until there are no new student applications, at which point it terminates and assigns each student to her tentatively assigned school seat.*

Definition 19 (Permuted Lottery Deferred Acceptance (PLDA) mechanisms)**.** *Let P be a permutation of Λ. 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. The permuted lottery deferred acceptance mechanism associated with P (PLDA(P)) is the mechanism that then computes a second-round assignment $\hat{\mu}^P$ by running DA on the same set of students Λ but with student preferences $\hat{\succ}$, a modified lottery $P \circ L$, and modified priorities $\hat{p}$ that give each student top priority at the school she was assigned to in the first round. Specifically, each school s_i's priorities $\hat{\succ}_i$ are defined by lexicographically ordering the students first by whether they were assigned to s_i in the first round, and then according to p_i. PLDA(P) is the two-round mechanism obtained from using the reassignment mechanism $DA\big((\hat{\succ}^\lambda, \hat{p}^\lambda)_{\lambda\in\Lambda}, P \circ L\big)$.*

We now formally define desirable properties from Section 2.2 in our discrete model. We remark that the definitions of respecting guarantees, strategy-proofness and anonymity do not reference school capacities and so carry over immediately. Similarly, the definitions for respecting priorities, non-wastefulness and constrained Pareto efficiency do not require non-atomicity and so our definition of η ensures that they also carry over. For completeness, we rewrite these properties without reference to η.

Definition 20. *A two-round mechanism M* ***respects priorities (subject to guarantees)*** *if (i) for every school $s_i \in S$ and student $\lambda \in \Lambda$ who prefers s_i to her assigned school $s_i \mathrel{\hat{\succ}^\lambda} \hat{\mu}(\lambda)$, we have $|\hat{\mu}(s_i)| = q_i$, and (ii) for all students λ' such that $\hat{\mu}(\lambda') = s_i \neq \mu(\lambda')$, we have $p_i^{\lambda'} \geq p_i^{\lambda}$.*

Definition 21. *A two-round mechanism is* ***non-wasteful*** *if no student is denied a seat at a school that has vacant seats; that is, for each student $\lambda \in \Lambda$ and school s_i, if $s_i \hat{\succ}^\lambda \hat{\mu}(\lambda)$, then $|\hat{\mu}(s_i)| = q_i$.*

Let $\hat{\mu}$ be a second-round assignment. A *Pareto-improving cycle* is an ordered set 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 $\tilde{s}_{i+1} \hat{\succ}^i \tilde{s}_i$ (where $\hat{\succ}^i$ denotes the second-round preferences of student λ_i, and we define $\tilde{s}_{m+1} = \tilde{s}_1$), and $\hat{\mu}(\lambda_i) = \tilde{s}_i$ for all 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 $\hat{p}_{\tilde{s}_{i+1}}^{\lambda_i} \geq \hat{p}_{\tilde{s}_{i+1}}^{\lambda_{i+1}}$ for all i (where we define $\lambda_{m+1} = \lambda_1$).

Definition 22. *A two-round mechanism is* ***constrained Pareto efficient*** *if the second-round assignment has no Pareto-improving cycles that respect second-round priorities.*

In a setting with a finite number of students, PLDA mechanisms exactly satisfy all these properties except for strategy-proofness.

Proposition 3. *Suppose student preferences are consistent. Then PLDA mechanisms respect guarantees and priorities, and are non-wasteful, constrained Pareto efficient, and anonymous.*

Proof. The proofs of all these properties are almost identical to those in the continuum setting.

As an illustration, we prove that PLDA is constrained Pareto efficient in the discrete setting by using the fact that both rounds use single tie-breaking and the output is stable with respect to the second-round priorities $\hat{p}$.

Fix a Pareto-improving cycle C. Since λ_i is assigned a seat at a school s_i when she prefers $s_{i+1} = \mu(\lambda_{i+1})$, by the stability of DA she must either be in a strictly worse priority group than λ_{i+1} at school s_{i+1}, or in the same priority group but have a worse lottery number. If C respects (second-round) priorities, then it must hold that for all i that students λ_i and λ_{i+1} are in the same priority group at school s_{i+1} and λ_i has a worse lottery number than λ_{i+1}. But since this holds for all i, single tie-breaking implies that we obtain a cycle of lottery numbers, which provides the necessary contradiction. $\square$

Proposition 3 states that in a setting with a finite number of students, PLDA mechanisms satisfy all our desired properties except for strategy-proofness. The following example illustrates that in a setting with a finite number of students, PLDA mechanisms may not satisfy two-round strategy-proofness. The intuition is that without non-atomicity, students are able to manipulate the first-round assignments of other students to change the guarantees, and hence change the second-round stability structure. In some cases in small markets, students are able to change the set of stable outcomes to benefit themselves.

Example 2 (PLDA with finite number of students is not strategy-proof.). *Consider a setting with $N = 2$ schools and $n = 4$ students. Each school has capacity 1 and a single priority class. For readability, we let $\emptyset$ denote the outside option, $\emptyset = s_{N+1} = s_3$. The students have the following preferences:*

1. $s_1 \succ_1 \emptyset \succ_1 s_2$ *and* $\emptyset \hat{\succ}_1 s_1 \hat{\succ}_1 s_2$,
2. $s_1 \succ_2 s_2 \succ_2 \emptyset$, *second-round preferences identical,*
3. $s_2 \succ_3 s_1 \succ_3 \emptyset$, *second-round preferences identical,*
4. $s_2 \succ_4 \emptyset \succ_4 s_1$, *second-round preferences identical.*

We show that the two-round mechanism where the second round is the reverse lottery deferred acceptance mechanism is not strategy-proof.

Consider the lottery that yields $L(1) > L(2) > L(3) > L(4)$. If the students report truthfully, the first-round assignment and second-round reassignment are

$$\begin{aligned} \mu(\Lambda) &= (\mu(1), \mu(2), \mu(3), \mu(4)) = (s_1, s_2, \emptyset, \emptyset), \text{ and} \\ \hat{\mu}(\Lambda) &= (\hat{\mu}(1), \hat{\mu}(2), \hat{\mu}(3), \hat{\mu}(4)) = (\emptyset, s_2, s_1, \emptyset) \end{aligned}$$

respectively. However, consider what happens if student 2 says that only school 1 is acceptable to her by reporting preferences $\succ^{\mathrm{r}}, \hat{\succ}^{\mathrm{r}}$ given by $s_1 \succ^{\mathrm{r}} \emptyset \succ^{\mathrm{r}} s_2$ and $s_1 \hat{\succ}^{\mathrm{r}} \emptyset \hat{\succ}^{\mathrm{r}} s_2$. Then

$$\mu(\Lambda) = (s_1, \emptyset, s_2, \emptyset), \quad \hat{\mu}(A) = (\emptyset, s_1, s_2, \emptyset),$$

which is a strictly beneficial change for student 2 in the second round (and, in fact, weakly beneficial for all students).

Note that this reassignment was not stable in the second round when students reported truthfully, since, in that case, school s_2 had second-round priorities $p_2^2 > p_2^4 > p_2^3 > p_2^1$ and so school s_2 and student 4 formed a blocking pair. In other words, for this particular realization of lottery numbers, student 2 is able to select a beneficial second-round assignment $\hat{\mu}$ that was previously unstable by changing student 3's first-round assignment so that student 4 cannot block $\hat{\mu}$.

In addition, the second-round outcome for student 2 under misreporting stochastically dominates her outcome from truthful reporting, when all other students report truthfully and the randomness is due the first-round lottery order. For if the lottery order is $L(1) > L(2) > L(3) > L(4)$ then *student 2 can change her second-round assignment from s_2 to s_1 by reporting s_2 as unacceptable, and this is the only lottery order for which student 2 receives a second-round assignment of s_2 under truthful reporting.*[27] *Moreover, for any lottery order where student 2 received s_1 in the first or second round under truthful reporting, she also received s_1 in the same round by misreporting.*[28] *Hence the second-round assignment student 2 receives by misreporting stochastically dominates the assignment she would have received under truthful reporting. This violates strategy-proofness.*

This example shows that, as noted in Section 6, PLDA mechanisms are not two-round strategy-proof in the finite setting. However, there are convergence results in the literature that suggest that PLDA mechanisms are almost two-round strategy-proof in large markets. We conjecture that the proportion of students who are able to successfully manipulate PLDA mechanisms decreases polynomially in the size of the market; a formal proof of such a result is beyond the scope of this paper.

Moreover, we believe that students will be unlikely to try to misreport under the PLDA mechanisms.[29] This is because, as Example 2 illustrates, successful manipulations require that students strategically change their first-round assignment and correctly anticipate that this changes the set of second-round stable assignments to their benefit. Such deviations are very difficult to plan and require sophisticated strategizing and detailed information about other students' preferences.

C Proofs

We begin with some general notation and definitions. Let μ be the initial assignment under DA-STB, and let P be a permutation. We say that a school s_i *reaches capacity* under a mechanism with output assignment μ if $\eta(\mu(s_i)) = q_i$.

We re-index the schools in $S \cup \{s_{N+1}\}$ so that $C_i \geq C_{i+1}$. Moreover, we assume that this indexing is done such that if the order condition is satisfied, then $\hat{C}_i^P \geq \hat{C}_{i+1}^P$ (where the cutoffs $\hat{\mathbf{C}}^P$ are as defined by PLDA(P)) holds simultaneously for all permutations P.

Recall that in DA each student is given a score $r_i^\lambda = p_i^\lambda + L(\lambda)$, and in PLDA$(P)$ this leads to

[27] This is because if $L(\lambda) > L(2)$ for $\lambda = 3, 4$ then student λ is assigned to school s_2 and stays there in both rounds, if $L(2) > L(1), L(3), L(4)$ then student 2 is assigned to school s_1 and stays there in both rounds, and finally if $L(1) > L(2) > L(4) > L(3)$ then student 2 is assigned to s_1 in the second round.

[28] This is because any stable matching in which student 2 is assigned s_1 remains stable after student 2 truncates. Indeed, student 2 is not part of any unstable pair, as she got her first choice, and any unstable pair not involving student 2 remains unstable under the true preferences, as only student 2 changes her preferences.

[29] as compared to the currently used DA mechanism.

a second-round score $\hat{r}_i^\lambda = \hat{p}_i^\lambda + P(L(\lambda)) = P(L(\lambda)) + n_i \mathbb{1}_{\{\mu(\lambda)=s_i\}} + p_i^\lambda \mathbb{1}_{\{\mu(\lambda)\neq s_i\}}$. Throughout the Appendix, for convenience, we slightly change the second-round score of a student λ under PLDA with permutation P to be $\hat{r}_i^\lambda = P(L(\lambda)) + n_i \mathbb{1}_{\{L(\lambda)\geq C_i\}} + p_i^\lambda \mathbb{1}_{\{L(\lambda)<C_i\}}$, meaning that we give each student a guarantee at any school for which she met the cutoff in the first round. By consistency of preferences, it is easily seen that this has no effect on the resulting assignment or cutoffs.

We say that a student can *afford* a school in a round if her score in that round is at least as large as the school's cutoff in that round. We say that the set of schools a student can afford in the second round (with her amended second-round score) is her *affordable set*.

Throughout the Appendix, we let $X_i = \{s_i, \dots, s_{N+1}\}$ be the set of schools at least as affordable as school s_i, and we let γ_i be the proportion of students whose first-round affordable set is X_i.

C.1 Proof of Proposition 1

Fix a permutation P and some PLDA with permutation P. We show that this particular PLDA satisfies all the desired properties. Let η be a distribution of students, and let $\hat{\mathbf{C}}^P$ be the second-round cutoffs corresponding to the assignment given by the PLDA for this distribution of student types.

PLDA respects guarantees because fewer students are guaranteed at each school than the capacity of the school. PLDA is non-wasteful because the second round terminates with a stable matching where all schools find all students acceptable, which is non-wasteful.

We now show that the PLDA mechanism is strongly two-round strategy-proof. Since students are non-atomic, no student can change the cutoffs $\hat{\mathbf{C}}^P$ by changing her first- or second-round reports. Hence it is a dominant strategy for all students to report truthfully in the second round. Moreover, for any student of type λ, the only difference between having a first-round guarantee at a school s_i and having no first-round guarantee is that in the former case, $\hat{r}_i^\lambda$ increases by $n_i - p_i^\lambda$. This means that having a guarantee at a school s_i changes the student's second-round assignment in the following way. She receives a seat in school s_i whereas without the guarantee she would have received a seat in some school s_j that she reported preferring less to s_i, and her second-round assignment is unchanged otherwise. Therefore, students want their first-round guarantee to be the best under their second-round preferences, and so it is a dominant strategy for students with consistent preferences to report truthfully in the first round.

PLDA is constrained Pareto efficient, since we use single tie-breaking and the output is the student-optimal stable matching with respect to the updated second-round priorities $\hat{p}$.

This is easily seen via the cutoff characterization. Let the second-round cutoffs be $\hat{P}$, where overloading notation we let $\hat{P}_i$ denote the cutoff for school $\tilde{s}_i$. Fix a Pareto-improving cycle $(\Theta^m, \Lambda^m, S^m)$. Without loss of generality we may assume that $\hat{p}^{\lambda}_{\tilde{s}_i} + L(\lambda) \geq P_i$ for all $\lambda \in \Lambda_i$, since the set of students for whom this is not true has measure 0. Moreover, since all students $\lambda \in \Lambda_i$ prefer school $\tilde{s}_{i+1}$ to their assigned school $\hat{\mu}(\lambda) = \tilde{s}_i$, without loss of generality we may also assume that $\hat{p}^{\lambda}_{\tilde{s}_{i+1}} + L(\lambda) < P_{i+1}$ for all $\lambda \in \Lambda_i$, since the set of students for whom this is not true has measure 0. This means that for all $\lambda_i \in \Lambda_i$ and $\lambda_{i+1} \in \Lambda_{i+1}$ it holds that $\hat{p}^{\lambda_i}_{\tilde{s}_{i+1}} + L(\lambda_i) < P_{i+1} \leq \hat{p}^{\lambda_{i+1}}_{\tilde{s}_{i+1}} + L(\lambda_{i+1})$, and so $\hat{p}^{\lambda_i}_{\tilde{s}_{i+1}} \leq \hat{p}^{\lambda_{i+1}}_{\tilde{s}_{i+1}}$.

Suppose for the sake of contradiction that the cycle $(\Theta^m, \Lambda^m, S^m)$ respects second-round priorities. Then for each $\lambda_i \in \Lambda_i$ and $\lambda_{i+1} \in \Lambda_{i+1}$ it holds that $\hat{p}^{\lambda_i}_{\tilde{s}_{i+1}} \geq \hat{p}^{\lambda_{i+1}}_{\tilde{s}_{i+1}}$, and so $L(\lambda_i) > L(\lambda_{i+1})$. But since this holds for all i we obtain a cycle of lottery numbers $L(\lambda_1) > L(\lambda_2) > \cdots > L(\lambda_m) > L(\lambda_1)$ for all $\lambda_i \in \Lambda_i$, which provides the necessary contradiction.

C.2 Proof of Theorem 1

We first prove Theorem 1 in the case where all schools have one priority group. We then show that if the order condition holds, all PLDA mechanisms assign the same number of seats at a given school s_i to students of a given priority class π. Hence, by restricting to the set of students with priority class π, we can reduce the general problem to the case where all schools have one priority group. This shows that all PLDA mechanisms produce type-equivalent assignments.

Lemma 1. *Assume that each school has a single priority group, $\mathbf{p} = \mathbf{1}$. If the order condition holds, all PLDA mechanisms produce type-equivalent assignments.*

Proof. Let P be a permutation.

Assume that the order condition holds. By Theorem 4, we may assume that the global order condition holds. Hence the schools in $S \cup \{s_{N+1}\}$ can be indexed so that $C_i \geq C_{i+1}$ and $\hat{C}^P_i \geq \hat{C}^P_{i+1}$ for all permutations P (simultaneously).

We first present the relevant notation that will be used in this proof. We are interested in sets of schools of the form $X_i = \{s_i, \ldots, s_{N+1}\}$. Let

$$\beta_{i,j} = \eta(\{\lambda \in \Lambda : s_i \text{ is the most desirable school in } X_j \text{ with respect to } \hat{\succ}^{\lambda}\})$$

be the measure of the students who, when their set of affordable schools is X_j, will choose s_i (when following their second-round preferences). Note that $\beta_{i,j} = 0$ for all $j > i$.

Let $E^\lambda(\mathbf{C})$ and $\hat{E}^\lambda_P(\hat{\mathbf{C}}^P)$ be the sets of schools affordable for type λ in the first and second round, respectively, when running PLDA with lottery P. Note that for each student $\lambda \in \Lambda$, there exists some i such that $E^\lambda(\mathbf{C}) = X_i$, and since the order condition is satisfied, there exists some $j \leq i$ such that $\hat{E}^\lambda_P(\hat{\mathbf{C}}^P) = X_j$. The fact that $\hat{E}^\lambda_P(\hat{\mathbf{C}}^P) = X_j$ for some j is a result of the order condition: students' amended second-round scores guarantee that $E^\lambda(\mathbf{C}) \subseteq \hat{E}^\lambda_P(\hat{\mathbf{C}}^P)$ (every school affordable in the first round is guaranteed in the second) and hence that $j \leq i$. Let $\gamma_i^P = \eta(\{\lambda \in \Lambda : \hat{E}^\lambda_P(\hat{\mathbf{C}}^P) = X_i\})$ be the fraction of students whose affordable set in the second round of PLDA with permutation P is[30] X_i. We note that by definition of PLDA, $\eta(\{\lambda \in \Lambda : \theta^\lambda = \theta, \hat{E}^\lambda_P(\hat{\mathbf{C}}^P) = X_i\}) = \zeta(\{\theta\})\gamma_i^P$; that is, the students whose affordable sets are X_i "break proportionally" into types. For a school i, this means that the measure of students assigned to s_i is therefore $\sum_{j \leq i} \beta_{i,j}\gamma_j^P$.

Let P' be another permutation, and define $\gamma_i^{P'}$ similarly. We will prove by induction that there exist PLDA(P') cutoffs $\hat{\mathbf{C}}^{P'}$ such that $\gamma_i^{P'} = \gamma_i^P$ for all $s_i \in S \cup \{s_{N+1}\}$. Note that by the proportional breaking into types of γ_i^P and $\gamma_i^{P'}$, this will imply type-equivalence.

Assume that the PLDA(P') cutoffs $\hat{\mathbf{C}}^{P'}$ are chosen such that $\gamma_j^{P'} = \gamma_j^P$ for all $j < i$, and i is maximal such that this is true. Then we have that $\sum_{j \leq i-1} \beta_{i,j}\gamma_j^P = \sum_{j \leq i-1} \beta_{i,j}\gamma_j^{P'}$. Assume w.l.o.g. that $\gamma_i^P > \gamma_i^{P'}$. It follows that $q_i \geq \sum_{j \leq i} \beta_{i,j}\gamma_j^P \geq \sum_{j \leq i} \beta_{i,j}\gamma_j^{P'}$, where the first inequality follows since s_i cannot be filled beyond capacity. If the second inequality is strict, then under P', s_i is not full, and therefore $\hat{C}_i^{P'} = 0$. However, this means that all students can afford s_i under P', and therefore $\gamma_i^{P'} = 1 - \sum_{j<i} \gamma_j^{P'} = 1 - \sum_{j<i} \gamma_j^P \geq \gamma_i^P$, a contradiction. If the second inequality is an equality, then $\beta_{i,i} = 0$ and no students demand school i under the given affordable set structure. It follows that we can define the cutoff $\hat{C}_i^{P'}$ such that $\gamma_i^{P'} = \gamma_i^P$. This provides the required contradiction, completing the proof. □

Now consider when schools have possibly more than one priority group. We show that if the order condition holds, then all PLDA mechanisms assign the same measure of students of a given priority type to a given school. It is not at all obvious that such a result should hold, since priority types and student preferences may be correlated, and the relative proportions of students of each priority type assigned to each school can vary widely. Nonetheless, the order condition (specifically, the equivalent global order condition) imposes enough structure so that any given priority type is

[30]Note that $\eta(\Lambda) = 1$, as η is a probability distribution over Λ.

treated symmetrically across different PLDA mechanisms.

Theorem 6. *If the order condition holds, then for all priority classes π and schools s_i all PLDA mechanisms assign the same measure of students of priority class π to school s_i.*

Proof. Fix a permutation P. By Theorem 4, we may assume that the global order condition holds.

We show that PLDA(P) assigns the same measure of students of each priority type to each school s_i as RLDA. The idea will be to define cutoffs on priority-type-specific economies, and show that these cutoffs are the same as the PLDA cutoffs. However, since cutoffs are not necessarily unique in the two-round setting, care needs to be taken to make sure that the individual choices for priority-type-specific cutoffs are consistent across priority types.

The proof runs as follows. We first define an economy $\mathcal{E}_\pi$ for each priority class π that gives only as many seats as are assigned to students of priority class π under RLDA. We then invoke the global order condition and Theorems 4 and 1 to show that all PLDA mechanisms are type-equivalent on each $\mathcal{E}_\pi$. We also use the global order condition to argue that it is sufficient to consider affordable sets, and also to select "minimal" cutoffs. Then we construct cutoffs $\mathcal{C}^P_{\pi,i}$ using the economies $\mathcal{E}_\pi$ and show that they are (almost) independent of priority type. Finally, we show that this means that $\mathcal{C}^P_{\pi,i}$ also define PLDA cutoffs for the large economy $\mathcal{E}$ and conclude that PLDA(P) assigns the same measure of students of each priority type to each school s_i as RLDA.

(1) Defining little economies $\mathcal{E}_\pi$ for each priority type.

Fix a priority class π. Let q_π be a restricted capacity vector, where $q_{\pi,i}$ is the measure of students of priority class π assigned to school s_i under RLDA. Let Λ_π be the set of students λ such that $p^\lambda = \pi$, and let η_π be the restriction of the distribution η to Λ_π. Let $\mathcal{E}_\pi$ denote the primitives $(S, q_\pi, \Lambda_\pi, \eta_\pi)$. Recall that $\hat{\mathbf{C}^{\mathbf{R}}}$ are the second-round cutoffs for RLDA on $\mathcal{E}$. It follows from the definition of $\mathcal{E}_\pi$ that $\hat{\mathbf{C}}^R_\pi$ are also the second-round cutoffs for RLDA on $\mathcal{E}_\pi$.

Let $\tilde{\mathbf{C}}^P_\pi$ be the second-round cutoffs of PLDA(P) on $\mathcal{E}_\pi$. We show that the cutoffs $\tilde{\mathbf{C}}^P_\pi$ defined for the little economy are the same as the consistent second-round cutoffs $\hat{\mathbf{C}}^P_\pi$ for PLDA with permutation P for the large economy $\mathcal{E}$, that is, $\tilde{\mathbf{C}}^P_\pi = \hat{\mathbf{C}}^P_\pi$.

(2) Implications of the global order condition.

We have assumed that the global order condition holds.

This has a number of implications for PLDA mechanisms run on the little economies $\mathcal{E}_\pi$. For all p, the local order condition holds for RLDA on $\mathcal{E}_\pi$. Hence, by Theorem 4, the little economies

$\mathcal{E}_\pi$ each satisfy the order condition. Moreover, by Theorem 1, all PLDA mechanisms produce type-equivalent assignments when run on $\mathcal{E}_\pi$. Finally, if we can show that for every permutation P, PLDA(P) assigns the same measure of students of each priority type π to each school s_i (namely $(q_\pi)_i$) as RLDA, then $\mathcal{E}$ satisfies the global order condition if and only if for all p the little economy $\mathcal{E}_\pi$ satisfies the global order condition.

The global order condition also allows us to determine aggregate student demand from the proportions of students who have each school in their affordable set. In general, if affordable sets break proportionally across types, and if for each subset of schools $S' \subseteq S$ we know the proportion of students whose affordable set is S', then we can determine aggregate student demand. The global order condition implies that for any pair of permutations P, P', the affordable sets from both rounds are nested in the same order under both permutations. In other words, for each priority class π there exists a permutation σ_π such that the affordable set of any student in any round of any PLDA mechanism is of the form $\{s_{\sigma_\pi(i)}, s_{\sigma_\pi(i+1)}, \dots, s_{\sigma_\pi(N)}, s_{N+1}\}$. Hence when the global order condition holds, to determine the proportion of students whose affordable set is S', it is sufficient to know the proportion of students who have each school in their affordable set.

Another more subtle implication of the global order condition is the following. In the second round of PLDA, for each permutation P and school s_i there will generically be an interval that $\hat{C}_i^P$ can lie in and still be market-clearing. The intuition is that there will be large empty intervals corresponding to students who had school s_i in their first-round affordable set, and whose second-round lottery changed accordingly. When the global order condition holds, we can without loss of generality assume that as many as possible of the cutoffs for a given priority type are 0 or 1, and the global order condition will still hold.

Formally, for cutoffs C we can equivalently define priority-type-specific cutoffs $C_{\pi,i} = (\lfloor C_i - \pi_i \rfloor)^+$. Note the cutoffs C_π are consistent across priority types, namely: (1) Cutoffs match for two priority types with the same priority group at a school, $\pi_i = \pi'_i \Rightarrow C_{\pi,i} = C_{\pi',i}$ and $\hat{C}_{\pi,i} = \hat{C}_{\pi',i}$; and (2) There is at most one marginal priority group at each school, $C_{\pi,i}, C_{\pi',i} \in (0,1) \Rightarrow \pi_i = \pi'_i$. Moreover, if cutoffs C_π are consistent across priority types, then there exist cutoffs C from which they arise.

Suppose that we set as many as possible of the priority-type-specific cutoffs $\hat{C}_\pi^P$ to be extremal; i.e., we let $\hat{C}_{\pi,i}^P$ be 1 if no students have s_i in their affordable set, and let $\hat{C}_{\pi,i}^P$ be minimal otherwise. We show that under this new definition, $C_\pi, \hat{C}_\pi^P$ satisfies the local order condition consistently with all other PLDAs.

Specifically, let

$$f^P_{\pi,i}(x) = |\{l : l \geq C_{\pi,i} \text{ or } P(l) \geq x\}|$$

be the proportion of students of priority class π who have school s_i in their affordable set if the first- and second-round cutoffs are $C_{\pi,i}$ and x respectively. Notice that f is decreasing in x. Define cutoffs $\tilde{C}^P_\pi$ as follows. If $f^P_{\pi,i}\left(\hat{C}^P_{\pi,i}\right) = 0$ we set $\tilde{C}^P_{\pi,i} = 1$, and otherwise we let $\tilde{C}^P_{\pi,i}$ be the minimal cutoff satisfying $f^P_{\pi,i}\left(\tilde{C}^P_{\pi,i}\right) = f^P_{\pi,i}\left(\hat{C}^P_{\pi,i}\right)$.

Since $\mathcal{E}$ satisfies the global order condition, for all π there exists an ordering σ_π such that $C_{\sigma_\pi(1)} \geq C_{\sigma_\pi(2)} \geq \cdots \geq C_{\sigma_\pi(N)}$ and $\hat{C}^{P'}_{\sigma_\pi(1)} \geq \hat{C}^{P'}_{\sigma_\pi(2)} \geq \cdots \geq \hat{C}^{P'}_{\sigma_\pi(N)}$ for all permutations P'. We show that the global order condition implies that the newly defined cutoffs $\hat{C}^P$ satisfy $\tilde{C}^P_{\pi,\sigma_\pi(1)} \geq \tilde{C}^P_{\pi,\sigma_\pi(2)} \geq \cdots \geq \tilde{C}^P_{\pi,\sigma_\pi(n)}$. This is because the global order condition implies that f^P_π is increasing in i; i.e., for each π, $i < j$, and x it holds that $f^P_{\pi,\sigma_\pi(i)}(x) \leq f^P_{\pi,\sigma_\pi(j)}(x)$. Hence for all $j > i$,

$$\begin{aligned}
f^P_{\pi,\sigma_\pi(j)}\left(\tilde{C}^P_{\pi,\sigma_\pi(j)}\right) &= f^P_{\pi,\sigma_\pi(j)}\left(\hat{C}^P_{\pi,\sigma_\pi(j)}\right) \\
&\geq f^P_{\pi,\sigma_\pi(j)}\left(\hat{C}^P_{\pi,\sigma_\pi(i)}\right) \text{ (since } f \text{ is decreasing)} \\
&\geq f^P_{\pi,\sigma_\pi(i)}\left(\hat{C}^P_{\pi,\sigma_\pi(i)}\right) \text{ (since } f \text{ is increasing in } i) \\
&= f^P_{\pi,\sigma_\pi(i)}\left(\tilde{C}^P_{\pi,\sigma_\pi(i)}\right)
\end{aligned}$$

and so since we set $\tilde{C}^P_{\pi,\sigma_\pi(j)}$ to be minimal and $f^P_{\pi,\sigma_\pi(j)}(\cdot)$ is decreasing it follows that $\tilde{C}^P_{\pi,\sigma_\pi(j)} \leq \tilde{C}^P_{\pi,\sigma_\pi(i)}$.

(3) Cutoffs $\tilde{C}^P_{\pi,i}$ are (almost) independent of priority type.

We now show that $\tilde{C}^P_{\pi,i}$ depends on π only via π_i, and for all $j \neq i$ does not depend on π_j. Since $\mathcal{E}_\pi$ satisfies the order condition, all PLDA mechanisms on $\mathcal{E}_\pi$ are type-equivalent, and the proportion of students who have each school in their affordable set is the same across all PLDA mechanisms. Hence for all permutations P, priority classes π, and schools i it holds that $f^P_{\pi,i}\left(\tilde{C}^P_{\pi,i}\right) = f^P_{\pi,i}\left(\hat{C}^P_{\pi,i}\right) = f^R_{\pi,i}\left(\hat{C}^R_{\pi,i}\right)$. This means that $\tilde{C}^P_{\pi,i}$ satisfies the following equation in terms of $\hat{C}^R_{\pi,i}$, $C_{\pi,i}$ and P:

$$f^P_{\pi,i}\left(\tilde{C}^P_{\pi,i}\right) = f^R_{\pi,i}\left(\hat{C}^R_{\pi,i}\right) = 2 - \hat{C}^R_{\pi,i} - C_{\pi,i}. \tag{4}$$

(We note that an application of the intermediate value theorem shows that this equation always has a solution in $[0,1]$, since $f^P_{\pi,i}(0) = 1 - C_{\pi,i}$, $f_{\pi,i}(1) = 1$, $f_{\pi,i}$ is continuous and decreasing on $[0,1]$, and we are in the case where $1 - C_{\pi,i} \leq \hat{C}^R_{\pi,i} \leq 1$. Hence $\tilde{C}^P_{\pi,i}$ is defined by $f^P_{\pi,i}$ and $f^R_{\pi,i}$.) In other words, the value of $\tilde{C}^P_{\pi,i}$ is defined by $f^P_{\pi,i}(\cdot), f^R_{\pi,i}(\cdot)$, and $\hat{C}^R_{\pi,i}$, which in turn are defined by $C_{\pi,i}$ and the permutations P and R. Since $C_{\pi,i}$ depends on π only through π_i, it follows that $\tilde{C}^P_{\pi,i}$ depends on π only through π_i. In other words the $\tilde{C}^P_{\pi,i}$ define cutoffs $\tilde{C}^P_i$ that are independent of priority type.

(4) $\tilde{C}^P_i$ are the PLDA cutoffs.

Finally, we remark that $\tilde{C}^P_i$ are market-clearing cutoffs. This is because we have shown that for each priority class π, the number of students assigned to each school s_i is the same under RLDA and under the demand induced by the cutoffs $\tilde{C}^P_i$, and we know that the RLDA cutoffs are market-clearing for $\mathcal{E}$.

Hence $\tilde{C}^P_i$ give the assignments for PLDA on $\mathcal{E}$, and since $\tilde{C}^P_i$ was defined individually for each priority class π on $\mathcal{E}_\pi$, it follows that PLDA(P) assigns the same measure of students of each priority type to each school s_i as RLDA. □

We are now ready to prove Theorem 1

Proof of Theorem 1. Fix a priority class π. By Theorem 6, for every school s_i, all PLDA mechanisms assign the same measure $q_{\pi,i}$ of students of priority class π to school s_i.

Consider the subproblem with primitives $\mathcal{E}_\pi = (S, q_\pi, \Lambda_\pi, \eta_\pi)$. By Lemma 1, for all $\theta \in \Theta$ and s_i,

$$\eta_\pi(\{\lambda \in \Lambda_\pi : \theta^\lambda = \theta, \hat{\mu}_P(\lambda) = s_i\}) = \eta_\pi(\{\lambda \in \Lambda_\pi : \theta^\lambda = \theta, \hat{\mu}_{P'}(\lambda) = s_i\}).$$

Since η_π is the restriction of η to λ_π, it follows that all PLDA mechanisms are type-equivalent. □

C.3 Proof of Theorem 4

Proof of Theorem 4. Suppose that the order condition holds. In what follows, we will fix a permutation P and show that the PLDA mechanism with permutation P satisfies the local order condition and is type-equivalent to the reverse lottery RLDA mechanism. As this holds for every P, it follows that the global order condition holds.

(1) Every school has a single priority group.

We first consider the case where $n_i = 1$ for all i; that is, every school has a single priority group. Recall that the schools are indexed according to the first-round overdemand ordering, so that $C_1 \geq C_2 \geq \cdots \geq C_N \geq C_{N+1}$. Since the local order condition holds for RLDA, let us assume that they are also indexed according to the second-round overdemand ordering under RLDA, so that $\hat{C}_1^R \geq \hat{C}_2^R \geq \cdots \geq \hat{C}_N^R \geq \hat{C}_{N+1}^R$.

The idea will be to construct a set of cutoffs $\tilde{\mathbf{C}}^P$ directly from the permutation P and the cutoffs $\hat{\mathbf{C}}^R$, show that the cutoffs are in the correct order $\tilde{C}_1^P \geq \tilde{C}_2^P \geq \cdots \geq \tilde{C}_N^P \geq \tilde{C}_{N+1}^P$, and show that the cutoffs $\tilde{\mathbf{C}}^P$ and resulting assignment are market-clearing when school preferences are given by the amended scoring function with permutation P.

(1a) Definitions.

As in the proof of Theorem 1, let $\beta_{i,j} = \eta(\{\lambda \in \Lambda : \operatorname{argmax}_{\succ^\lambda} X_j = s_i\})$ be the measure of students who, when their set of affordable schools is X_j, will choose s_i. Let $E^\lambda(\mathbf{C})$ be the set of schools affordable for type λ in the first round under PLDA with any permutation, let $\hat{E}^\lambda(\hat{\mathbf{C}}^R)$ be the set of schools affordable for type λ in the second round under RLDA, and let $\hat{E}^\lambda(\hat{\mathbf{C}}^P)$ be the set of schools affordable for type λ in the second round under PLDA with permutation P.

Let $\gamma_i^R = \eta(\{\lambda \in \Lambda : \hat{E}^\lambda(\hat{\mathbf{C}}^R) = X_i\})$ be the fraction of students whose affordable set in the second round of RLDA is X_i, and let $\gamma_i^P = \eta(\{\lambda \in \Lambda : \hat{E}^\lambda(\hat{\mathbf{C}}^P) = X_i\})$ be the fraction of students whose affordable set in the second round of PLDA with permutation P is X_i.

Let $\hat{n}$ be the smallest index such that $s_{\hat{n}}$ does not reach capacity when it is not offered to all the students. In other words, $\hat{n}$ is the smallest index such that every student has school $s_{\hat{n}}$ in her affordable set under RLDA, i.e., $s_{\hat{n}} \in \hat{E}^\lambda(\hat{\mathbf{C}}^R)$. Since the local order condition holds for RLDA, we may equivalently express $\hat{n}$ in terms of cutoffs as the smallest index such that $(1-C_{\hat{n}})+(1-\hat{C}_{\hat{n}}^R) \geq 1$. Such an $\hat{n}$ always exists, since every student has the outside option s_{N+1} in her total affordable set.

(1b) Defining cutoffs for PLDA.

Let us define cutoffs $\tilde{\mathbf{C}}^P$ as follows. For $i \geq \hat{n}$ let $\tilde{C}_i^P = 0$. For each permutation P, define a function

$$f_i^P(x) = |\{l : l \geq C_i \text{ or } P(l) \geq x\}|$$

representing the proportion of students who have s_i in their (second-round) affordable set with first- and second-round cutoffs C_i, x under the amended scoring function with permutation P. Since P is measure-preserving, $f_i^P(x)$ is continuous and monotonically decreasing in x.

For $i < \hat{n}$, we inductively define $\tilde{C}_i^P$ to be the largest real smaller than $\tilde{C}_{i-1}^P$ satisfying

$$f_i^P(\tilde{C}_i^P) = f_i^R\left(\hat{C}_i^R\right) \tag{5}$$

(where we define $\tilde{C}_0^P = 1$). Now $f_i^P(0) = 1 \geq f_i^R\left(\hat{C}_i^R\right)$, and

$$\begin{aligned} f_i^P\left(\tilde{C}_{i-1}^P\right) &= f_{i-1}^P\left(\tilde{C}_{i-1}^P\right) + |\{l \mid l \in [C_i, C_{i-1}) \text{ and } P(l) \geq \tilde{C}_{i-1}^P\}| \\ &\leq f_{i-1}^R\left(\hat{C}_{i-1}^R\right) + (C_{i-1} - C_i) \\ &= (1 - C_i) + (1 - \hat{C}_{i-1}^R) \\ &\leq f_i^R\left(\hat{C}_i^R\right) = f_i^P\left(\tilde{C}_i^P\right) \end{aligned}$$

where in the first equality we are using that $C_{i-1} \geq C_i$, the first inequality follows from the definition of $\tilde{C}_{i=1}^P$, and the last inequality holds since $\hat{C}_{i-1}^R \geq \hat{C}_i^R$.

It follows from the intermediate value theorem that the cutoffs $\tilde{\mathbf{C}}^P$ are well defined and satisfy $\tilde{C}_1^P \geq \tilde{C}_2^P \geq \cdots \geq \tilde{C}_N^P \geq \tilde{C}_{N+1}^P$.

(1c) The constructed cutoffs clear the market.

We show that the cutoffs $\tilde{\mathbf{C}}^P$ and resulting assignment (from letting students choose their favorite school out of those for which they meet the cutoff) are market-clearing when the second-round scores are given by $\hat{r}_i^\lambda = P(L(\lambda)) + n_i \mathbb{1}_{\{L(\lambda) \geq C_i\}} + p_i^\lambda \mathbb{1}_{\{L(\lambda) \geq C_i\}}$. We call the mechanism with this second-round assignment M^P.

The idea is that since the cutoffs $\tilde{C}_i^P$ are decreasing in the same order as C_i and $\hat{C}_i^R$, the (second-round) affordable sets are nested in the same order under both sets of second-round cutoffs. It follows that aggregate student demand is uniquely specified by the proportion of students with each school in their affordable set, and we have defined these to be equal, $f_i^P\left(\tilde{C}_i^P\right) = f_i^R\left(\hat{C}_i^R\right)$. It follows that $\tilde{\mathbf{C}}^P$ are market-clearing and give the PLDA(P) cutoffs, and so PLDA(P) satisfies the local order condition (with the indices indexed in the same order as with RLDA). We make the affordable set argument explicit below.

Consider the proportion of lottery numbers giving a second-round affordable set X_i. Since $\hat{C}_1^R \geq \hat{C}_2^R \geq \cdots \geq \hat{C}_N^R$, under RLDA this is given by

$$\gamma_i^R = f_{i+1}^R\left(\hat{C}_{i+1}^R\right) - f_i^R\left(\hat{C}_i^R\right),$$

if $i < \hat{n}$ and by 0 if $i > \hat{n}$, where we define $f_0^P(x) = 1$ for all P and x. Similarly, since $\tilde{C}_1^P \geq \tilde{C}_2^P \geq \cdots \geq \tilde{C}_N^P$, under M^P this is given by

$$f_{i+1}^P\left(\tilde{C}_{i+1}^P\right) - f_i^P\left(\tilde{C}_i^P\right)$$

if $i < \hat{n}$, which is precisely γ_i^R, and by 0 if $i > \hat{n}$.

Hence, for all $i < \hat{n}$, the measure of students assigned to school s_i under both RLDA and M^P is $\sum_{j \leq i} \beta_{i,j}\gamma_j^R = q_i$, and for all $i \geq \hat{n}$, the measure of students assigned to school s_i is $\sum_{j \leq \hat{n}} \beta_{i,j}\gamma_j^R < q_i$. It follows that the cutoffs $\tilde{\mathbf{C}}^P$ are market-clearing when the second-round scores are given by $\hat{r}_i^\lambda = P(L(\lambda)) + n_i \mathbb{1}_{\{L(\lambda) \geq C_i\}} + p_i^\lambda \mathbb{1}_{\{L(\lambda) \geq C_i\}}$, and so $PLDA(\mathrm{P}) = M^P$ satisfies the local order condition.

(2) Some school has more than one priority group.

Now consider when schools have possibly more than one priority group. We show that if RLDA satisfies the local order condition, then PLDA with permutation P assigns the same number of students of each priority type to each school s_i as RLDA, and within each priority type assigns the same number of students of each preference type to each school as RLDA. We do this by first assuming that PLDA with permutation P assigns the same number of students of each priority type to each school s_i as RLDA, and showing that this gives consistent cutoffs.

We note that this proof uses very similar arguments to the proof of Theorem 6.

(2a) Defining little economies $\mathcal{E}_\pi$ for each priority type.

Fix a priority class π. Let q_π be a restricted capacity vector, where $q_{\pi,i}$ is the measure of students of priority class π assigned to school s_i under RLDA. Let Λ_π be the set of students λ such that $p^\lambda = \pi$, and let η_π be the restriction of the distribution η to Λ_π. Let $\mathcal{E}_\pi$ denote the primitives $(S, q_\pi, \Lambda_\pi, \eta_\pi)$.

Let $\tilde{\mathbf{C}}_\pi^P$ be the second-round cutoffs of PLDA(P) on $\mathcal{E}_\pi$. By definition, $\hat{\mathbf{C}}_\pi^R$ are the second-round cutoffs of RLDA on $\mathcal{E}_\pi$. We show that the cutoffs $\tilde{\mathbf{C}}_\pi^P$ defined for the little economy are the same as the consistent second-round cutoffs $\hat{\mathbf{C}}_\pi^P$ for PLDA(P) run on the large economy $\mathcal{E}$, that is, $\tilde{\mathbf{C}}_\pi^P = \hat{\mathbf{C}}_\pi^P$.

(2b) Implications of RLDA satisfying the local order condition.

Since RLDA satisfies the local order condition for $\mathcal{E}$, RLDA also satisfies the local order condition for $\mathcal{E}_\pi$ for all π. It follows from (1) that the global order condition holds on each of the little economies $\mathcal{E}_\pi$. Hence by Theorem 1 all PLDA mechanisms produce type-equivalent assignments when run on $\mathcal{E}_\pi$. Moreover, as in the proof of Theorem 6, the global order condition on $\mathcal{E}_\pi$ also allows us to determine aggregate student demand in $\mathcal{E}_\pi$ from the proportions of students who have each school in their affordable set.

Finally, as in the proof of Theorem 6, we may assume that for each π and school s_i the cutoff $\tilde{C}^P_{\pi,i}$ is the minimal real satisfying

$$f^P_{\pi,i}\left(\tilde{C}^P_{\pi,i}\right) = f^R_{\pi,i}\left(\hat{C}^R_{\pi,i}\right)$$

where for each permutation P,

$$f^P_{\pi,i}(x) = |\{l : l \geq C_{\pi,i} \text{ or } P(l) \geq x\}|$$

is the proportion of students of priority class π who have school s_i in their affordable set if the first- and second-round cutoffs are $C_{\pi,i}$ and x respectively.

It follows that $\tilde{C}^P_{\pi,i}$ depends on π only via π_i, and does not depend on π_j for all $j \neq i$. This is because $\tilde{C}^P_{\pi,i}$ is defined by $f^P_{\pi,i}(\cdot)$, $f^R_{\pi,i}(\cdot)$, and $\hat{C}^R_{\pi,i}$, which are in turn defined by $C_{\pi,i}$ and the permutations P and R. Moreover, $C_{\pi,i}$ depends on π only through π_i. Hence, if π, π' are two priority vectors such that $\pi_i = \pi'_i$, then $\tilde{C}^P_{\pi,i} = \tilde{C}^P_{\pi',i}$, and so the $\tilde{C}^P_{\pi,i}$ are consistent across priority types and define cutoffs $\tilde{C}^P_i$ that are independent of priority type.

(3) $\tilde{C}^P_i$ are the PLDA cutoffs.

Finally, we show that $\tilde{C}^P_i$ are market-clearing cutoffs. By (1), for each priority class π, the number of students assigned to each school s_i is the same under RLDA as under the demand induced by the cutoffs $\tilde{C}^P_i$, and we know that the RLDA cutoffs are market-clearing for $\mathcal{E}$.

Hence $\tilde{C}^P_i$ give the assignments for PLDA on $\mathcal{E}$, and since $\tilde{C}^P_i$ was defined individually for each priority class π for $\mathcal{E}_\pi$ it follows that PLDA(P) assigns the same measure of students of each priority type to each school s_i as RLDA.

□

C.4 Proof of Theorem 3

Proof of Theorem 3. We first note that with a single priority class, the first round corresponds to the random serial dictatorship (RSD) mechanism of Abdulkadiroglu and Sönmez (1998), where the (random) order of students is the single order of tie-breaking. Hence instead of referring to the first-round mechanism as DA-STB, we will sometimes refer to it as RSD.

Recall the cutoff characterization of the set of stable matchings for given student preferences and responsive school preferences (encoded by student scores $r_i^\lambda = p_i^\lambda + L(\lambda)$) Azevedo and Leshno (2016). Namely, if $\mathbf{C} \in \mathbb{R}_+^{\mathbf{N}}$ is a vector of cutoffs, let the assignment μ defined by $\mathbf{C}$ be given by assigning each student of type λ to her favorite school among those where her score weakly exceeds the cutoff, $\mu(\lambda) = \max_{\succ^\lambda}(\{s_i \in S : r_i^\lambda \geq C_i\} \cup \{s_{N+1}\})$. The cutoffs $\mathbf{C}$ are market-clearing if under the assignment μ defined by $\mathbf{C}$, every school with a positive cutoff is exactly at capacity, $\eta(\mu(s_i)) \leq q_i$ for all $s_i \in S$, with equality if $C_i > 0$. The set of all stable matchings is precisely given by the set of assignments defined by market-clearing vectors (Azevedo and Leshno, 2016).

Under PLDA(P), a student of type λ has a second-round score $\hat{r}_i^\lambda = P(L(\lambda)) + \mathbb{1}_{\{L(\lambda) \geq C_i\}}$ at school s_i for each school $s_i \in S \cup \{s_{N+1}\}$ (assuming that scores are modified to give guarantees to students who had a school in their first-round affordable set, instead of just students assigned to the school in the first round). In a slight abuse of notation, we will sometimes let $\hat{\mathbf{C}}^P$ refer to the second-round cutoffs from some fixed PLDA(P) (not necessarily corresponding to the student-optimal stable matching given by PLDA).

The proof that any PLDA satisfies the axioms essentially follows from Proposition 1. We note that averaging follows from the continuum model, which preserves the relative proportion of students with different reported types under random lotteries and permutations of random lotteries. Hence it suffices to show that any mechanism M satisfying the axioms is a PLDA.

We will show that the reassignment produced by M is type-equivalent to the reassignment produced by some PLDA. If we assume that, conditional on their reports, students' assignments under M are uncorrelated, we are able to explicitly construct a PLDA that provides the same joint distribution over assignments and reassignments as M. We provide a sketch of the proof before fleshing out the details.

Fix a distribution of student types ζ. Since the first round of our mechanism M is DA-STB and M is anonymous, this gives a distribution η of students that is the same (up to relabeling of students) at the end of the first round. For a fixed labeling of students, it also gives a distribution over first-round assignments μ and a distribution over second round assignments $\tilde{\mu}$.

We first invoke averaging to assume that all ensuing constructions of aggregate cutoffs and measures of students assigned to pairs of schools in the two rounds are deterministic. Specifically, since the first-round assignment μ is given by STB, and the mechanism satisfies the averaging axiom, we may assume that each pathwise realization of the mechanism gives type-equivalent (two-round) assignments. Hence, for the majority of the proof we perform our constructions of aggregate cutoffs and measures of students pathwise, and assume that any realization of the lottery numbers produces the same cutoffs and measures of students. (In particular, the quantities $\hat{C}_i, \rho_{i,j}, \gamma_{i,j}$ that we will later define will be the same across all realizations.)

Outline of Proof. We use constrained Pareto efficiency to construct a first-round overdemand ordering $s_1, s_2, \ldots, s_N, s_{N+1}$ and a permutation σ giving the second-round overdemand ordering $s_{\sigma(1)}, s_{\sigma(2)}, \ldots, s_{\sigma(N+1)}$, as in (Ashlagi and Shi, 2014), where school s comes before s' in an ordering for the first (second) round if there exists a non-zero measure of students who prefer school s to s' in the first (second) round but who are assigned to s' in the first (second) round. (In the case of the second-round ordering, we require that these students' second-round assignments s' not be the same as their first-round guarantees.) The existence of these orderings follows from the facts that the first-round mechanism, DA with a single priority class and uniform-at-random single tie-breaking, is Pareto efficient, that the two-round mechanism is constrained Pareto efficient. We let $X_i = \{s_i, s_{i+1}, \ldots, s_{N+1}\}$ denote the set of schools after s_i in the first-round overdemand ordering, and let $\tilde{X}_i = \{s_{\sigma(i)}, s_{\sigma(i+1)}, \ldots, s_{\sigma(N+1)}\}$ denote the set of schools after $s_{\sigma(i)}$ in the second-round overdemand ordering.

We next note that instead of assignments μ and $\hat{\mu}$, we can think of giving students first- and second-round affordable sets $E(\lambda), \hat{E}(\lambda)$ so that μ and $\hat{\mu}$ are given by letting each student choose her favorite school in her affordable set for that round. We use weak two-round strategy-proofness and anonymity to show that two students of different types face the same joint distribution over first- and second-round affordable sets. This allows us to construct the permutation P by constructing proportions $\gamma_{i,j}$ of students whose first-round affordable set was X_i and whose (second-round) affordable set was $\tilde{X}_j$. This is the most technical step in the proof, and so we separate it into several steps. The crux of the analysis is the fact that for any school s and set $S' \not\ni s$ of schools, two students with top choices S' who are assigned to a school they weakly prefer to s the first round have the same conditional probability of being assigned to a school in S' in the second round.[31]

[31] The formal statement also takes into account how demanded the schools they weakly prefer to s are, and is given

We term this the "prefix property" and prove it in Lemma 2.

Finally, we construct the lottery L and verify that if second-round scores are given by first prioritizing all guaranteed students over non-guaranteed students and subsequently breaking ties according to the permuted lottery $P \circ L$, then PLDA(P) gives every student the same pair of first- and second- round assignments as M.

Formal Proof. We now present the formal proof. Since we are assuming that the considered mechanism M is weakly strategy-proof, we assume that students report truthfully and so we consider preferences instead of reported preferences. We will explicitly specify when we are considering the possible outcomes from a single student misreporting.

(2a) Definitions

Let the schools be numbered $s_1, s_2, \ldots, s_N$ such that $C_i \geq C_{i+1}$ for all i. The intuition is that this is the order in which they reach capacity in the first round. We observe that all reassignments are index-decreasing. That is, for all s, s', if there exists a non-zero measure of students who are assigned to s in the first round and to s' in the second round, and $s' \neq s_{N+1}$, then $s = s_i$ and $s' = s_j$ for some $i \geq j$. This follows since the mechanism respects guarantees, student preferences are consistent, and the schools are indexed in order of increasing first-round affordability. Throughout this section we will denote the outside option s_{N+1} either by s_0 or $\emptyset$, to make it more evident that indices are decreasing.

Next, we define a permutation σ on the schools. We think of this as giving a second-round overdemand (or inverse affordability) ordering, where in the second round the schools fill in the order $s_{\sigma(1)}, s_{\sigma(2)}, \ldots, s_{\sigma(N)}$. We will eventually show that M gives the same outcome as a PLDA with cutoffs that are ordered $\hat{C}_{\sigma(1)} \geq \hat{C}_{\sigma(2)} \geq \cdots \geq \hat{C}_{\sigma(N)}$. We require that σ satisfies the following property. For all s, s', if there exists a non-zero measure of students with consistent preferences who have second-round preference reports $\succ$ such that $s \succ s'$, and who are not assigned to s' in the first round, but are assigned to s' in the second round, then $s = s_{\sigma(i)}$ and $s' = s_{\sigma(j)}$ for some $i < j$. We assume that σ is the unique permutation satisfying this property that is maximally order-preserving. That is, for all pairs of schools s_i, s_j for which no non-zero measure of students of the above type exists, $\sigma(i) < \sigma(j)$ iff $i < j$. We also define $\sigma(N+1) = N+1$. An ordering σ with the required properties exists since the mechanism is constrained Pareto efficient. In particular, if there is a cycle of schools $s_{i_1}, s_{i_2}, \ldots, s_{i_m}$ where for each j there is a set of students Λ_j with non-zero

in terms of student types who were assigned to s, and lottery numbers.

measure who prefer $s_{i_{j+1}}$ to their second-round assignment s_{i_j} and who are not assigned to s_{i_j} in the first round, then $\hat{p}_{s_{i_j}}^{\lambda_j} = p_{s_{i_j}}^{\lambda_j}$ for each $\lambda_j \in \Lambda_j$, and so there is a Pareto-improving cycle that respects second-round priorities.

Let S' be a set of schools, and let $\succ$ be a preference ordering over all schools. We say that S' is a *prefix* of $\succ$ if $s' \succ s$ for all $s' \in S', s \notin S'$. For a set of schools S', let $i(S') = \max\{j : s_j \in S'\}$ be the maximum index of a school in S'. We may think of $i(S')$ as the index of the most affordable school in S' in the first round.

For a student type $\theta = (\succ, \hat{\succ})$, an interval $I \subseteq [0,1]$, and a set of schools S', let $\rho^\theta(I, S')$ be the proportion of students with type θ who, under the mechanism M, have a first-round lottery in the interval I and are assigned to a school in S' in the second round. When $S' = \{s'\}$ we will sometimes write $\rho^\theta(I, s')$ instead of[32] $\rho^\theta(I, \{s'\})$. In this section, for brevity, when defining preferences $\succ$ we will sometimes write $\succ$: $[s_{i_1}, s_{i_2}, \ldots, s_{i_k}]$ instead of $s_{i_1} \succ s_{i_2} \succ \cdots \succ s_{i_k}$.

(2b) Constructing the permutation P.

We now construct the permutation P as follows. For all pairs of indices i, j, we define a scalar $\gamma_{i,j}$, which we will show can be thought of as the proportion of students (of any type) whose first-round affordable set is X_i and whose second-round affordable set is $\tilde{X}_j$.

Now, for all pairs of indices i, j such that $\sigma(j) < i$, we define student preferences $\theta_{i,j} = (\succ_{i,j}, \hat{\succ}_{i,j})$ such that

$$\succ_{i,j}: [s_{\sigma(j)}, s_{i-1}, s_i, s_{N+1}] \quad \text{and} \quad \hat{\succ}_{i,j} : [s_{\sigma(j)}, s_{N+1}],$$

with all other schools unacceptable. (We remark that in the case where $\sigma(j) = i - 1$, the first two schools in this preference ordering coincide.) We note that the full-support assumption implies that there is a positive measure of such students. Let $\rho_{i,j}$ be the proportion of students of type $\theta_{i,j}$ whose first-round assignment is s_i and whose second-round assignment is school $s_{\sigma(j)}$. Intuitively, $\rho_{i,j}$ is the proportion of students who can deduce that their lottery number is in the interval $[C_i, C_{i-1}]$, and whose second-round affordable set contains $\tilde{X}_j$.

For a fixed index i, we define $\gamma_{i,j}$ for $j = 1, 2, \ldots, N$ to be the unique solutions to the following n equations:

$$\gamma_{i,j} = 0 \quad \text{for all } j \text{ such that } \sigma(j) \geq i$$

$$\gamma_{i,1} + \cdots + \gamma_{i,j} = \rho_{i,j} \quad \text{for all } j \text{ such that } \sigma(j) < i.$$

[32]Here we are assuming that this proportion is the same for every realization of the first round of M. This requires non-atomicity and anonymity.

Note that by this definition it holds that $\gamma_{1,j} = 0$ for all j. We may intuitively think of $\gamma_{i,j}$ as the proportion of students of type $\theta_{i,j}$ whose first-round lottery is in $[C_i, C_{i-1}]$ and whose second-round affordable set contains $s_{\sigma(j)}$ but not $s_{\sigma(j-1)}$. (This is not quite the case, as we let $\gamma_{i,j} = 0$ for all j such that $\sigma(j) \geq i$. More precisely, if $\sigma(j) < i$ then $\gamma_{i,j}$ is the proportion of students of type $\theta_{i,j}$ whose first-round lottery is in $[C_i, C_{i-1}]$ and whose second-round affordable set contains $s_{\sigma(j)}$, but not $s_{\sigma(j')}$, where $j' = \max\{j'' : \sigma(j'') < i\}$.) Note that if $\sigma(j) \geq i$ then school $s_{\sigma(j)}$ will be in the first-round affordable set for all students whose first-round lottery is in $[C_i, C_{i-1}]$, and we define $\gamma_{i,j} = 0$ and keep track of these students separately.

We also define $\gamma_{i,N+1}$ to be

$$\gamma_{i,N+1} = C_{i-1} - C_i - \sum_{j=1}^{n} \gamma_{i,j}.$$

Since transfers are index-decreasing, we may intuitively think of $\gamma_{i,N+1}$ as the proportion of students of type $\theta_{i,j}$ assigned to school s_i in the first round whose only available school in the second round comes from their first-round guarantee.

We define the lottery P from $\gamma_{i,j}$ as follows. We break the interval $[0,1]$ into $(N+1)^2$ intervals, $\tilde{I}_{i,j}$, where the interval $\tilde{I}_{i,j}$ has length $\gamma_{i,j}$, and the intervals are ordered in decreasing order of the first index[33] i,

$$\tilde{I}_{N+1,N+1}, \tilde{I}_{N+1,N}, \ldots, \tilde{I}_{1,2}, \tilde{I}_{1,1}.$$

The permutation P maps the intervals back into $[0,1]$ in decreasing order of the second index[34] j,

$$P(\tilde{I}_{N+1,N+1}), P(\tilde{I}_{N,N+1}), \ldots P(\tilde{I}_{2,1}), P(\tilde{I}_{1,1}).$$

In Figure 5, we show an example with two schools.

We note that $\sum_{j=1}^{N+1} \gamma_{i,j} = C_{i-1} - C_i$, which is the proportion of students whose first-round affordable set is X_i. We may interpret $\gamma_{i,j}$ to be the proportion of students who can deduce that their lottery number is in the interval $[C_i, C_{i-1}]$, and whose second-round affordable set is $\tilde{X}_j$, and so $\sum_{i=1}^{N+1} \gamma_{i,j}$ is the proportion of students whose second-round affordable set is $\tilde{X}_j$. We remark that there may be multiple values of i, j for which $\gamma_{i,j} = 0$ (i.e. there are no students whose first-round affordable set is X_i and second-round affordable set is X_j), but that this does not affect our ability

[33] Specifically, let $\tilde{I}_{i,j} = [C_{i-1} - \sum_{j' \leq j} \gamma_{i,j'}, C_{i-1} - \sum_{j' < j} \gamma_{i,j'}]$.

[34] Specifically, let $\hat{C}_{\sigma(j)} = 1 - \sum_{i',j':j' \leq j} \gamma_{i',j'}$, and let $P(\tilde{I}_{i,j}) = [\hat{C}_{\sigma(j-1)} - \sum_{i' \leq i} \gamma_{i',j}, \hat{C}_{\sigma(j-1)} - \sum_{i' < i} \gamma_{i',j}]$.

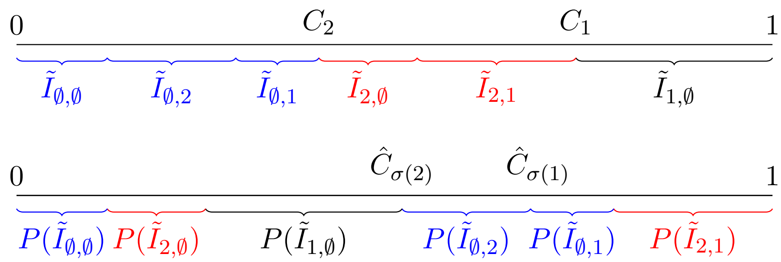

Figure 5: Constructing the permutation P for $n = 2$ schools, where σ is the identity permutation. The intervals $\tilde{I}_{i,j}$ for $i \leq \sigma(j) = j < N + 1$ are empty by definition, as all transfers are index-decreasing.

to assign students to all possible pairs of schools that are consistent with consistent preferences and the first- and second-round overdemand orderings. For example $\gamma_{1,j} = 0$ for all j, but any student whose first-round affordable set is X_1 is assigned to her top choice school in both rounds, and hence her second-round affordable set is inconsequential.

We show that there exists a PLDA mechanism with permutation P, where the students with first-round scores in $\tilde{I}_{i,j}$ are precisely the students with a first-round affordable set X_i and a second-round affordable set $\tilde{X}_j$, and that this PLDA mechanism gives the same joint distribution over first- and second-round assignments as M. To do this, we first show that this distribution of first- and second-round affordable sets gives rise to the correct joint first- and second-round assignments over all students. We then use anonymity to construct L in such a way as to have the correct first- and second-round assignment joint distributions for each student. Finally, we verify that these second-round affordable sets give the student-optimal stable matching under the second round school preferences given by P.

(2c) Equivalence of the joint distribution of assignments given by affordable sets and M.

Fix student preferences $\theta = (\succ, \hat{\succ})$. We show that if we let $\gamma_{i,j}$ be the proportion of students with preferences θ who have first- and second-round affordables X_i and $\tilde{X}_j$ respectively, then we obtain the same joint distribution over assignments in the first and second rounds for students with preferences θ as under mechanism M. In doing so, we will use the following "prefix lemma".

The "prefix lemma" states that for every set of schools S', there exist certain intervals of the form $I_i^j = [C_i, C_j]$ such that for any two student types whose top set of acceptable schools under second-round preference reports is S', the proportion of students with lotteries in I_i^j who are upgraded to a school in S' in the second round is the same for each type.

We define a *prefix* of preferences $\succ$ to be a set of schools S' that is a top set of acceptable schools under $\succ$; that is, for all $s' \in S'$ and $s \notin S'$, it holds that $s' \succ s$.

Lemma 2. *[Prefix Property] Let $s = s_j$ be a school, and let $S' \not\ni s$ be a set of schools such that $i(S') < j$. Let $\theta = (\succ, \hat{\succ})$ and $\theta' = (\succ', \hat{\succ}')$ be consistent preferences such that S' is a prefix of $\succ, \hat{\succ}$ and some students with preferences θ are assigned to school s in the first round, and similarly S' is a prefix of $\succ', \hat{\succ}'$ and some students with preferences θ' are assigned to school s in the first round.*

Then

$$\rho^{\theta}([C_j, C_{i(S')}], S') = \rho^{\theta'}([C_j, C_{i(S')}], S').$$

That is, the proportion of students of type θ whose first-round lotteries are in the interval $[C_j, C_{i(S')}]$ and who are assigned to a school in S' in the second round is the same as the proportion of students of type θ' whose first-round lotteries are in the interval $[C_j, C_{i(S')}]$ and who are assigned to a school in S' in the second round.

Sketch of proof of Lemma 2. The idea of the proof is to use weak strategy-proofness and first-order stochastic dominance to show that the probabilities of being assigned to S' (conditional on certain first-round assignments) are the same for students of type θ or θ'. We then invoke anonymity to argue that proportions of types of students assigned to a certain school are given by the conditional probabilities of individual students being assigned to that school. We present the full proof at the end of Section 3.1. $\square$

We now show that the mechanism M and the affordable set distribution $\gamma_{i,j}$ produce the same joint distribution of assignments.

(2c.i.) Students with two acceptable schools.

To give a bit of the flavor of the proof, we first consider student preferences θ of the form $\succ$: $[s, s', s_{N+1}]$ and $\hat{\succ}$: $[s, s_{N+1}]$, where all other schools are unacceptable. We let k, l be the indices such that $s = s_k$ and $s' = s_l$.

There are five ordered pairs of schools that students of this type can be assigned to in the two rounds. Namely, if we let (s, s') denote assignment to s in the first round and to s' in the second round, then the ordered pairs are (s, s), (s', s), (s', s_{N+1}), (s_{N+1}, s), and (s_{N+1}, s_{N+1}). Since the proportion of students with each first-round assignment is fixed, it suffices to show that the mechanism M and the mechanism that assigns first- and second-round affordable set distributions according to $\gamma_{i,j}$ produce the same proportion of students assigned to (s', s) and the same proportion of students assigned to (s_{N+1}, s).

Let $I_l^k = [C_l, C_k]$, and let $I_{N+1}^{\max\{k,l\}} = [0, C_{\max\{k,l\}}]$. The proportions of students with preferences θ who are assigned to (s', s) and (s_{N+1}, s) under M are given by $\rho^\theta([C_l, C_k], s)$ and $\rho^\theta([0, C_{\max\{k,l\}}], s)$ respectively. We want to show that this is the same as the proportion of students with preferences θ who are assigned to (s', s) and (s_{N+1}, s) respectively when first- and second-round affordable sets are given by the affordable set distribution $\gamma_{i,j}$. We remark that when $k > l$ this holds vacuously, since all the terms are 0. Hence, since for any school s the proportion of students with preferences θ who are assigned to s in the first round does not depend on θ, it suffices to consider the case where $k < l$.

Let $\theta' = (\succ', \hat{\succ}')$ be the preferences given by $\succ'$: $[s = s_k, s_{k+1}, \ldots, s_{l-1}, s' = s_l, s_{N+1}]$ and $\hat{\succ}'$: $[s = s_k, s_{N+1}]$, where only the schools with indices between k and l are acceptable in the first round, only $s = s_k$ is acceptable in the second round, and all other schools are unacceptable.

For all pairs of indices i, j such that $j < i$, let $\theta'_{i,j} = (\succ_{i,j}, \hat{\succ}_{i,j})$ be the student preferences such that $\succ_{i,j}$: $[s_j, s_{i-1}, s_i, s_{N+1}]$ and $\hat{\succ}_{i,j}$: $[s_j, s_{N+1}]$, with all other schools unacceptable. (In the case where $i = j + 1$, we let the first two schools under the preference ordering $\succ_{i,j}$ coincide.) We note that $\theta'_{i,j} = \theta_{i,\sigma^{-1}(j)}$, where $\theta_{i,j}$ was defined in (2b), and that for $i > \sigma(j)$ we previously defined $\rho_{i,j} = \sum_{l \le j} \gamma_{i,l}$ to be the proportion of students of type $\theta_{i,j}$ whose first-round assignment is s_i and whose second-round assignment is school s_j.

The proportion of students with preferences θ who are assigned to (s', s) under M is given by

$$
\begin{aligned}
\rho^\theta([C_l, C_k], s) &= \rho^{\theta'}([C_l, C_k], s) \text{ (by the prefix property (Lemma 2))} \\
&= \sum_{k < i \le l} \rho^{\theta'}([C_i, C_{i-1}], s) \\
&= \sum_{k < i \le l} \rho^{\theta'_{i,k}}([C_i, C_{i-1}], s) \\
&\quad \text{(since the second-round assignment does not depend on the first-round report)} \\
&= \sum_{k < i \le l} \rho_{i,\sigma^{-1}(k)} \text{ (by the definition of } \rho_{i,\sigma^{-1}(k)}) \\
&= \sum_{k < i \le l} \sum_{j \le \sigma^{-1}(k)} \gamma_{i,j} \text{ (by the definition of } \gamma_{i,j}),
\end{aligned}
$$

which is precisely the proportion of students with preferences θ who are assigned to (s', s) if the first- and second-round affordable sets are given by $\gamma_{i,j}$.

Similarly, let $\theta'' = (\succ'', \hat{\succ}'')$ be the preferences given by $\succ''$: $[s = s_k, s' = s_l, s_{l+1}, \ldots, s_N, s_{N+1}]$ and $\hat{\succ}''$: $[s = s_k, s_{N+1}]$, where only $s = s_k$ and the schools with indices greater than l are acceptable in the first round, only $s = s_k$ is acceptable in the second round, and all other schools are unacceptable. Then the proportion of students with preferences θ who are assigned to (s_{N+1}, s) under M is given by

$$\begin{aligned}
\rho^{\theta}([0, C_l], s) &= \rho^{\theta''}([0, C_l], s) \text{ (by the prefix property (Lemma 2))} \\
&= \sum_{l<i\leq N} \rho^{\theta''}([C_i, C_{i-1}], s) \\
&= \sum_{l<i\leq N} \rho^{\theta'_{i,k}}([C_i, C_{i-1}], s) \\
&\quad \text{(since the second-round assignment does not depend on the first-round report)} \\
&= \sum_{l<i\leq N} \rho_{i,\sigma^{-1}(k)} \text{ (by the definition of } \rho_{i,\sigma^{-1}(k)}) \\
&= \sum_{i,j:l<i\leq N, j\leq\sigma^{-1}(k)} \gamma_{i,j} \text{ (by the definition of } \gamma_{i,j}),
\end{aligned}$$

which is precisely the proportion of students with preferences θ who are assigned to (s_{N+1}, s) if the first- and second-round affordable sets are given by $\gamma_{i,j}$.

(2c.ii.) Students with general preferences.

We now consider general (consistent) student preferences θ of the form $(\succ, \hat{\succ})$, where

$$\succ: [s_{i_1}, s_{i_2}, \ldots, s_{i_k}, s_{N+1}] \text{ and } \hat{\succ} : [s_{i_1}, s_{i_2}, \ldots, s_{i_l}, s_{N+1}],$$

for some $k > l$ and where all other schools are unacceptable. We wish to show that for every pair of schools $s, s' \in \{s_{i_1}, s_{i_2}, \ldots, s_{i_k}, s_{N+1}\}$, the mechanism M and the mechanism that assigns first- and second-round affordable set distributions according to $\gamma_{i,j}$ produce the same proportion of students assigned to (s, s'). It suffices to show that for every prefix S' of the preferences $\hat{\succ}$ and every school $s \in \{s_{i_2}, \ldots, s_{i_k}, s_{N+1}\}$, the mechanism M and the mechanism that assigns first- and second-round affordable set distributions according to $\gamma_{i,j}$ produce the same proportion of students assigned to s in the first round and some school in S' in the second round. We say that the students are assigned to (s, S').

Fix a prefix S' of $\hat{\succ}$ and a school $s = s_{i_j}, 1 < j \leq k$. Let $m \leq k$ be such that $S' = \{s_{i_1}, s_{i_2}, \ldots, s_{i_m}\}$. If $j \leq m$ then $s \in S'$, and so in any mechanism that respects guarantees, the proportion of students assigned to (s, S') is the same as the proportion of students assigned to s in the first round.

Recall that $i(S')$ is the largest index of a school in S', i.e. if $i(S') = max\{i' : s_{i'} \in S'\}$. (Note that this is not necessarily i_m, the index of the school in S' that is least preferred by a student of type θ.) If $j > m$ and $i_j \leq i(S')$, then in the first round, whenever the school i_j is available in the first round, so is the preferred school $i(S')$; thus, for any school s', the proportion of students assigned to s in the first round is 0. It follows that in any mechanism that respects guarantees, the proportion of students assigned to (s, S') is 0.

From here on, we may assume that $j > m$ (i.e., $s \notin S'$) and $i_j > i(S')$. Since $i_j > i(S')$, the proportion of students with preferences θ who are assigned to (s, S') under M is given by $\rho^\theta([C_{i_j}, C_{i(S')}], S')$. Let $i(\sigma(S'))$ be the index i such that $s_i \in S'$ and $\sigma^{-1}(i)$ is maximal, that is, the index of the school in S' that is most affordable in the second round.

Let $\theta' = (\succ', \hat{\succ}')$ be the preferences given by

$$\succ': [s_{i(\sigma(S'))}, s', s_{i(S')+1}, s_{i(S')+2}, \cdots, s_{i_j-1}, s_{i_j}, s_{N+1}] \text{ and } \hat{\succ}': [s_{i(\sigma(S'))}, s', s_{N+1}]$$

for all $s' \in S' \setminus \{s_{i(\sigma(S'))}\}$. It may be helpful to think of this as all preferences of the form

$$\succ': [S', s_{i(S')+1}, s_{i(S')+2}, \cdots, s_{i_j-1}, s_{i_j}, s_{N+1}] \text{ and } \hat{\succ}': [S', s_{N+1}],$$

where the school $s_{i(\sigma(S'))}$ comes first and otherwise the schools in S' are ordered arbitrarily, and where all schools between $s_{i(S')}$ and s_{i_j} are acceptable in the same order as first round overdemand.

We remark that only the schools in S' and the schools with indices between $i(S')$ and i_j are acceptable in the first round, only the schools in S' are acceptable in the second round, and all other schools are unacceptable. Since $j > m$, $i_j > i(S')$, and the preferences θ are consistent, the preferences θ' are well defined. Let $\theta'' = (\succ'', \hat{\succ}'')$ be the preferences given by $\succ'' = \succ'$ and $\hat{\succ}'': [s_{i(\sigma(S'))}, s_{N+1}]$.

Recall that for all $i > i(\sigma(S'))$, $\theta'_{i,i(\sigma(S'))} = (\succ_{i,i(\sigma(S'))}, \hat{\succ}_{i,i(\sigma(S'))})$ are the student preferences such that

$$\succ_{i,i(\sigma(S'))}: [s_{i(\sigma(S'))}, s_{i-1}, s_i, s_{N+1}] \text{ and } \hat{\succ}_{i,i(\sigma(S'))}: [s_{i(\sigma(S'))}, s_{N+1}],$$

with all other schools unacceptable. Additionally, recall that $\rho_{i,\sigma^{-1}(i(\sigma(S')))}$ is the proportion of students of type $\theta_{i,i(\sigma(S'))}$ whose first-round assignment is s_i and whose second-round assignment is school $s_{i(\sigma(S'))}$.

Let $\hat{S} = \{s_{i_1}, s_{i_2}, \ldots, s_{i_{j-1}}\}$, and let $i(\hat{S})$ be the index i such that $i \in \hat{S}$ and $\sigma^{-1}(i)$ is maximal, that is, the index of the school preferable to s under $\succ$ that is most affordable in the second round.

Then the proportion of students with preferences θ who are assigned to (s, S') under M is given by $\rho^{\theta}([C_{i_j}, C_{i(\hat{S})}], S')$, where

$$\begin{aligned}
\rho^{\theta}([C_{i_j}, C_{i(\hat{S})}], S') &= \rho^{\theta'}([C_{i_j}, C_{i(\hat{S})}], S') \text{ (by the prefix property (Lemma 2) with prefix } S'\text{)} \\
&= \sum_{i(\hat{S})<i\le i_j} \rho^{\theta'}([C_i, C_{i-1}], S') \\
&= \sum_{i(\hat{S})<i\le i_j} \rho^{\theta'}([C_i, C_{i-1}], s_{i(\sigma(S'))}) \\
&\qquad \text{(by the definition of the second-round overdemand ordering)} \\
&= \sum_{i(\hat{S})<i\le i_j} \rho^{\theta''}([C_i, C_{i-1}], s_{i(\sigma(S'))}) \text{ (by the prefix property with prefix } \{s_{i(\sigma(S'))}\}\text{)} \\
&= \sum_{i(\hat{S})<i\le i_j} \rho^{\theta'_{i,i(\sigma(S'))}}([C_i, C_{i-1}], s_{i(\sigma(S'))}) \\
&\quad \text{(since the second-round assignment does not depend on the first-round report)} \\
&= \sum_{i(\hat{S})<i\le i_j} \rho_{i,\sigma^{-1}(i(\sigma(S')))} \text{ (by the definition of } \rho_{i,\sigma^{-1}(i(\sigma(S')))}\text{)} \\
&= \sum_{i(\hat{S})<i\le i_j} \sum_{j'\le\sigma^{-1}(i(\sigma(S')))} \gamma_{i,j'} \text{ (by the definition of } \gamma_{i,j'}\text{)},
\end{aligned}$$

which is precisely the proportion of students with preferences θ who are assigned to (s, S') if the first- and second-round affordable sets are given by $\gamma_{i,j'}$. Note that all $\theta'_{i,i(\sigma(S'))}$ and $\rho_{i,\sigma^{-1}(i(\sigma(S')))}$ in the summation are well-defined, since the sum is over indices satisfying $i > i\left(\hat{S}\right)$, and since $j > m$ it follows that $\hat{S} \supseteq S'$ and hence $i > i\left(\hat{S}\right) \ge i\left(\sigma\left(S'\right)\right)$.

(2d) Constructing the lottery L.

Fix a student λ who reports first- and second-round preferences $\theta = (\succ, \hat{\succ})$. Suppose that λ is assigned to schools (s_i, s_j) in the first and second rounds respectively. We first characterize all first- and second-round budget sets consistent with the overdemand orderings that could have led to this assignment. Let $\underline{i}$ be the smallest index i' such that $\max_{\succ} X_{i'} = s_i$, let $\underline{j}$ be the smallest index j' such that $\max_{\hat{\succ}} \tilde{X}_{j'} \cup \{s_i\} = s_j$, and let $\overline{j}$ be the largest index j' such that $\max_{\hat{\succ}} \tilde{X}_{j'} \cup \{s_i\} = s_j$. Then the set of first- and second-round budget sets that student λ could have been assigned by the mechanism is given by $\{X_{i'}, X_{j'} \cup \{s_i\} : \underline{i} \le i' \le i, \underline{j} \le j' \le \overline{j}\}$. (We remark that the asymmetry in these definitions is due to the existence of the first-round guarantee in the second-round budget sets.)

Conditional on λ being assigned to schools (s_i, s_j) in the first and second rounds respectively, we assign a lottery number $L(\lambda)$ to λ distributed uniformly over the union of intervals $\cup_{i',j' : \underline{i} \le i' \le i, \underline{j} \le j' \le \overline{j}} \tilde{I}_{i',j'}$,

$$(\ L(\lambda) \mid (\mu(\lambda), \tilde{\mu}(\lambda)) = (s_i, s_j)) \ \sim \ \text{Unif}\left(\cup_{i',j' : \underline{i} \le i' \le i, \underline{j} \le j' \le \overline{j}} \tilde{I}_{i',j'}\right),$$

independent of all other students' assignments.

We show that this is consistent with the first round of the mechanism being RSD. We have shown in *(1)* that if for each pair of reported preferences $\theta = (\succ, \hat{\succ}) \in \Theta$, a uniform proportion $\gamma_{i',j'}$ of students with reported preferences θ are given first- and second-round budget sets $X_{i'}, \{s^\theta\} \cup \tilde{X}_{j'}$ (where $s^\theta = \max_{\succ} X_i$ is the first-round assignment of such students), we obtain the same distribution of assignments as M. Since M is anonymous and satisfies the averaging axiom, and since $|\tilde{I}_{i',j'}| = \gamma_{i',j'}$, it follows that each student's first-round lottery number is distributed as Unif$[0, 1]$.

Given the constructed lottery L, we construct the second-round cutoffs $\hat{C}_i$ for the PLDA and verify that the assignment $\tilde{\mu}$ is feasible and stable with respect to the schools' second-round preferences, as defined by $P \circ L$ and the guarantee structure. Specifically, in PLDA, each student with a first-round score l and a first-round assignment s has a second-round score $\hat{r}_i = P(l) + \mathbb{1}(s = s_i)$ at each school $s_i \in S$, and students are assigned to their favorite school s_i at which their second-round score exceeds the school's second-round cutoff, $\hat{r}_i \ge \hat{C}_i$ (or to the outside option s_{N+1}).

Recall that the schools are indexed so that $C_1 \ge C_2 \ge \cdots \ge C_{N+1}$, and that the permutation σ is chosen so that the second-round overdemand ordering is given by $s_{\sigma(1)}, s_{\sigma(2)}, \ldots, s_{\sigma(N+1)} = s_{N+1}$, and so it should follow that the second-round cutoffs $\hat{C}_i$ satisfy $\hat{C}_{\sigma(1)} \ge \hat{C}_{\sigma(2)} \ge \cdots \ge \hat{C}_{\sigma(N+1)}$.

By the characterization of stable assignments given by Azevedo and Leshno (2016), it suffices to show that if each student with a first-round assignment s and second-round lottery number in $[\hat{C}_{\sigma^{-1}(i)}, \hat{C}_{\sigma^{-1}(i-1)}]$ is assigned to her favorite school in $\{s\} \cup \tilde{X}_i$, where we define $\tilde{X}_i = \{s_{\sigma(i)}, s_{\sigma(i+1)}, \ldots, s_{\sigma(N+1)}\}$, then the resulting assignment $\hat{\mu}$ is equal to the second-round assignment $\tilde{\mu}$ of our mechanism M, and satisfies that $\eta(\hat{\mu}^{-1}(s_i)) \le q_i$ for any school s_i, and $\eta(\hat{\mu}^{-1}(s_i)) = q_i$ if $\hat{C}_i > 0$.

For fixed i, j, let $\hat{C}_{\sigma(j)} = 1 - \sum_{i',j':j'\leq j} \gamma_{i',j'}$ and let $\hat{C}_{i,\sigma(j)} = \hat{C}_{\sigma(j-1)} - \sum_{i'\leq i} \gamma_{i',j}$. (We remark that since $\gamma_{i,j}$ refers to the i-th school to fill in the first round, s_i, and the j-th school to fill in the second round, $s_{\sigma(j)}$, the $\hat{C}$ are indexed slightly differently than $\gamma_{i,j}$ is.)

We use the averaging assumption and the equivalence of assignment probabilities that we have shown in (1) to conclude that if $\hat{\mu}$ is the assignment given by running DA with round scores $\hat{r}$ and cutoffs $\hat{\mathbf{C}}$, then $\tilde{\mu} = \hat{\mu}$.

This is fairly evident, but we also show it explicitly below. Specifically, consider a student $\lambda \in \Lambda$ with a first-round lottery number $L(\lambda)$ and reported preferences $\theta = (\succ, \hat{\succ})$. Let i, j be such that $L(\lambda) \in \cup_{i',j':\underline{i}\leq i'\leq i, \underline{j}\leq j'\leq \overline{j}} \tilde{I}_{i',j'}$, where $\underline{i}$ is the smallest index i' such that $\max_{\succ} X_{i'} = s_i$, $\underline{j}$ is the smallest index j' such that $\max_{\hat{\succ}} \tilde{X}_{j'} \cup \{s_i\} = s_j$, and $\overline{j}$ is the largest index j' such that $\max_{\hat{\succ}} \tilde{X}_{j'} \cup \{s_i\} = s_j$. Then, because of the way in which we have constructed the lottery L, $(\mu(\lambda), \tilde{\mu}(\lambda)) = (s_i, s_j)$.

Moreover, since

$$
\begin{aligned}
P(L(\lambda)) \quad &\in \quad P(\cup_{i',j':\underline{i}\leq i'\leq i, \underline{j}\leq j'\leq \overline{j}} \tilde{I}_{i',j'}) \\
&= \quad \cup_{i',j':\underline{i}\leq i'\leq i, \underline{j}\leq j'\leq \overline{j}} P(\tilde{I}_{i',j'}),
\end{aligned}
$$

where $P(\tilde{I}_{i',j'}) \in [\hat{C}_{\sigma(j')}, \hat{C}_{\sigma(j'-1)}]$, it holds that under $\hat{\mu}$, student λ receives her favorite school in $\{s_i\} \cup \tilde{X}_{j'}$ for some $\underline{j} \leq j' \leq \overline{j}$, which is the school s_j. Hence $\tilde{\mu}(\lambda) = \hat{\mu}(\lambda) = s_j$.

It follows immediately that the assignment $\hat{\mu}$ is feasible, since it is equal to the feasible assignment $\tilde{\mu}$.

Finally, let us check that the assignment is stable. Suppose that $\hat{C}_j > 0$. We want to show that $\eta(\tilde{\mu}^{-1}(s_j)) = q_j$. First note that it follows from the definition of $\hat{C}_j$ that

$$
1 > \sum_{i',j':j'\leq \sigma^{-1}(j)} \gamma_{i',j'} = \sum_{i'} \rho_{i',\sigma^{-1}(j)}.
$$

Consider student preferences $\theta = (\succ, \succ)$ given by $\succ$: $[s_j, s_1, s_2, \ldots, s_{j-1}, s_{j+1}, \ldots, s_{N+1}]$. Then $\sum_{i'} \rho_{i',\sigma^{-1}(j)}$ is the proportion of students of type θ who are assigned to school s_j in the second round, which, by assumption, is also the probability that a student with preferences θ is assigned to s_j in the second round. But since M is non-wasteful, this means that $\eta(\tilde{\mu}^{-1}(s_j)) = q_j$. It follows from constrained Pareto efficiency that the output of M is the student-optimal stable matching.

$\square$

Proof of Lemma 2. Here, we prove the prefix property. We first observe that any schools reported to be acceptable but ranked below s in the first round are inconsequential. Moreover, since M respects guarantees, weak two-round strategy-proofness implies that any schools reported to be acceptable but ranked below s in the second round are inconsequential. Hence it suffices to prove the lemma for first-round preference orderings $\succ$ and $\succ'$ for which s is the last acceptable school.

Suppose that the lemma holds for $i = i(\mathcal{C}')$. Then if $i(\mathcal{C}') = i' < i$ it holds that

$$\begin{aligned}
\rho^{\theta}\left([C_j, C_i], \mathcal{C}'\right) &= \frac{\rho^{\theta}\left(\left[C_j, C_{i(\mathcal{C}')}\right], \mathcal{C}'\right)\left(C_{i(\mathcal{C}')} - C_j\right) - \rho^{\theta}\left(\left[C_i, C_{i(\mathcal{C}')}\right], \mathcal{C}'\right)\left(C_{i(\mathcal{C}')} - C_i\right)}{C_i - C_j} \\
&= \frac{\rho^{\theta'}\left(\left[C_j, C_{i(\mathcal{C}')}\right], \mathcal{C}'\right)\left(C_{i(\mathcal{C}')} - C_j\right) - \rho^{\theta'}\left(\left[C_i, C_{i(\mathcal{C}')}\right], \mathcal{C}'\right)\left(C_{i(\mathcal{C}')} - C_i\right)}{C_i - C_j} \\
&= \rho^{\theta'}\left([C_j, C_i], \mathcal{C}'\right),
\end{aligned}$$

where the first and last equalities follow from Bayes' rule, and the second equality holds since the lemma holds for $i = i(\mathcal{C}')$, and the theorem follows. Hence it suffices to prove the lemma for $i = i(\mathcal{C}')$.

Let $i_1, \ldots, i_k$ be the indices of the schools in S', in increasing order. We observe that $i_k = i(S')$. Recall that $s = s_j$, where $i_k < j$.

Since we wish to prove that the lemma holds for all pairs θ, θ' satisfying the assumptions, it suffices to show that the lemma holds for a fixed preference θ when we vary only θ'. Therefore, we may, without loss of generality, fix the preferences θ to satisfy that

$$\succ: [s_{i(S')}, s_{i_1}, \ldots, s_{i_{k-1}}, s = s_j, s_{N+1}] \text{ and } \hat{\succ}: [s_{i(S')}, s_{i_1}, \cdots, s_{i_{k-1}}, s_{N+1}],$$

and all other schools are unacceptable. That is, the worst school in S' is top ranked, then all other schools in S' in order. In the first round $s = s_j$ is also acceptable, and in the second round only schools in S' are acceptable.

We remark that given the first-round ordering, the worst school in S' and the school s (namely, $s_{i(S')}$ and s_j) are the only acceptable schools to which students of type θ will be assigned in the first round. Moreover, the proportion of students with preferences θ (or θ') who can deduce that their score is in $[C_j, C_{i(S')}]$ is precisely $C_{i(S')} - C_j$, since such students are assigned in the first round to some school not in S' that they weakly prefer to s, and all such schools are between $s_{i(S'')}$ and s_j

in the overdemand ordering. Similarly, the proportion of students with preferences θ (or θ') who can deduce that their lottery number is in $[C_{i(S')}, 1]$ is precisely $1 - C_{i(S')}$, since such students are assigned in the first round to a school in S'. (Note that students with preferences θ' may be able to deduce that their lottery number falls in a subinterval of the interval we have specified. However, this does not affect our statements.)

To compare the proportion of students of types θ and θ' whose scores are in $[C_j, C_{i(S')}]$ and who are assigned to S' in the second round, we define a third student type θ'' as follows. Let $\theta'' = (\succ', \hat{\succ})$ be a set of preferences where the first-round preferences are the same as the first-round preferences of θ', and the second-round preferences are the same as the second-round preferences of type θ.

Let λ be a student with preferences θ, and similarly let λ' be a student with preferences θ'. We use the two-round strategy-proofness of the mechanism to show that λ has the same probability of being assigned to some school in S' in the second round as if she had reported type θ'', and similarly for λ'. Since the proportion of students of either type being assigned to a school in S' in the first round is the same and the mechanism respects guarantees, this is sufficient to prove the prefix property.

Formally, let ρ be the probability that λ is assigned to some school in S' in the second round if she reports truthfully, conditional on being able to deduce that her first-round score is in $[C_j, C_{i(S')}]$, and let ρ' be the probability that λ' is assigned to some school in S' in the second round if she reports truthfully, conditional on being able to deduce that her first-round score is in $[C_j, C_{i(S')}]$. (We note that given her first-round assignment $\mu(\rho')$, the student ρ' may actually be able to deduce more about her first-round score, and so the interim probability after knowing her assignment that ρ' is assigned to some school in S' in the second round if she reports truthfully is not necessarily ρ'.) Let ρ'' be the probability that a student with preferences θ'' and a first-round score in $[C_j, C_{i(S')}]$ chosen uniformly at random is assigned to some school in S' in the second round. It follows from the design of the first round and from anonymity that ρ is the probability that a student with preferences θ and a lottery number in $[C_j, C_{i(S')}]$ chosen uniformly at random is assigned to some school in S' in the second round, and similarly for ρ'.

Proving the lemma is equivalent to proving $\rho = \rho'$. We show that $\rho = \rho'' = \rho'$. Note that the first equality is between preferences that are identical in the second round, and the second equality is between preferences that are identical in the first round.

We first show that $\rho = \rho''$; that is, changing just the first-round preferences does not affect the probability of assignment to S'. This is almost immediate from first-order stochastic dominance of truthful reporting, since the second-round preferences under θ and θ'' are identical. (This also illustrates the power of the assumption that the second-round assignment does not depend on first-round preferences. It implies that manipulating first-round reports to obtain a more fine-grained knowledge of the lottery number does not help, since assignment probabilities are conditionally independent of the lottery number.) We present the full argument below.

Let π be the probability that a student with preferences θ who is unassigned in the first round is assigned to a school in S' in the second round. We note that since the last acceptable school under preferences θ and θ' is $s = s_j$, the set of students with preferences θ who are unassigned in the first round is equal to the set of students with preferences θ with lottery number in $[0, C_j]$, and similarly the set of students with preferences θ'' who are unassigned in the first round is equal to the set of students with preferences θ'' with lottery number in $[0, C_j]$. Hence, the fact that θ and θ'' have the same second preferences gives us that π is also the probability that a student with preferences θ'' who is unassigned in the first round is assigned to a school in S' in the second round.

The probability of being assigned in the second round to a school in S' when reporting θ is given by:

$$(1 - C_{i(S')}) + (C_{i(S')} - C_j)\rho + C_j\pi,$$

The probability of being assigned in the second round to a school in S' when reporting θ'' is given by:

$$(1 - C_{i(S')}) + (C_{i(S')} - C_j)\rho'' + C_j\pi,$$

It follows from first-order stochastic dominance of truthful reporting for types θ and θ' that $\rho = \rho''$.

We now show that $\rho' = \rho''$. This is a little more involved, but essentially relies on breaking the set of students with first-round score in $[C_j, C_{i(S')}]$ into smaller subsets, depending on their first-round assignment, and using first-order stochastic dominance of truthful reporting to show that in each subset, the probability of an arbitrary student being assigned to a school in S' in the second round is the same for students with either set of preferences θ' or θ''.

We first introduce some notation for describing the first-round preferences of θ' and θ''. Let $\{j_1 \leq \cdots \leq j_m\}$ be the indices between $i(S')$ and j corresponding to schools that a student with preferences θ' and a lottery number in $[C_j, C_{i(S')}]$ could have been assigned to in the first round.

Formally, we define them to be the indices k for which $s_k \notin S'$, $i(S') < k \leq j$, $s_k \succeq' s_j$ and s_k is relevant in the first-round overdemand ordering, that is, $k' < k$ for all k' such that $s_{k'} \succ' s_k$. We observe that $j_m = j$. For $l = 1, \ldots, m$, let ρ'_l be the probability that a student with preferences θ' who was assigned to school s_{j_l} is assigned to a school in S' in the second round.

The set of students with preferences θ' assigned to school s_{j_l} in the first round is precisely the set of students with preferences θ' whose first-round lottery number is in $[C_{j_l}, C_{j_{l-1}}]$ and similarly the set of students with preferences θ'' assigned to school s_{j_l} in the first round is precisely the set of students with preferences θ'' whose first-round lottery number is in $[C_{j_l}, C_{j_{l-1}}]$. If we define $j_0 = i$, it follows that $(C_{i(S')} - C_j) = \sum_{l=1}^{m}(C_{j_{l-1}} - C_{j_l})$, and that

$$(C_{i(S')} - C_j)\rho' = \sum_{l=1}^{m}(C_{j_{l-1}} - C_{j_l})\rho'_l.$$

Let ρ''_l be the probability that a student with preferences θ'' who was assigned to school s_{j_l} is assigned to a school in S' in the second round. Then it also holds that

$$(C_{i(S')} - C_j)\rho'' = \sum_{l=1}^{m}(C_{j_{l-1}} - C_{j_l})\rho''_l.$$

We show now that $\rho''_l = \rho'_l$ for all l, which implies that $\rho' = \rho''$.

Consider a student λ_l who reported $\succ'=\succ''$ in the first round and was assigned to school s_{j_l}. Note that such a report is consistent with either reporting θ' or θ'', and since the first-round reports of these types are the same and the first-round mechanism is DA-STB there exists some set of lottery numbers L_l such that students of type θ' and θ'' are assigned to j_l in the first round if and only if their lottery lies in L_l. The probabilities that this student is assigned in the second round to a school in $\mathcal{C}'$ when reporting θ' and θ'' are given by $\rho'{}_l$ and $\rho''{}_l$ respectively. Now for any fixed lottery $L(\lambda)$, truthful reporting is a dominant strategy in the second round for types θ and θ'. It follows that $\rho'{}_l = \rho''_l$.

This completes the proof of the lemma. $\square$