

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

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

April 13, 2017

Abstract

In many centralized school admission systems, a significant fraction of allocated seats are later vacated, often due to students obtaining better outside options. We consider the problem of reassigning these seats in a fair and efficient manner while also minimizing the movement of students between schools. Centralized admissions are typically conducted using the Deferred Acceptance (DA) algorithm, with a lottery used to break ties caused by indifferences in school priorities. We introduce the Permuted Lottery Deferred Acceptance (PLDA) mechanisms, which reassign vacated seats using a second round of Deferred Acceptance with a lottery given by a suitable *permutation* of the first round lottery numbers. We show that a mechanism based on a simple *reversal of the first round lottery order* performs the best among all PLDA mechanisms. We also characterize PLDA mechanisms as the class of truthful mechanisms satisfying some natural efficiency and fairness properties. 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 many public school systems throughout the United States, students and families are required to submit preferences over the schools for which they are eligible. As this is done fairly early in the academic year, students typically do not know their options outside of the public school system at the time they submit preferences. As a consequence, a significant fraction of the students do not use their allotted seat in a public school, leading to significant inefficiency. For example, in the NYC public high school system, over 80,000 students are assigned a public school seat each year in March, and about 10% of these students choose not to attend a public school in September, possibly opting instead to attend a private or charter school.¹ A well-designed reassignment process, run after students learn about their outside options, could lead to significant gains in overall welfare. Yet there is no known systematic way of reassigning unused seats. Our goal is to design an explicit reassignment mechanism run at a late stage of the matching process that reassigns these seats well.

*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

¹In 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.

During the past fifteen years, insights from matching theory have informed the design of school choice programs in several 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, such as changes in student preferences, in designing assignment mechanisms 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-stage model of school assignment with finitely many schools. Students initially submit their ordinal preferences over schools, and receive a first round assignment based on these preferences via the standard Deferred Acceptance mechanism (DA). School preferences are given by weak priorities, and ties are broken via a single lottery ordering across all schools.² Afterwards, some students may be presented with better 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 re-submit their (new) ordinal preferences over schools. The goal is to reassign the seats so that the resulting assignment is efficient, fair, and so that the overall (two-stage) mechanism is strategy-proof and does not penalize students for participating in the second round.

A natural starting point for reallocating seats is to simply re-run DA on the new preferences, using the same school preferences (priorities and tie-breaking) as in the first round. However, this approach may result in a cascade of reassignments of students from one school to another, which is costly for school administration and for students. Alternatively, the second round mechanism could allocate the vacant seats first to those students who were unassigned in the first round so as to reduce reassignment costs. Unfortunately, students can readily manipulate such a reassignment mechanism.³ The challenge is to design a reallocation mechanism that retains the good properties of DA while also avoiding its potentially high reassignment costs.

We suggest a class of mechanisms with good incentive properties — the *permuted lottery deferred acceptance mechanisms* (PLDA)— in which the assignments in both stages are given by DA, the initial assignment serves as a guarantee in the second round, and the lottery numbers across the two rounds are correlated. The mechanisms first break ties in school priorities by a single lottery ordering of the students, and a first-round assignment is computed by running DA; in the second round, each school first prioritizes students who were assigned to it in the first round over those who were not, and within each of the resulting two classes, students are prioritized according to their initial priorities at the school; finally, further ties across all schools are broken via a *permutation* of the (first-round) lottery numbers, and a second-round assignment is computed by running DA.

Our key insight is that the mechanism designer can leverage the correlation between tie-breaking lotteries to achieve operational goals. In particular, we show that *reversing* the lottery between the two rounds reduces the number of reassigned students. Our main theoretical result is that under a simple and intuitive condition, which we term the *order condition*, all PLDAs

²This model for the first round assignment is consistent with current practice at a number of school choice systems, including at New York City, Chicago and Denver (see NYC Department of Education [2017], Chicago Public Schools [2017], Denver Public Schools [2017], Pathak and Sönmez [2013]).

³Specifically, by submitting truncated preference lists initially, students are either assigned one of their top choices in the first round, or receive high priority in the second round.




Figure 1: Running DA with the same lottery creates a cascade of reassignments, whereas reversing the lottery minimizes reassignment.

In this example, there are 6 students with common preferences and 6 schools with a single priority group. All students prefer schools in the order $s_1 \succ s_2 \succ \dots \succ s_6$. 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 all students to the school one better on their preference list, whereas reversing the tie-breaking lottery reassigns only the student initially assigned to s_6 .

produce the same distribution over the final assignment, and the *reverse lottery DA (RLDA)*⁴ minimizes reassignment. In other words, when the order condition holds, RLDA is provably optimal among PLDAs with respect to both ex ante allocative efficiency and minimizing reassignment. The order condition can be interpreted as all schools having the same relative overdemand in the two rounds, despite changes in student preferences, and is satisfied when dropouts are uniform across first round student preferences, or when students have common preferences, as in Figure 1. (Our theoretical result holds in more general settings with heterogeneous student preferences and arbitrary priorities at schools.) We also give an axiomatic justification for the class of PLDA mechanisms. In the case with no school priorities, they are equivalent to the class of mechanisms that are two-round strategy-proof while satisfying natural efficiency and fairness requirements.

We empirically assess the performance of RLDA using data from the New York City high school system. We investigate a class of PLDAs that include 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 that all these mechanisms perform fairly similarly in terms of allocative efficiency, but RLDA reduces the number of reassigned students significantly. For instance, on the NYC public school system data set from 2004-2005, we find that FLDA results in about 7,200 reassignments out of a total of about 75,000 who remained in the public school system, whereas RLDA results in fewer than 3,100 reassignments.

1.1 Related Work

The mechanism design approach to school choice was first formulated by [Balinski and Sonmez, 1999] and [Abdulkadiroglu and Sönmez, 2003]. Since then, many economists have worked closely with school authorities to redesign school choice systems.⁵ These centralized mechanisms appear

⁴The student with the worst first round lottery number is given the highest priority in the second round.

⁵See for example New York City by [Abdulkadirolu et al., 2005a] and Boston by [Abdulkadirolu et al., 2005b] in 2003 and 2005 respectively, followed by New Orleans (2012), Denver (2012), and Washington DC (2013), among others.

to outperform the uncoordinated and ad-hoc assignment systems that they replaced ([Abdulkadiroğlu et al., 2015]). A significant portion of the theoretical literature has focused on the relative merits of two canonical mechanisms, Deferred Acceptance (DA), first introduced by [Gale and Shapley, 1962], and Top Trading Cycles (TTC), as well as their variations. We refer the reader to recent surveys by [Pathak, 2011] and [Abdulkadiroğlu and Sonmez, 2011] of the rich and growing (theoretical and empirical) literature on school choice problems.

One strand of literature explores how ties are broken within priority groups in DA. [Abdulkadiroğlu et al., 2009] empirically compared single tie-breaking and multiple tie-breaking for DA, and recent papers of [Arnosti, 2015], [Ashlagi and Nikzad, 2016] and [Ashlagi et al., 2015] study these tie-breaking rules analytically. We note that in our model, PLDA mechanisms use tie-breaking to reduce the number of reassignments while still satisfying natural efficiency and fairness properties.⁶

A second strand of literature that is relevant to our model is the work of [Abdulkadiroğlu and Sonmez, 1999] on house allocation models with existing tenants (or housing endowments). The second round in our model can be thought of as school seat allocation with existing endowments. [Abdulkadiroğlu and Sonmez, 1999] prove that a generalization of TTC is strategy-proof, Pareto efficient, and individually rational⁷ in this setting. However, directly applying the generalized TTC mechanism in the two-round setting suffers from some drawbacks. First, in our model the endowments are computed endogenously from the preferences. This means that a strategy-proof mechanism in the housing model is not necessarily strategy-proof in our two-round model, as a student has an incentive to manipulate their first round endowment. Secondly, while the TTC mechanism is natural in a model with endowments, it is less appropriate for assigning public school seats to students, where it is perfectly reasonable for a student to have priorities giving them a right to *attend* a school, but not to be able to *trade* these rights with each other.

Our work is among a growing number of papers that consider a dynamic model for school admissions. [Compte and Jehiel, 2008] consider using DA to reassign agents who each hold a position at an organization. They show that a modified version of DA that prioritizes agents currently holding a position is strategy-proof, stable, and respects individual guarantees.⁸ [Combe et al., 2016] study a more general model that also incorporates new agents, and show that the modified version of DA can be improved significantly in terms of overall welfare while also reducing the degree of instability. However, a critical distinction between this stream of work and ours is that in our model, the initial endowment is determined endogenously by preferences, and so students could manipulate their first round endowment to improve their final assignment.

A number of recent papers focus on the strategic issues in dynamic reassignment, such as [Dur, 2012, Kadam and Kotowski, 2014, Pereyra, 2013], and also propose using modified versions of DA in each round. These works develop appropriate solution concepts in finite markets with specific cross-period constraints and propose DA-like mechanisms that implement them. In recent work that is complementary to ours, [Narita, 2016] analyzes the preference data from NYC school choice system and observes that a significant proportion of preferences change after the initial match. Narita also considers a modified version of DA in this setting and establishes that it has good incentive and efficiency properties. We similarly propose PLDA mechanisms for their desirable incentive and efficiency properties. However, we focus on the problem of reassignment in school choice, where the emphasis is on the final assignment after two rounds, and the cost of reassignment. In contrast to [Narita, 2016], we consider a setting where outside options are realized after the initial assignment, and preferences are otherwise unchanged. In

⁶Prior work has exploited indifferences to achieve other ends. For example, [Erdil and Ergin, 2008] show how to improve allocative efficiency, and [Ashlagi and Shi, 2014a] show how to increase community cohesion.

⁷An allocation is individually rational if every agent weakly prefers his assignment to his original endowment.

⁸A two-round mechanism respects individual guarantees if each agent is assigned to a weakly better position in the second round.

addition, we provide an axiomatic justification for the class of PLDA mechanisms. This provides us with a natural class of mechanisms over which to optimize. We also exploit indifferences in school priorities to reduce the number of reassignments, which is not addressed in prior literature.

Our work connects with several strands in the operations management literature. [Ashlagi and Shi, 2014a] consider the problem of improving community cohesion in school choice, and find that correlating the lottery number of students belonging to the same community can be used to improve cohesion. [Ashlagi and Shi, 2014b] consider the problem of optimal allocation without money, motivated also by the school choice problem. Both these papers provide an axiomatic characterization of a class of candidate mechanisms and optimize over the class, which is similar to our approach. 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 the number of reassignments within this class. This is similar to the choice of a service policy in a queueing system, e.g., first-in-first-out (FIFO), last-in-first-out (LIFO), shortest-remaining-processing-time (SRPT), etc., wherein a particular policy is chosen in order to minimize an appropriate cost function such as the expected waiting time, e.g., see [Lee and Srinivasan, 1989]. “Work-conserving” service policies such as these result in different expected waiting times though the achievable throughput is identical, similar in spirit to our finding that different PLDA mechanisms differ in the number of reassignments but have identical allocative efficiency (under some conditions). 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 the OM literature.

2 Model

We begin by informally describing our setting. We consider the problem of assigning seats at a finite set S of schools to a set of students. Each school partitions the students into priority groups that are exogenously determined and publically known. Each student submits a strict preference ordering over the schools that she finds acceptable. A *single* lottery ordering of the students 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. Seats are initially assigned according to the student optimal Deferred Acceptance (DA) algorithm, as follows. In each step, unassigned students apply to their favorite school that has not yet rejected them. Each school s_i tentatively assign seats to the top q_i students who have applied to it, where q_i is the capacity of school s_i , and rejects any remaining students who have applied to it. 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. We remark that the strict student preferences, weak school priorities, and use of DA with single tie-breaking are consistent with many school choice systems, such as those in New York City, Chicago and Denver (see e.g. [Abdulkadiroglu and Sönmez, 2003]).

After this initial assignment, some students are subsequently presented with better outside options—such as admission to a private school that is not in S —and may no longer be interested in the seats assigned to them, effectively vacating those seats. After these outside options are revealed, each student submits her new ordinal preferences over schools, and a reassignment is computed. Since the reassignment occurs at a relatively late stage, moving students from one school to another is costly, potentially for both schools and students. Our goal is to design a procedure to reassign students to schools that minimizes the amount of student movement with respect to the initial assignment, while satisfying appropriate notions of efficiency, fairness, and incentive compatibility.

The timeline for the mechanism design problem considered here is as follows. Students submit




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

first round preference reports $\succ$,⁹ the mechanism designer obtains the first round assignment μ by running DA with uniform-at-random single tie-breaking, and the mechanism designer announces μ . Then, students observe their outside option, and update their preferences accordingly. Finally, students submit their updated second round preference reports $\hat{\succ}$, the mechanism designer obtains the second round assignment $\hat{\mu}$ by running a reassignment mechanism M , and the mechanism designer announces $\hat{\mu}$. We illustrate this in Figure 2.

We describe two models for the problem of assigning seats to students. The *discrete* model (Section 2.1), which assumes a finite set of students, can be easily translated into implementable mechanisms and is used in the empirical analysis (Section 5). The *continuum* model (Section 2.2), which assumes a finite number of schools and a continuum of students, is used in most of the theoretical analysis (Section 3.1). Intuitively, one could think of the continuum model as a reasonable approximation of the discrete model when the number of students is large.¹⁰ Our continuum model can be viewed as a two-round version of the model introduced by Azevedo and Leshno [2014].

2.1 Discrete Model

A finite set $\Lambda = \{1, 2, \dots, m\}$ 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. For every school $s_i \in S$, let $q_i \in \mathbb{N}_+$ be the *capacity* of school s_i . Let $s_{n+1} \notin S$ denote the outside option. We assume that the outside option has infinite capacity $q_{n+1} = \infty$. Each student $\lambda \in \Lambda$ has strict first round preferences $\succ^\lambda$ and second round preferences $\hat{\succ}^\lambda$ over $S \cup \{s_{n+1}\}$. Each school s_i has weak priorities $\succeq_i^S$ over Λ , which partition the students into priority groups. Equivalently, each student λ has a priority group $p_i^\lambda \in \mathbb{N}$ at school s_i , and students in higher priority groups are preferred.

Definition 1. Preferences $(\succ, \hat{\succ})$ are **consistent** if the second round preferences $\hat{\succ}$ are obtained from the first round preferences $\succ$ via truncation: (1) for every $s_i, s_j \in S$, $s_i \succ s_j$ iff $s_i \hat{\succ} s_j$, and (2) for every $s_i \in S$, $s_{n+1} \succ s_i$ implies $s_{n+1} \hat{\succ} s_i$.

In other words, preferences are consistent if the only change in preferences across the two rounds is in each student's relative ranking of s_{n+1} , which weakly improves from the first round to the second, corresponding to an outside option (possibly) being realized between the two rounds. We say that a student λ is *consistent* if she has consistent preferences $(\succ^\lambda, \hat{\succ}^\lambda)$.

⁹Since we will be considering mechanisms which are strategy-proof in the large, we assume that students report truthfully and do not distinguish between reported preferences and true preferences.

¹⁰We do not establish a formal relationship between the discrete and continuum models, as that is beyond the scope of our paper. Continuum models have been used in a number of papers on school choice, see [Agarwal and Somaini, 2014, Ashlagi and Shi, 2014a, Azevedo and Leshno, 2014]

Each student $\lambda \in \Lambda$ is given a *lottery* number $L(\lambda)$, drawn independently and uniformly from $[0, 1]$. We remark that every ordinal ranking of students can be realized by some cardinal lottery function L , and these random lottery numbers generate a permutation of the students uniformly at random.

An *assignment* specifies a school for each student. For a generic assignment μ , we let $\mu(\lambda)$ denote the school to which student λ is assigned, and $\mu(s_i)$ denote the *set* of students assigned to school s_i . If $\mu(\lambda) = s_i$ then student λ is said to be assigned to school s_i ; if $\mu(\lambda) = s_{n+1}$, student λ is matched with her outside option and is said to be unassigned. An assignment is *feasible* if $|\mu(s_i)| \leq q_i$ for all s_i . In the rest of the paper, the term “assignment” is always used to mean a feasible assignment.

Next, we define what we mean by a *reassigned* student. All else being equal, we want to minimize the number of such students, as student movement is costly both for students and for school administration.

Definition 2. A student $\lambda \in \Lambda$ is a **reassigned student** if she leaves a school in S for another school in S . That is, λ is a reassigned student¹¹ if $\mu(\lambda) \neq \hat{\mu}(\lambda)$ and $\mu(\lambda) \succ^{\lambda} s_{n+1}$.

2.2 Continuum Model

In the continuum model, a *continuum* of students Λ is to be assigned to a set $S = \{s_1, \dots, s_n\}$ of schools. Each student can attend at most one school. The set of students Λ has an associated measure η , i.e., for any (measurable) $A \subseteq \Lambda$, $\eta(A)$ gives the mass of students in A . The outside option is $s_{n+1} \notin S$. As before, the capacities of the schools are $q_1, \dots, q_n \in \mathbb{R}_+$, and $q_{n+1} = \infty$. A set of students of η -measure at most q_i can be assigned to school s_i .

Each student $\lambda \in \Lambda$ has a type θ^λ and a first round lottery number $L(\lambda) \in [0, 1]$, which encode both student preferences and school priorities. A student’s *type* θ is given by the tuple $\theta = (\succ^\theta, \hat{\succ}^\theta, p^\theta)$, where $\succ^\theta$ and $\hat{\succ}^\theta$ are strict preferences over $S \cup \{s_{n+1}\}$, which are, respectively, the student’s first and second round preferences, and p^θ is an n -dimensional priority vector with p_i^θ indicating (the number of) the priority group of student θ at school s_i . Each school s_i has n_i priority groups. We assume that larger priority groups are preferred, and that a student in the k th most preferred priority group at s_i is in priority group $p_i = n_i - k$, so that $p_i \in \{0, 1, \dots, n_i - 1\}$.

Let Θ be the set of all student types. For each $\theta \in \Theta$ let $\zeta(\theta) = \eta(\{\lambda \in \Lambda : \theta^\lambda = \theta\})$ be the measure of all students with type θ . We say that θ is *consistent* if the preferences $(\succ^\theta, \hat{\succ}^\theta)$ are consistent (see Definition 1), and otherwise we say that θ is *inconsistent*. We assume that all students have consistent preferences.

Assumption 1 (Consistent preferences). *If $\zeta(\theta) > 0$ then θ is consistent.*

As in the discrete model, we assume that the first round lottery numbers are i.i.d variables drawn uniformly from $[0, 1]$ and do not depend on preferences. This means that for all $\theta \in \Theta$ and intervals (a, b) with $0 \leq a \leq b \leq 1$, the proportion of students with type θ who have lottery number in (a, b) is equal to the length of the interval (a, b) ,¹²

$$\eta(\{\lambda \in \Lambda : \theta^\lambda = \theta, L(\lambda) \in (a, b)\}) = (b - a)\zeta(\theta).$$

An *assignment* specifies the set of students admitted at each school. For any assignment μ , we let $\mu(\lambda)$ denote the school to which student λ is assigned, and $\mu(s_i)$ denote the *set* of

¹¹Several alternative definitions of reassigned students, such as counting students who are initially in s_{n+1} and end up at a school in S , and/or counting students who no longer find their initial assignment acceptable, could also be considered. We note that our results continue to hold for all these alternative definitions.

¹²This can be justified via an axiomatization of the kind obtained in Al-Najjar [2004].

students assigned to school s_i . We assume that $\mu(s_i)$ is η -measurable and $\eta(\mu(s_i)) \leq q_i$, for all $s_i \in S \cup \{s_{n+1}\}$. We will again let μ denote the first round assignment, and let $\hat{\mu}$ denote the second round assignment.

2.3 Mechanisms

In this section, we formally define the class of mechanisms that we will be considering.

Let the students' first and second round preference reports be denoted by $\succ$ and $\hat{\succ}$ respectively. A *reassignment mechanism* is a function that maps the realization of first round lotteries $L(\lambda)$, first round assignment μ and students' second round reports $\hat{\succ}$ into a reassignment $\hat{\mu}$.¹³ A *two-round mechanism* obtained from a reassignment mechanism M is a two-round mechanism where the first round mechanism is DA with uniform-at-random single tie-breaking, and the second round mechanism is M . We emphasize that we intentionally keep the first round consistent with currently used mechanisms by fixing it to be DA with the same uniform-at-random lottery used for tie-breaking at all schools. It follows that the only freedom afforded the planner is the design of the reassignment mechanism.

We next describe some desirable properties of reassignment mechanisms. Any reassignment that requires taking away a student's initial assignment against her will is impractical. Thus, we require our reassignment to respect first round guarantees:

Definition 3. A reassignment $\hat{\mu}$ **respects guarantees** if every student prefers her second round assignment to her first round assignment, that is $\hat{\mu}(\lambda) \hat{\succ}^\lambda \mu(\lambda)$ for every $\lambda \in \Lambda$.

One of the main reasons for the success of DA in practice is the way it respects priorities: if a student is not assigned to a school she wants, it is because that school filled up with assigned students of higher priority. We require reassignments to respect priorities in this sense:

Definition 4. A reassignment $\hat{\mu}$ **respects priorities (subject to guarantees)** if for every school $s_i \in S$ and student $\lambda \in \Lambda$ such that $s_i \hat{\succ}^\lambda \hat{\mu}(\lambda)$, we have $\eta(\hat{\mu}(s_i)) = q_i$, and if in addition a student λ' satisfies $\hat{\mu}(\lambda') = s_i \neq \mu(\lambda')$ then $\lambda' \succeq_i^S \lambda$.

We are interested in reassignment mechanisms that respect guarantees and priorities. We next define a class of reassignment mechanisms that respect guarantees and priorities. These mechanisms play a prominent role in our analysis, as we will show that, in the continuum model, they are uniquely characterized by respecting guarantees and priorities, and appropriate notions of two-round strategyproofness, efficiency and fairness in the continuum.

We say P is a *permutation* if it is a Lebesgue measure preserving bijection from $[0, 1]$ to $[0, 1]$.

Definition 5 (Permuted Lottery Deferred Acceptance (PLDA) Mechanisms). *Let P be a permutation (that may depend on $\zeta(\cdot)$). 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 (PLDA) associated with P is a function mapping $(L, \mu, \succ, \hat{\succ})$ into a reassignment $\hat{\mu}^P$, which is obtained by running DA on Λ with student preferences $\hat{\succ}$ and S and with school preferences $\hat{\succ}^S$ that are determined as follows. For each s_i :*

- *Student $\lambda \in \Lambda$ for whom $\mu(\lambda) = s_i$ are termed guaranteed at s_i , and all other students are termed non-guaranteed at s_i .*

¹³Here we make the restriction that the second round assignment depends on the first round report 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 will help the mechanism appear more fair from the student perspective.

- Schools prefer all guaranteed students to all non-guaranteed students, that is, for every student $\lambda \in \Lambda$ guaranteed at s_i and student $\lambda' \in \Lambda$ non-guaranteed at s_i , we have $\lambda \succ_i^S \lambda'$.
- Ties within each of the two groups—guaranteed and non-guaranteed—are broken first according to $\succeq_i^S$, and then according to the permuted lottery $P \circ L$ (in favor of the student with the larger permuted lottery number).

Intuitively, these mechanisms involve running deferred acceptance twice. They use single tie-breaking in both rounds, explicitly correlate the lotteries used in the two rounds via P , and modify school priorities in order to guarantee that each student receives a (weakly) better assignment in the second round. Two special cases are worth highlighting. 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 $F(x) = x$. In this paper, we provide strong evidence to support the use of the RLDA mechanism.

We now describe some desirable incentive and efficiency properties for two-round mechanisms in the continuum model, and define PLDA mechanisms in this setting. In the school choice problem, and in the reassignment problem in particular, it is reasonable to assume that students will strategize for all the rounds for which they know their preferences for schools in S . Hence, it is desirable that at every point when a student (with consistent preferences) reports preferences, she receives the best expected utility, conditional on everything that has happened up to that point, by reporting truthfully.¹⁴ Specifically, let *ex ante* expected utility refer to the expected utility from the second round assignment at the beginning of the first round, conditional on first round reports, and let *interim* expected utility refer to the expected utility from the second round assignment at the beginning of the second round, conditional on first round guarantees and second round reports.¹⁵

Definition 6. A two-round mechanism is **two-round strategy-proof** (in the continuum model)¹⁶ if the following hold:

- Knowing first round preferences and the joint distribution of first round assignments and second round preferences, and assuming other students are truthful, no student (with consistent preferences) can obtain a better *ex ante* expected utility from her second round assignment by using any two-round strategy that involves lying in one or both rounds.
- Knowing the specific realization of first round assignments and the conditional distribution of second round preferences, and assuming other students are truthful, no student can obtain a better *interim* expected utility from her second round assignment by lying in the second round.

Note that a mechanism that uses a top-trading cycles like mechanism (which uses the first round allocation as an initial endowment) in the second round will *not* be two-round strategy proof, because students will benefit from manipulating their first round reports to obtain a more popular initial allocation which they can leverage to their advantage in the second round.

In order to be efficient, a reassignment mechanism should not waste any unused seats that are desired by students. A reassignment mechanism is *non-wasteful* if no student is assigned a

¹⁴Here we assume a cardinal utility model underlying student preferences. We make this formal in Section 3.1.

¹⁵We remark that our model makes the most sense with a fixed outside option value in the first round that is then updated in the second round. We may interpret this as students knowing their valuations for each school both inside and outside the system, and assuming that the probability they will receive an outside school that they are not guaranteed is 0, and reporting accordingly in the first round. We note that it is a dominant strategy for them to report in this manner.

¹⁶We remark that PLDA mechanisms are not generally strategy-proof (see Example C in the appendix). However, there is reason to expect that students will not strategize in a sufficiently large discrete market, since PLDAs are two-round strategy-proof in the continuum model, and hence satisfy a “strategy-proofness in the large” condition defined by Azevedo and Budish [2013]

school she prefers less to a school not at capacity, that is for each realization of $\hat{\mu}$, student $\lambda \in \Lambda$ and schools s_i, s_j , if $\hat{\mu}(\lambda) = s_i$ and $s_j \succ^\lambda s_i$, 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 pair of students to be able to improve their utility by swapping probability shares in second round allocations. However, we also require that our reassignment mechanism respects guarantees and priorities (see Definitions 3 and 4). This motivates the following definitions. A *Pareto-improving cycle among reassigned students* is an ordered set of types $(\theta_1, \theta_2, \dots, \theta_m) \in \Theta^m$, sets of students $(\Lambda_1, \Lambda_2, \dots, \Lambda_m), \Lambda_i \subseteq \Lambda$ and schools $(\tilde{s}_1, \tilde{s}_2, \dots, \tilde{s}_m) \in S^m$ such that for all i it holds that $\eta(\Lambda_i) > 0$, $\tilde{s}_{i+1} \succ^{\theta_i} \tilde{s}_i$ (where we define $\tilde{s}_{m+1} = \tilde{s}_1$), and for all $\lambda \in \Lambda_i$ it is the case that $\theta^\lambda = \theta_i, \hat{\mu}(\lambda) = s_i, \mu(\lambda) \neq s_i$. We say that a Pareto-improving cycle among reassigned students *respects priorities* if for all i and $\lambda \in \Lambda_i, \hat{\lambda} \in \Lambda_{i+1}$ it holds that $p_{i+1}^\lambda \geq p_{i+1}^{\hat{\lambda}}$ (where we define $\Lambda_{m+1} = \Lambda_1$ and $p_{m+1} = p_1$).

Definition 7. A two-round mechanism is **constrained Pareto efficient among reassigned students** if there are no Pareto-improving cycles among reassigned students that respect priorities.

We formally define permuted lottery deferred acceptance mechanisms in the continuum model in terms of cutoffs, in the style of Azevedo and Leshno [2014], as follows.

Definition 8 (PLDA Mechanisms in the Continuum). Let P be a permutation (that may depend on $\zeta(\cdot)$). The first round assignment μ is defined by a vector of cutoffs $\mathbf{C} \in \mathbb{R}_+^{n+1}$, where a student λ is given a first round score $r_i^\lambda = p_i^\lambda + L(\lambda)$ at school s_i and assigned to her favorite school among those where her first round score exceeds the cutoff:

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

In the second round, for each student λ with priority vector $p^\lambda = p$, lottery number $L(\lambda) = l$ and first round assignment $\mu(\lambda) = s_i$, we associate a second round score of $\hat{r}_i^\lambda = n_i + P(l)$ at school s_i , and a second round score of $\hat{r}_j^\lambda = p_j + P(l)$ at all other schools $s_j, j \neq i$. The second round assignment $\hat{\mu}^P$ is defined via a vector of cutoffs $\hat{\mathbf{C}}^P \in \mathbb{R}_+^{n+1}$, where a student λ is assigned to her favorite school among those where her second round score weakly exceeds the cutoff:

$$\hat{\mu}^P(\lambda) = \max_{\hat{\succ}^\lambda}(\{s_i \in S \cup \{s_{n+1}\} : \hat{r}_i^\lambda \geq \hat{C}_i^P\}).$$

The validity of such a definition in describing DA mechanisms follows from the arguments in [Azevedo and Leshno, 2014], and it is easy to verify that our model satisfies the technical conditions required in that paper.

PLDA mechanisms are an attractive class of two-round mechanisms, as they satisfy all our desired incentive and efficiency properties.

Proposition 1. Suppose student preferences are consistent. Then PLDA mechanisms respect guarantees and priorities, and are two-round strategy-proof (in the continuum model), non-wasteful, and constrained Pareto efficient among reassigned students.

Proposition 1 demonstrates that PLDA mechanisms have desirable properties. They are two-round strategy-proof in the continuum model, so students have an incentive to truthfully report their preferences when the number of students is large. They are also efficient in that they are non-wasteful and constrained Pareto efficient among reassigned students.

We will further show (Theorem 4) that in a setting without priorities, the PLDA mechanisms are in a sense the only mechanisms that satisfy all these properties. Specifically, suppose we

additionally require that our reassignment mechanism is *anonymous*, in that students with the same first round assignments and second round preference reports have the same distributions over reassignments, and is *averaging*, in that the aggregate proportion of students of a given type assigned to each pair of schools in the two rounds is deterministic. Then every mechanism that satisfies these properties is a PLDA mechanism. We make this formal in Section 3.1.

3 Main Results

We have introduced the PLDA mechanisms as a class of mechanisms with desirable incentive and efficiency properties. We will show that the defining characteristic of a mechanism in this class –the permutation of lotteries between the two rounds– can be chosen to achieve desired operational goals.

We first provide a simple and intuitive *order condition* under which all PLDA mechanisms give the same ex ante allocative efficiency. When the primitives of the market satisfy the order condition, it is possible to pursue secondary operational goals without sacrificing allocative efficiency. Then, in the context of reassigning school seats before 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 PLDA mechanisms. In Section 3.1 we provide axiomatic justification for PLDAs (in the case of no school priorities), showing that PLDAs are the two-round strategyproof mechanisms that are Pareto efficient among reassigned students and satisfy certain natural fairness properties.

In Section 5 we demonstrate that our theoretical results hold empirically even in settings where the order condition does not hold. In simulations using data from the NYC high school system, the reverse lottery significantly reduces the number of reassigned students (relative to reusing the same lotteries as in the initial round, or other natural permutations), and does not significantly affect allocative efficiency.

Our results suggest that RLDA is a good choice of mechanism when the primary goal is to provide an efficient second round assignment (in terms of the allocation produced) while minimizing the number of reassigned students. In Section 6 we discuss how the choice of lottery can be used to achieve other operational goals, such as maximizing the number of students with improved assignments. Our results indicate that it is possible to do so without sacrificing allocative efficiency.

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

Definition 9. *The **order condition** holds on a set of primitives (S, q, p, Λ, η) if for every priority class p , the first and second round school cutoffs under RLDA within that priority class are in the same order. That is, for all $s_i, s_j \in S \cup \{s_{n+1}\}$,*

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

We emphasize that the order condition is a condition on the market primitives, namely, school priorities and capacities and student preferences. 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 relative overdemand for the schools. One important setting where the order condition holds is the case of uniform dropouts and a single priority type. In this setting, each student either remains in the system and retains her first round preferences in the second round, or drops out of the system entirely. If dropouts occur in an i.i.d. fashion, the order condition is satisfied (for arbitrary school priorities and capacities and initial student preferences).¹⁷ We provide direct

¹⁷Formally, this is the case where there exists $\rho \in [0, 1]$ s.t. for every strict preference $\succ$ over schools, it

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 them in the general setting.

Next, we define type equivalence of mechanisms. In words, this refers to equivalence between mechanisms in terms of the masses of different student types θ assigned to each school in the second round allocation.

Definition 10. *Two allocations $\hat{\mu}$ and $\hat{\mu}'$ are said to be **type equivalent** if for all types $\theta \in \Theta$ and schools s_i ,*

$$\eta(\{\lambda \in \Lambda : \theta^\lambda = \theta, \hat{\mu}(\lambda) = s_i\}) = \eta(\{\lambda \in \Lambda : \theta^\lambda = \theta, \hat{\mu}'(\lambda) = s_i\}).^{18}$$

We remark that type equivalence depends only on the second round allocation, while reassignment measures the difference between the first and second round allocations.

We are now ready to state the main results of this section.

Theorem 1 (Order condition implies type equivalence). *If the order condition (Definition 9) holds, all PLDA mechanisms produce type equivalent allocations.*

Thus, if the order condition holds, the measure of each student type $\theta \in \Theta$ assigned to each school in the second round is independent of the permutation chosen for the PLDA.

Theorem 2 (Reverse lottery minimizes reassignment). *If all PLDA mechanisms produce type equivalent allocations, then RLDA minimizes the measure of reassigned students among PLDA mechanisms.*

Our results present a strong case for using the 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, RLDA is unequivocally optimal in the class of PLDA mechanisms, since all PLDA mechanisms give type-equivalent allocations¹⁹ and RLDA minimizes the number of reassigned students. We remark that running RLDA provides a simple way to check for whether the order condition holds for a given set of primitives.²⁰

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. *Consider a setting with 2 schools, each with a single priority group. School s_1 has lower capacity and is initially more overdemand. Student preferences are such that when all the students who only want s_2 drop out, s_1 is still more overdemand and the order condition holds, and when all the students who only want s_1 drop out, s_2 now becomes more overdemand under RLDA and the order condition does not hold.*

Specifically, 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

holds that $\zeta(\{\theta = (\succ^\theta, \hat{\succ}^\theta, \mathbf{1}) \in \Theta : \succ^\theta = \succ, \hat{\succ}^\theta = s_{n+1} \succ \dots\}) = \rho \zeta(\{\theta = (\succ^\theta, \hat{\succ}^\theta, \mathbf{1}) \in \Theta : \succ^\theta = \succ\})$, and $\zeta(\{\theta = (\succ^\theta, \hat{\succ}^\theta, \mathbf{1}) \in \Theta : \succ^\theta = \succ, \hat{\succ}^\theta = \succ\}) = (1 - \rho) \zeta(\{\theta = (\succ^\theta, \hat{\succ}^\theta, \mathbf{1}) \in \Theta : \succ^\theta = \succ\})$. We note that this is essentially the case where dropouts are i.i.d. with probability p . We remark that there are well-known technical measurability issue w.r.t. a continuum of random variables, but it should be noted that this issue can be handled—see, for example, Al-Najjar [2004].

¹⁸We remark that the type equivalence condition is well-defined in the space of interest. Specifically, although for general random mechanisms these measures are random variables, in the case of PLDA mechanisms, these measures are a deterministic function of priorities and preferences, and the equality is well-defined.

¹⁹A reasonable utility model in the continuum would yield that type equivalence implies welfare equivalence.

²⁰We are not suggesting that the mechanism should involve checking the order condition and then using RLDA only if this condition is satisfied (based on the guarantee in Theorems 1 and 2). However, one could, in principle, check the order condition on historical data and accordingly decide whether to use the RLDA mechanism or not. See Section 5 for discussion on considerations if the order condition is not exactly satisfied.

school s_i acceptable, and $\theta_{i,j}$ denote the type that finds both acceptable and prefers s_i to s_j . As we will be considering only students who either leave the system completely or keep the same preferences, we will also let these denote student types with the same preferences in the second round. If we run DA with single tie-breaking, the first round cutoffs are $(C_1, C_2) = (\frac{3}{4}, \frac{1}{2})$.

Suppose that all type θ_2 students leave the system, and no other students receive an outside option. This frees up 2 units at s_2 . Under RLDA, the second round cutoffs are $(\hat{C}_1^R, \hat{C}_2^R) = (1, \frac{3}{4})$. In this case, the order condition holds and FLDA and RLDA are type equivalent. It is simple to verify that both FLDA and RLDA assign

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

where the vector denotes the measure of students of type $(\theta_1, \theta_{12}, \theta_{21})$ assigned to the school.

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

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

where the vector denotes the measure of students of type $(\theta_{12}, \theta_{21}, \theta_2)$ assigned to the school. We note that in this case, FLDA is Pareto efficient, whereas RLDA is not.

We end this section by giving a little intuition as to why our results are true. A key conceptual insight is that we can simplify the analysis by shifting away from student assignments, which depend on student preferences, and considering instead the options that a student is allowed to choose from, which are independent of preferences. Specifically, if we define the *affordable sets* 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, 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 11. We say that a PLDA with permutation P satisfies the **local order condition** on a set of primitives (S, q, p, Λ, η) if, for every priority class p , the first and second round school cutoffs within that priority class are in the same order under this PLDA. That is, for all $s_i, s_j \in S \cup \{s_{n+1}\}$,

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

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

1. (Consistency across rounds) For every permutation P , the local order condition holds for PLDA with permutation P on (S, q, p, Λ, η) ; and
2. (Consistency across permutations) For every priority class p , for all pairs of permutations P, P' and schools $s_i, s_j \in S \cup \{s_{n+1}\}$, it holds that $\hat{C}_{p,i}^P > \hat{C}_{p,j}^P \Rightarrow \hat{C}_{p,i}^{P'} \geq \hat{C}_{p,j}^{P'}$.

In other words, the global order condition requires that not just RLDA but all PLDA mechanisms result in a school cutoff ordering that is preserved from the first to the second round, and




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 top and bottom right, we show the second round assignments under FLDA and RLDA when type θ_2 students (who only want s_2) drop out, and when type θ_1 students (who only want s_1) drop out. Students toward the left have larger lottery numbers. The patterned boxes above each column of students indicate the affordable sets for students in that column. When students who only want s_2 drop out, the order condition holds, and FLDA and RLDA are type equivalent. When students who only want 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.

further that first-round ties in cutoffs are broken consistently across PLDAs in the second round. Surprisingly, this global requirement is guaranteed to hold if the order condition (Definition 9) holds, i.e., if RLDA satisfies the local order condition.

Theorem 3. *The order condition holds on a set of primitives (S, q, p, Λ, η) if and only if the global order condition holds on (S, q, p, Λ, η) .*

We provide some intuition as to why Theorem 3 holds 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 unaligned. 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 Theorem 3. Fix a mechanism and suppose that the first and second round cutoffs are in the same order. Then each student λ 's affordable set is of the form $X_i = \{s_i, s_{i+1}, \dots, s_n\}$ for some $i = i(\lambda)$, where schools are indexed in decreasing order of their cutoffs for the relevant priority group p^{θ^λ} , and the probability that a student receives some affordable set is independent of her preferences. Moreover, since affordable sets are nested $X_1 \supseteq X_2 \supseteq \dots \supseteq X_n$, and since the lottery order is independent of student types, aggregate demand 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.

3.1 Characterization

In this section, we consider a setting where each school has exactly one priority group. We show that the class of PLDA mechanisms is equivalent to the class of two-round mechanisms that respect guarantees and priorities and are two-round strategy-proof, non-wasteful, Pareto efficient among reassigned students, and satisfy a few natural fairness conditions. Note that, in the case with exactly one priority type, the first round corresponds to the random serial dictatorship (RSD) mechanism of Abdulkadiroğlu and Sönmez [1998], where the (random) order of students is given by the order of tie-breaking. We will only consider mechanisms that are *non-atomic*, that is, any single student changing her preferences has no effect on the assignment probabilities of other students.

In order to let students trade off between probabilities of assignment, we will need to assume some underlying cardinal utilities. Let $\Upsilon \subseteq \mathbb{R}^{n+1} \times \mathbb{R}^{n+1}$ be the set of pairs of first and second round cardinal utility vectors $(u, \hat{u})$ that correspond to consistent preferences, where $u_i, \hat{u}_i$ are the student's utility from being assigned to s_i in the first and second round respectively. Formally,²¹

$$\Upsilon = \{(u, \hat{u}) \in \mathbb{R}^{n+1} \times \mathbb{R}^{n+1} : u_i \geq u_j \Leftrightarrow \hat{u}_i \geq \hat{u}_j \ \forall i, j \leq n, \text{ and } u_i \leq u_{n+1} \Rightarrow \hat{u}_i \leq \hat{u}_{n+1} \ \forall i \leq n\}.$$

We assume that students' consistent cardinal utilities have full support.

Assumption 2 (Full support). *Student consistent cardinal utilities have full support. That is, for all sets $U \subseteq \Upsilon$ with positive Lebesgue measure, there exists a positive measure of students whose first and second round cardinal utilities pairs are in U .*

We note that this is a strong assumption, and is stated in this manner for simplicity and brevity.²²

We have already shown that PLDA mechanisms satisfy a number of desirable properties. We restate them here for completeness. A reassignment mechanism $\hat{\mu}$ *respects guarantees* if every student prefers her second round assignment to her first round assignment, that is $\hat{\mu}(\lambda) \succeq^\lambda \mu(\lambda)$ for every $\lambda \in \Lambda$. A reassignment mechanism is *non-wasteful* if no student is assigned a school she prefers less to a school not at capacity, that is for each realization of $\hat{\mu}$, student $\lambda \in \Lambda$ and schools s_i, s_j , if $\hat{\mu}(\lambda) = s_i$ and $s_j \succ^\lambda s_i$, then $\eta(\hat{\mu}(s_j)) = q_j$. A two-round mechanism is *two-round strategy-proof* (in the continuum model) if: knowing the joint distribution of first round preferences, first round assignments and second round preferences, and assuming other students are truthful, no student (with consistent preferences) can obtain a better ex ante expected utility from her second round assignment by using any two-round strategy that involves lying in one or both rounds; and knowing the specific realization of first round assignments and the conditional distribution of second round preferences, and assuming other students are truthful, no student can obtain a better interim expected utility from her second round assignment by lying in the second round.

We remark that we require two-round strategy-proofness only for students whose true preference type is consistent. This is because reversals in student preferences 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. Moreover, it may be reasonable to assume that students who are sophisticated enough to strategize about misreporting in the first round to affect the guarantee structure in the second round will also know their

²¹We note that Υ is isomorphic to $\mathbb{R}^{n+2}$ and inherits the Lebesgue measure via the isomorphism.

²²We will only need support on certain specific subsets of cardinal utilities in order to prove our characterization result. Moreover, if we let the measure of students of any subset of types tend to 0, our characterization result will still hold.

second round preferences over schools in S at the beginning of the first round, and hence will have consistent preferences.²³

We also specify the appropriate notion of efficiency in a setting without priorities. A *Pareto-improving cycle among reassigned students* is an ordered set of types $(\theta_1, \theta_2, \dots, \theta_m) \in \Theta^m$, sets of students $(\Lambda_1, \Lambda_2, \dots, \Lambda_m)$, $\Lambda_i \subseteq \Lambda$ and schools $(\tilde{s}_1, \tilde{s}_2, \dots, \tilde{s}_m) \in S^m$ such that for all i it holds that $\eta(\Lambda_i) > 0$, $\tilde{s}_{i+1} \succ^{\theta_i} \tilde{s}_i$ (where we define $\tilde{s}_{m+1} = \tilde{s}_1$), and for all $\lambda \in \Lambda_i$ it is the case that $\theta^\lambda = \theta_i$, $\hat{\mu}(\lambda) = s_i$, $\mu(\lambda) \neq s_i$.

Definition 12. A two-round mechanism is *Pareto efficient among reassigned students*²⁴ if there are no Pareto-improving cycles among reassigned students.

We now additionally define some desirable fairness properties. We want our mechanism to treat students equally. A two-round mechanism is *anonymous* if it entirely ignores student identities. In particular, students with the same first round assignment and first and second round preference reports have the same distribution over second round assignments. A two-round mechanism where the first round is RSD satisfies the *averaging* axiom, if for a fixed first round allocation μ , every realization of the second round allocation $\hat{\mu}$ leads to the same proportion of students with type θ being assigned to a pair of schools (s, s') in the two rounds. That is, for all θ, s, s' , there exists a constant $c_{\theta, s, s'}$ such that $\eta(\{\lambda \in \Lambda : \theta^\lambda = \theta, \mu(\lambda) = s, \hat{\mu}(\lambda) = s'\}) = c_{\theta, s, s'}$.

Our characterization result is the following.

Theorem 4. Suppose student preferences are consistent and student cardinal utilities have full support (Assumption 2). A non-atomic two-round assignment mechanism where the first round is DA with uniform-at-random single tie-breaking²⁵ respects guarantees and is:

1. non-wasteful;
2. two-round strategy-proof;
3. Pareto efficient among reassigned students;
4. anonymous; and
5. averaging;

if and only if the second round assignment is given by PLDA, or obtainable as a stable matching under the same second round school preferences as in PLDA. (Here the permutation P can depend on the measure of student preference types $\zeta(\cdot)$.)

Our result mirrors similar large market cutoff characterizations for single-round mechanisms by Liu and Pycia [2016] and Ashlagi and Shi [2014a], which show, in settings with a single and multiple priority types respectively, that a mechanism is non-atomic, strategyproof, symmetric, and efficient (in each priority class) if and only if it is lottery-plus-cutoff. Our contribution is that we consider a two-round setting, where we use an affordable set argument, and the fact that the mechanism respects guarantees, to isolate the second round from the first. We then employ arguments similar to those used in these previous results to show that the first and second round mechanisms can individually be characterized using lottery-plus-cutoff mechanisms.

²³One obvious objection is that students may also obtain extra utility from staying at a school between rounds, or equivalently have a disutility for moving, creating inconsistent preferences where the school they are assigned in the first round becomes preferred to previously more desirable schools. We remark that Theorem 4 extends to the case when we consider students whose preferences incorporate an additional utility for staying put, provided that the utility is the same at every school for a given student, or satisfies a similar non-crossing property.

²⁴An alternative would be to require Pareto efficiency including all students. This would lead to a smaller class of mechanisms that includes the forward lottery DA (FLDA) mechanism, but typically does *not* include the reverse lottery DA (RLDA) mechanism (see the end of Section 2 for definitions of FLDA and RLDA).

²⁵This is also random serial dictatorship.

The main focus of our result is the effect of cross-round constraints. First, two-round strategyproofness, together with respecting guarantees, constrains the second round to be DA, with each student assigned in the first round given a guarantee at the school she was assigned in the first round. In addition, we show that two-round strategyproofness and anonymity together allow the cutoffs in the two rounds to be explicitly related, in a way that is independent of student preferences.

The intuition behind the proof again uses affordable sets, which are analogous to cutoffs. We first construct overdemand orderings in both rounds, as in Ashlagi and Shi [2014a]. We next note that the affordable sets in each round are given by the n sets of least overdemanded schools according to the overdemand ordering in that round, which gives some structure to the possible joint distribution over first and second round affordable sets. The bulk of the proof is dedicated to showing a technical lemma (Lemma 2), which states that two-round strategyproofness and anonymity together imply that two students of different types face the same joint distribution of first and second round affordable sets. The formal proof can be found in the appendix.

4 A Special Case: Uniform Dropouts

In this section, we consider the special case of our model in which there is exactly one priority group at each school, students leave the system uniformly at random with some fixed probability ρ , and the students who remain in the system retain their preferences. For this special case, we give a direct proof that the global order condition holds, and show that all PLDA mechanisms give type equivalent allocations. These proofs give the interested reader a taste of the analysis and the proof techniques used in the paper, but in a simpler and more transparent setting. This section may be skipped at a first reading without loss of continuity.

In the uniform dropouts model, each student drops out of the system with probability ρ . This can be interpreted as students leaving the city for reasons that are independent of the school choice system. In this case, the second round problem can 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 factor $1 - \rho$, the residual capacity of each school is ρq_i , and the measure of students of each type θ who are still in the system is scaled down by factor $1 - \rho$. Thus, the relative overdemand remains the same and schools fill in the same order regardless of the choice of permutation.

Throughout 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 $\theta = (\succ^\theta, \hat{\succ}^\theta)$.

Formally, we define uniform dropouts with probability ρ as follows. For every strict preference $\succ$ over schools, students with first round preferences $\succ$ with probability ρ find the outside option s_{n+1} the most attractive in the second round, and with probability $1 - \rho$ retain the same preferences in the second round,

$$\begin{aligned} \zeta(\{\theta = (\succ^\theta, \hat{\succ}^\theta) \in \Theta : \succ^\theta = \succ, \hat{\succ}^\theta = s_{n+1} \succ \dots\}) &= \rho \zeta(\{\theta = (\succ^\theta, \hat{\succ}^\theta) \in \Theta : \succ^\theta = \succ\}), \\ \zeta(\{\theta = (\succ^\theta, \hat{\succ}^\theta) \in \Theta : \succ^\theta = \succ, \hat{\succ}^\theta = \succ\}) &= (1 - \rho) \zeta(\{\theta = (\succ^\theta, \hat{\succ}^\theta) \in \Theta : \succ^\theta = \succ\}). \end{aligned}$$

We show first that the order condition (Definition 9) holds, and that Theorem 3 holds in the setting with uniform dropouts. Specifically, we show that the local order condition holds for the reverse lottery, and that the global order condition holds. The proof of Theorem 3 in the general setting has a similar flavor, and can be found in the appendix.

Theorem 5. *In the case of uniform dropouts, the global order condition holds.*

Proof of Theorem 5. We start with some intuition and a proof outline before providing the formal proof. We assume without loss of generality that all schools reach capacity in the second round of PLDA.²⁶ Each student s has an *affordable set*, which is the set of schools for which she meets either the first or second round cutoffs. Suppose that the first and second round cutoffs are in the same order. Then each student λ 's affordable set is of the form $X_i = \{s_i, s_{i+1}, \dots, s_n\}$ for some $i = i(\lambda)$, and the proportions of students with each affordable set is independent of student preferences. Hence the measure of students assigned to a school is uniquely identified by the proportion of students with affordable set X_i for each i .

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. It follows that the constructed cutoffs are precisely the PLDA(P) cutoffs.

Let the first rounds cutoffs be $C_1, C_2, \dots, C_n$, where without loss of generality we index the schools such that $C_1 \geq C_2 \geq \dots \geq C_n$.

(1) 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 we want the proportion of students with an affordable set containing s_j to be $\frac{1}{1-\rho}$ times the original proportion.

(2) We translate this into cutoffs in the following way. 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 scoring functions under permutation P if the first and second round cutoffs are C_i and x respectively. Notice that $f_i(x)$ is decreasing 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 defined by the equation

$$f_i(\tilde{C}_i^P) = \frac{1}{1-\rho}(1 - C_i),$$

and we let $\tilde{C}_i^P = 0$ if $C_i < \rho$.

(3) We now show that the cutoffs $\tilde{\mathbf{C}}$ are in the right order. Suppose $i < j$. 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, we have

$$\begin{aligned} \tilde{C}_j^P &\leq \tilde{C}_i^P \\ \Leftrightarrow f_j(\tilde{C}_j^P) - (1 - C_j) &\geq f_j(\tilde{C}_i^P) - (1 - C_j) \\ \Leftrightarrow \frac{\rho}{1-\rho}(1 - C_j) &\geq \left| \left\{ l : l < C_j, P(l) \geq \tilde{C}_i^P \right\} \right|. \end{aligned}$$

Now

$$\left| \left\{ l : l < C_j, P(l) \geq \tilde{C}_i^P \right\} \right| \leq \left| \left\{ l : l < C_i, P(l) \geq \tilde{C}_i^P \right\} \right| = \frac{\rho}{1-\rho}(1 - C_i) \leq \frac{\rho}{1-\rho}(1 - C_j),$$

where the inequalities follow from the fact that $C_i \geq C_j$, and so $\tilde{C}_i^P \geq \tilde{C}_j^P$ as required.

(4) We now show that $\tilde{\mathbf{C}}^P$ is the set of market-clearing DA cutoffs for the second round of PLDA with permutation P . This step essentially verifies that affordable sets uniquely determine

²⁶Suppose that not all schools reach capacity in the second round of PLDA. Then we may add a mass of auxiliary students with the same type distribution as the existing students, and the result follows by considering all PLDA mechanisms on this amended economy where the first round lottery L provides all auxiliary students with worse lottery numbers than those of the original students in the first round, and the permutation P again provides all the auxiliary students with worse lottery numbers in the second round.

demand. Note that $\gamma_i = C_i - C_{i+1}$ is the proportion of students whose first round affordable set is X_i . 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 the dropouts occur.

Now $f_i(\tilde{C}_i^P)$ is the proportion of students whose second round affordable set contains s_i , and since $C_1 \geq C_2 \geq \dots \geq C_n$ and $\tilde{C}_1^P \geq \tilde{C}_2^P \geq \dots \geq \tilde{C}_n^P$, it follows that the affordable sets are nested. Hence the proportion of students (remaining after dropouts) whose second round affordable set is X_i is given by

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

For each student type $\theta = (\succ, \succ)$ 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 for all i , a set of students of measure

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

choose to go to school s_i in the second round under the second round cutoffs $\tilde{\mathbf{C}}^P$, which is the same as the measure of the set of students who choose to go to school s_i in the first round under first round cutoffs $\mathbf{C}$. Since $\mathbf{C}$ are market-clearing cutoffs, it follows that $\tilde{\mathbf{C}}^P$ are too.

Hence in PLDA(P), the second round cutoffs are exactly the constructed cutoffs $\tilde{\mathbf{C}}^P$ and they satisfy $\tilde{C}_1^P \geq \dots \geq \tilde{C}_n^P$. In other words, the local order condition is satisfied for all P , with schools indexed such that $C_1 \geq C_2 \geq \dots \geq C_n$, and so the global order condition holds. $\square$

Remark. The general proof of Theorem 3 uses the cutoffs for RLDA in steps (1) and (2) above to guess the proportions of students who receive an affordable set that contains s_j , and requires careful accounting for each student priority type. However, the general structure for the proof is similar, and the tools used are straightforward generalizations of the ones used in the proof above.

We next show that Theorem 1 holds with uniform dropouts. Specifically, we show that all PLDA mechanisms give type equivalent allocations.

Proposition 2. *In the case of uniform dropouts, all PLDA mechanisms produce type equivalent allocations.*

Proof of Proposition 2. As in the proof of Theorem 5, we assume without loss of generality that all schools reach capacity in the second round of PLDA. We showed in the proof of Theorem 5 that for all i , a set of students of measure

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

choose to go to school s_i in the second round under the second round cutoffs $\hat{\mathbf{C}}^P$, where

$$\gamma_i^P = f_{i+1}(\hat{C}_{i+1}^P) - f_i(\hat{C}_i^P) = \frac{C_i - C_{i+1}}{1 - \rho} = \frac{\gamma_i}{1 - \rho}$$

is the proportion of students with affordable set X_i in the second round with permutation P and cutoffs $\hat{\mathbf{C}}^P$.

It follows from this analysis that all PLDAs are “type equivalent” to the first round assignment, in the following sense. For each preference order $\succ$, let $\hat{\succ}$ be the preferences obtained

from $\succ$ by making the outside option the most desirable, $s_n \tilde{\succ} \dots$. Let $T = \{(\succ, \succ), (\succ, \tilde{\succ})\}$ be a set containing a pair of consistent types that either retain the same preferences or drop out. Then for all schools s_i , it holds that

$$\eta\left(\left\{\lambda \in \Lambda : \left(\succ^\lambda, \tilde{\succ}^\lambda\right) \in T, \hat{\mu}^P(\lambda) = s_i\right\}\right) = \eta\left(\left\{\lambda \in \Lambda : \left(\succ^\lambda, \tilde{\succ}^\lambda\right) \in T, \mu(\lambda) = s_i\right\}\right).$$

Since, under uniform dropouts, exactly one type from each set T remains in the system in round 2, and this covers all student types that remain in the system, it follows that $\hat{\mu}^P$ is type equivalent to $\hat{\mu}^{P'}$ for all permutations P, P' . $\square$

5 Empirical Analysis

In this section, we use data from New York City’s (NYC) school choice system to simulate and evaluate the performance of 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 very similar for different permutations P . At the same time, the difference in the number of reassigned students is significant and is minimized under RLDA.

5.1 The Data and Simulations

We use data from the high school admissions process in NYC for the academic years 2004-05, 2005-06 and 2006-07, as follows:²⁷

1. First round preferences: In our simulation, we take the first round preferences $\succ$ of every student to be the preferences they submitted in the main round of admissions. The algorithm used in practice is essentially strategy-proof (see Abdulkadirolu et al. [2005a]),²⁸ so it is reasonable to assume that reported preferences are true preferences.²⁹
2. Second round preferences: In our simulation, students either drop out from 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).³⁰
3. School capacities: Each school’s capacity is set equal to the number of students assigned to it in the data.³¹ This is only a lower bound on the actual capacity, but allows us to compute the true final allocation under PLDA since under our model, the occupancy of schools with vacant seats only decreases in the second round.
4. School priorities over students are obtained directly from the data.³²

²⁷We performed an initial cleanup of the data, such as removing preference entries which did not correspond to an existing school code.

²⁸The 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.

²⁹However, 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], Narita [2016]

³⁰For a minority of the students (9.2% – 10.45%), their 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 rate of dropouts for the rest of the students (8.9% – 9.2%).

³¹As per the final assignment produced by centralized allocation.

³²Unlike in the theoretical analysis, where we assumed no priorities, we take them into consideration here. We obtained similar results to the ones described below in simulations with no school priorities.

α	Reassignments (number)	Reassignments (%)	$k = 1$ (%)	$k = 2$ (%)	$k = 3$ (%)	Unassigned (%)
FLDA: ∞	7235	8.84	50.27	13.21	7.53	4.79
7.39	7035	8.59	50.26	13.22	7.53	4.79
2.72	6551	8.00	50.21	13.24	7.55	4.79
1.00	5830	7.12	50.10	13.31	7.57	4.78
0.37	5240	6.40	49.98	13.38	7.61	4.76
0.00	4792	5.85	49.86	13.45	7.64	4.76
-0.37	4336	5.30	49.75	13.49	7.68	4.74
-1.00	3751	4.58	49.59	13.55	7.72	4.73
-2.72	3253	3.97	49.47	13.57	7.76	4.71
-7.39	3106	3.79	49.44	13.57	7.76	4.70
RLDA: $-\infty$	3079	3.76	49.43	13.56	7.76	4.70

Table 1: Simulation Results, 2004-2005 NYC High School admissions with priorities.

We show the mean number of students being reassigned, the mean percentage of students being reassigned, and the mean percentage getting their k -th choice or remaining unassigned, averaged across 50 realizations for each value of α . All percentages are out of the total number of students, including those that drop out. The data contained 81,884 students and 653 schools. The percentage of students who dropped out was 9.22%. Variation across realizations in number of reassignments was approximately 100 students.

We consider the following family of 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. The second round score 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. Note that $\alpha = -\infty$ corresponds to RLDA and $\alpha = \infty$ corresponds to FLDA. 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 remark that two different iterations using the same α do not generally give the same PLDA, since in a PLDA mechanism the permutation used is a deterministic function of the input, while $\tilde{L}(\lambda)$ is random. We quote averages across simulations.)

5.2 Results

The results of our computational experiments based on 2004-05 NYC high school admissions data appear in Figure 4 and Table 1.³³ Allocative efficiency appears not to vary much across values of α ; the number of students receiving their k -th choice for each $1 \leq k \leq 12$, as well as the number of unassigned students, vary by less than 1% of the total number of students. (Larger values of α give more students their first choice, but only slightly.) We further find that for most students, the likelihood of getting one of their top k choices under FLDA and under RLDA are very close to each other. (Specifically, for 87% of students, these likelihoods differ empirically by less than 3% for all k . Here, the bound of 3% was chosen to suppress statistical variation in our simulations, that involved 19,800 trials each for FLDA and RLDA. FLDA and RLDA were compared since they represent the two extreme PLDAs. The same divergence metric was under 10% for 98.2% of students.) This is consistent with what we would expect based on our theoretical finding of type equivalence (Theorem 1) of the final allocation under different PLDA mechanisms.

Figure 4 shows that the mean number of reassignments is minimized at $\alpha = -\infty$ (RLDA)

³³The results for 2005-06 and 2006-07 were similar.




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

and increases with α , which is consistent with our theoretical result in Theorem 2. The mean number of reassignments is as large as 7,200 under FLDA compared to just 3,100 under RLDA. Overall our findings suggest that RLDA would be the best choice of mechanism in the family of PLDA mechanisms considered.

6 Discussion

We have shown that the reverse lottery deferred acceptance mechanism (RLDA) is an attractive choice when the objective is to minimize the number of reassigned students. It attains similar allocative efficiency to all PLDA mechanisms when the order condition is satisfied (Theorem 1) and minimizes reassigned students (Theorem 2). In addition, RLDA also has the nice property of being equitable in an observable way. Specifically, under RLDA, 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, the MIT dorm Random Hall uses a mechanism for assigning rooms that resembles the reverse lottery mechanism we have proposed (see Ran). Freshmen get their rooms through a serial dictatorship mechanism. At the end of the year (after seniors leave), they can claim the rooms vacated by the seniors, and each student is guaranteed to at least keep her current room. The initial lottery numbers (from their first match) are reversed.³⁴

We also remark that our results suggest that PLDA mechanisms are an attractive class of mechanisms in more general settings, and the choice of mechanism within this class 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, then the FLDA mechanism that simply runs DA again would optimize over this objective. Moreover, 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.

We axiomatically justified the class of PLDA mechanisms in settings where schools do not

³⁴The 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.

have priorities (Theorem 4). A natural question is whether a similar result holds for a model with school priorities. We find that in a model with priorities, natural extensions of our axioms continue to describe PLDA mechanisms, but also include undesirable generalizations of PLDA mechanisms. Specifically, suppose we add an axiom requiring that for each school s , the probability that a student who reports a top choice of s then receives it in the first or second round is independent of their priority at other schools. This new set of axioms describes a class of mechanisms that includes the PLDA mechanisms, but there also exists an example market and a mechanism satisfying this new set of axioms such that the joint distribution over the two rounds of allocations does not match any PLDA. We believe that this new set of axioms describes a generalization of PLDA where the permutation of a student’s lottery number depends on their priority type. However, characterizing the class of mechanisms satisfying this natural set of 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.

It is natural to ask what implications our results have for finite markets. Azevedo and Leshno [2014] 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 the continuum PLDA mechanisms are a good approximation for the discrete PLDA mechanisms on sufficiently large discrete economies, and that our theoretical results should approximately hold on large discrete economies. As an example, we provide a heuristic argument for why PLDA mechanisms are strategyproof in the large. By definition, PLDA mechanisms satisfy the efficiency and anonymity requirements in finite markets as well. We believe that PLDA mechanisms are approximately strategyproof in large finite markets. In the second round it is clearly a dominant strategy to be truthful, and intuitively, for student a to benefit from a first-round manipulation, his report should affect the second-round cutoffs in a manner that gives him a second-round allocation he would not have received otherwise. If the market is large enough, the cutoffs should converge to their limiting values, and the probability that student a could benefit from such a manipulation would be negligible. A similar argument suggests that an approximate version of our characterization result (Theorem 4) should hold for finite markets with no priorities, as PLDA mechanisms satisfy an approximate version of the averaging axiom in large finite markets. Our results concerning type equivalence (Theorem 1), RLDA minimizing transfers (Theorem 2), and our sufficiency check for the order condition (Theorem 3) should also be approximately valid in the large market limit.³⁵

Another natural question is how to deal with inconsistent student preferences. Narita 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 Narita [2016]. 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 there is a systematic change in preferences, for instance if students derive utility from obtaining a slot in a school s in the first round itself relative to getting a slot in s in the second round.³⁶ If student preferences are modified in this fashion, our

³⁵Specifically, consider a sequence of markets of increasing size. If the order condition holds in the continuum limit, this should lead to approximate type equivalence under all PLDAs and that RLDA approximately minimizes transfers among PLDAs in the finite markets as market size grows. Moreover, if the local order condition holds for the reverse lottery, then in large finite economies and for every permutation P , the set of students that violate the local order condition on $\text{PLDA}(P)$ will be small relative to the size of the market.

³⁶For example, a student who was allocated her fourth choice in the first round may have a second round preference list that lists her top two choices followed by her fourth choice, i.e., her current allocation, because moving to her erstwhile third choice would not be worth the effort.

characterization result (Theorem 4) justifying PLDA mechanisms still holds, but an example shows that our results regarding type-equivalence and optimality of RLDA break down.

Finally, what insights do our results provide for situations in which allocation 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

- The bylaws of random hall. <http://web.mit.edu/random-hall/www/Government/bylaws/bylaws.html>. Adopted November 9, 2003.
- A. Abdulkadiroğlu 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. Sonmez. House allocation with existing tenants. *Journal of Economic Theory*, 88(2):233–260, October 1999.
- A. Abdulkadiroglu and T. Sönmez. School choice: A mechanism design approach. *American Economic Review*, 93(3):729–747, 2003. doi: 10.1257/000282803322157061.
- A. Abdulkadiroglu and T. Sonmez. *Matching markets: Theory and practice*, pages 1–47. Cambridge University Press, 1 2011. ISBN 9781139060011. doi: 10.1017/CBO9781139060011.002.
- A. Abdulkadiroglu, P. A. Pathak, and A. E. Roth. Strategy-proofness versus efficiency in matching with indifference: Redesigning the NYC high school match. *American Economic Review*, 99(5):1954–78, December 2009.
- A. Abdulkadiroğlu, 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.
- A. Abdulkadirolu, P. A. Pathak, and A. E. Roth. The New York City high school match. *American Economic Review*, 95(2):364–367, May 2005a.
- A. Abdulkadirolu, P. A. Pathak, A. E. Roth, and T. Sönmez. The Boston public school match. *American Economic Review*, 95(2):368–371, May 2005b.
- 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.
- 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.
- 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, 2014a. doi: 10.1287/opre.2014.1319.
- I. Ashlagi and P. Shi. Optimal allocation without money: an engineering approach. In *ACM Conference on Economics and Computation, EC '14, Stanford, CA, USA, June 8-12, 2014*, pages 351–352, 2014b.
- 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.
- E. M. Azevedo and E. B. Budish. Strategy-proofness in the large. *Chicago Booth Research Paper*, (13-35), 2013.
- E. M. Azevedo and J. D. Leshno. A supply and demand framework for two-sided matching markets. *Available at SSRN 2260567*, 2014.
- M. Balinski and T. Sonmez. A tale of two mechanisms: Student placement. *Journal of Economic Theory*, 84(1):73–94, January 1999.
- Chicago Public Schools. Office of access and enrollment, 2017. URL <http://cps.edu/AccessAndEnrollment/Pages/OAE.aspx>.
- J. Combe, O. Tercieux, and C. Terrier. The design of teacher assignment: Theory and evidence. 2016.
- O. Compte and P. Jehiel. Voluntary participation and re-assignment in two-sided matching. Technical report, working paper, 2008.
- Denver Public Schools. School choice, 2017. URL <http://schoolchoice.dpsk12.org/schoolchoice/>.
- U. Dur. Dynamic school choice problem. Technical report, Mimeo, 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. ISSN 00029890. URL <http://www.jstor.org/stable/2312726>.
- A. Hassidim, D. Marciano-Romm, A. Romm, and R. I. Shorrer. Strategic behavior in a strategy-proof environment. Technical report, working paper, 2015.
- S. V. Kadam and M. H. Kotowski. Multi-period matching. Technical report, working paper, 2014.
- 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. 2016.
- Y. Narita. Match or mismatch: Learning and inertia in school choice. 2016.
- NYC Department of Education. High school admission resources, 2017. URL <http://schools.nyc.gov/ChoicesEnrollment/High/Resources/default.htm>.

- P. A. Pathak. The mechanism design approach to student assignment. *Annual Review of Economics*, 3(1):513–536, 09 2011.
- P. A. Pathak and T. Sönmez. School admissions reform in chicago and england: Comparing mechanisms by their vulnerability to manipulation. *American Economic Review*, 103(1):80–106, February 2013.
- J. S. Pereyra. A dynamic school choice model. *Games and economic behavior*, 80:100–114, 2013.
- 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.

A List of Notation

A.1 Discrete Model

- $S = \{s_1, \dots, s_n\}$: schools
- s_{n+1} : outside option
- q_i : capacity of school i
- Λ : finite set of students
- $\succ^\lambda$: first round preferences of student λ over schools
- $\hat{\succ}^\lambda$: second round preferences of student λ over schools
- $\succeq_i^S$: weak priorities of school s_i over students
- p^λ : vector giving the number of the priority group of student λ at each school
- L : student lottery numbers

A.2 Continuum Model

- Λ : mass 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 a given type
- n_i : the number of priority groups at school s_i

A.3 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
- $\mathbf{C}_p$: first round cutoffs restricted to priority type p
- $\hat{\mathbf{C}}_p^P$: second round cutoffs from PLDA with permutation P restricted to priority type 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 scoring function of PLDA
- $X_i = \{s_i, \dots, s_{n+1}\}$: 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 3

- $\beta_{i,j} = \eta(\{\lambda \in \Lambda : \arg\max_{\lambda} X_j = s_i\})$: 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_p : restricted capacity vector for priority type p
- Λ_p : set of students with priority type p
- η_p : restriction of η to students with priority type p
- $\mathcal{E}_p = (S, q_p, \Lambda_p, \eta_p)$: restricted primitives for priority type p
- $s_{\sigma_p(i)}$: i th school under second round overdemand ordering for $\mathcal{E}_p$
- $\tilde{\mathbf{C}}^P$: second round cutoffs defined for PLDA with the amended scoring function from the RLDA cutoffs $\hat{\mathbf{C}}^R$
- $\tilde{\mathbf{C}}_p^P$: second round cutoffs defined for PLDA on $\mathcal{E}_p$ with the amended scoring function from the RLDA cutoffs $\hat{\mathbf{C}}^R$
- $\hat{n}$: smallest index of a school affordable to everyone

A.7 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.8 Proof of Theorem 4

- $s_{\sigma(i)}$: i th school under second round overdemand ordering in a non-atomic mechanism M satisfying axioms (1)-(4)
- $\tilde{X}_i = \{s_{\sigma(i)}, s_{\sigma(i+1)}, \dots, s_{\sigma(n+1)}\}$: 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 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 first round score in the interval I and are assigned to a school in S' in the second round

B Proofs

We begin with some general notation and definitions. Let μ be the initial allocation under RSD, and let P be a permutation. We say that a school s_i *reaches capacity* under a mechanism with output allocation 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 $\hat{C}^P$ are as defined by PLDA(P)) holds simultaneously for all permutations P .

Throughout the appendix, for convenience, we slightly change the second round scoring function of a 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 they 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 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 the amended scoring function) is their *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 .

We say that a cutoff and assignment pair $(\mathbf{C}, \mu)$ are *market-clearing* if $\eta(\mu(s_i)) \leq q_i$ for all $s_i \in S \cup \{s_{n+1}\}$, with equality if $C_i > 0$. It follows from the results in Azevedo and Leshno [2014] that the PLDA cutoffs and assignments $(\mathbf{C}, \mu)$ and $(\hat{\mathbf{C}}^P, \hat{\mu}^P)$ are market-clearing.

Given cutoffs $\mathbf{C}$, we will find it useful to define cutoffs $\mathbf{C}_p$ for each priority type p by $C_{p,i} = \lfloor C_i - p_i \rfloor$, where $\lfloor x \rfloor$ is the largest integer smaller than or equal to x . Intuitively, $C_{p,i}$

is the lottery number a student λ with priority type $p^\lambda = p$ needs in order to be able to go to school s_i .

B.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 finds a stable matching where all schools find all students acceptable, which is non-wasteful.

We now show that the PLDA mechanism is two-round strategyproof. Since students are non-atomic, no student can change the cutoffs $\hat{\mathbf{C}}^P$ by changing their 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, $\hat{r}_i^\lambda$ 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 school s_i when she would have otherwise 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, so it is a dominant strategy for students with consistent preferences to report truthfully in the first round.

PLDA is constrained Pareto efficient among reassigned students, since we use single tie-breaking and the output is stable with respect to the second round lotteries $\hat{r}$. This is easily seen via the cutoff characterization. Fix a Pareto improving cycle that respects priorities. Without guarantees, if a student gets assigned a seat at a school s_i when they prefer s_{i+1} , then they must be in the same priority group as the students in the cycle assigned to s_{i+1} and have a worse lottery number. Hence $\hat{C}_i^P < \hat{C}_{i+1}^P$, which creates a cycle in the cutoffs.

We make the above intuition formal. Fix the first round allocation μ . Suppose there exists a Pareto-improving cycle among reassigned students that respects priorities, that is an ordered set of types $(\theta_1, \theta_2, \dots, \theta_m) \in \Theta^m$, sets of students $(\Lambda_1, \Lambda_2, \dots, \Lambda_m)$, $\Lambda_i \subseteq \Lambda$ and schools $(\tilde{s}_1, \tilde{s}_2, \dots, \tilde{s}_m) \in S^m$ such that for all i it holds that $\eta(\Lambda_i) > 0$, $\tilde{s}_{i+1} \succ^{\theta_i} \tilde{s}_i$ (where we define $\tilde{s}_{m+1} = \tilde{s}_1$), for all $\lambda \in \Lambda_i$ it is the case that $\theta^\lambda = \theta_i$, $\hat{\mu}(\lambda) = s_i$, $\mu(\lambda) \neq s_i$, and for all i and $\lambda \in \Lambda_i$, $\hat{\lambda} \in \Lambda_{i+1}$ it holds that $p_{i+1}^\lambda \geq p_{i+1}^{\hat{\lambda}}$ (where we define $\Lambda_{m+1} = \Lambda_1$ and $p_{m+1} = p_1$).

Let $\lambda_i \in \Lambda_i$. Then, if l_i is the first round lottery number of student λ_i at school s_i (which is a random variable), and $\hat{r}_i$ is the second round lottery number of student λ_i at school s_i , then $\hat{r}_i = p_i^\lambda + P(l_i)$ for all i . Moreover, since $s_{i+1} \succ_{\lambda_i} s_i$ but $\eta(\Lambda_{i+1}) > 0$, where for all $\hat{\lambda} \in \Lambda_{i+1}$ it holds that $p_{i+1}^{\hat{\lambda}} \geq p_{i+1}^{\lambda_i}$, it follows that λ_i and all the students in Λ_{i+1} are in the same priority group at school s_i , say p_i . Moreover, the second round lottery number of student λ_i must be in the interval $[\hat{C}_i^P - p_i, \hat{C}_{i+1}^P - p_i] \subseteq [0, 1]$ (where we define $\hat{C}_{m+1}^P := \hat{C}_1^P$) with positive probability. Since P is fixed, we may also assume that the cutoffs $\hat{C}_i^P$ are fixed and so this interval is non-empty for all i . But as this is true for all i , it follows that $\hat{C}_1^P < \hat{C}_2^P < \dots < \hat{C}_n^P < \hat{C}_1^P$, which is a contradiction.

B.2 Proof of Theorem 1

We first prove Theorem 1 in the case when 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 type p . Hence we can reduce the general problem to the case when all schools have one priority group by restricting to the set of students with priority type p . This shows that all PLDA mechanisms produce type equivalent allocations.

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

Proof. Let P be a permutation.

Assume the order condition holds. By Theorem 3, 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 $\hat{C}_i^P \geq \hat{C}_{i+1}^P$ 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 $X_i = \{s_i, \dots, s_{n+1}\}$. Let

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

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 $\hat{E}_P^\lambda(\hat{\mathbf{C}}^P)$ 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 $\hat{E}_P^\lambda(\hat{\mathbf{C}}^P) = X_j$. The fact that $\hat{E}_P^\lambda(\hat{\mathbf{C}}^P) = X_j$ for some j is a result of the order condition; our modified scoring function guarantees that $E^\lambda(\mathbf{C}) \subseteq \hat{E}_P^\lambda(\hat{\mathbf{C}}^P)$ (every school affordable in the first round is guaranteed in the second) and hence that $j \leq i$. Let $\gamma_i^P = \eta(\{\lambda \in \Lambda : \hat{E}_P^\lambda(\hat{\mathbf{C}}^P) = X_i\})$ be the fraction of students whose affordable set in the second round of PLDA with permutation P is X_i .³⁷ We note that by definition of PLDA, $\eta(\{\lambda \in \Lambda : \theta^\lambda = \theta, \hat{E}_P^\lambda(\hat{\mathbf{C}}^P) = X_i\}) = \zeta(\{\theta\})\gamma_i^P$, 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. $\gamma_i^P > \gamma_i^{P'}$. 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 $\gamma_i^{P'} = 1 - \sum_{j < i} \gamma_j^{P'} = 1 - \sum_{j < i} \gamma_j^P \geq \gamma_i^P$, 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. $\square$

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 type 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 treated symmetrically across different PLDA mechanisms.

³⁷Note that $\eta\Lambda = 1$, as η is a probability distribution over Λ .

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

Proof. Fix a permutation P . By Theorem 3, we may assume that the global order condition holds.

We show that $\text{PLDA}(P)$ 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 $\mathcal{E}_p$ for each priority type p that gives only as many seats as assigned to students of priority type p under RLDA. We then invoke the global order condition and Theorems 3 and 1 to show that all PLDA mechanisms are type equivalent on each $\mathcal{E}_p$. 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 $\tilde{\mathbf{C}}_{p,i}^P$ using the economies $\mathcal{E}_p$ and show that they are (almost) independent of priority type. Finally, we show that this means $\tilde{\mathbf{C}}_{p,i}^P$ also define PLDA cutoffs on the large economy $\mathcal{E}$ and conclude that $\text{PLDA}(P)$ assigns the same measure of students of each priority type to each school s_i as RLDA.

(1) Defining little economies $\mathcal{E}_p$ on each priority type.

Fix a priority type p . Let q_p be a restricted capacity vector, where $q_{p,i}$ is the measure of students of priority type p assigned to school s_i under RLDA. Let Λ_p be the set of students λ such that $p^\lambda = p$, and let η_p be the restriction of the distribution η to Λ_p . Let $\mathcal{E}_p$ denote the primitives $(S, q_p, \Lambda_p, \eta_p)$. Recall that $\hat{\mathbf{C}}^{\mathbf{R}}$ are the second round cutoffs of RLDA on $\mathcal{E}$. It follows from the definition of $\mathcal{E}_p$ that $\hat{\mathbf{C}}_p^{\mathbf{R}}$ are also the second round cutoffs for RLDA on $\mathcal{E}_p$.

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

(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 on the little economies $\mathcal{E}_p$. For all p , the local order condition holds for RLDA on $\mathcal{E}_p$. Hence by Theorem 3 the little economies $\mathcal{E}_p$ each satisfy the order condition. Moreover, by Theorem 1 all PLDA mechanisms on $\mathcal{E}_p$ are type equivalent. Finally, if we can show that for every permutation P , $\text{PLDA}(P)$ assigns the same measure of students of each priority type to each school s_i as RLDA, then $\mathcal{E}$ satisfies the global order condition if and only if for all p the little economy $\mathcal{E}_p$ 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 type p there exists a permutation σ_p such that the affordable set of any student in any round of any PLDA mechanism is of the form $s_{\sigma_p(i)}, s_{\sigma_p(i+1)}, \dots, s_{\sigma_p(n)}, s_{n+1}$. 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 of the cutoffs for a given priority type are 0 or 1 as possible, and the global order condition will still hold.

Formally, for cutoffs C we can equivalently define priority-type-specific cutoffs $C_{p,i} = \lfloor C_i - p_i \rfloor$. Note the cutoffs C_p are consistent across priority types, namely

- Cutoffs match for two priority types with the same priority group at a school,

$$p_i = p'_i \Rightarrow C_{p,i} = C_{p',i} \text{ and } \hat{C}_{p,i} = \hat{C}_{p',i}; \text{ and}$$

- There is at most one marginal priority group at each school,

$$C_{p,i}, C_{p',i} \in (0, 1) \Rightarrow p_i = p'_i.$$

Moreover, if cutoffs C_p are consistent across priority types, then there exist cutoffs C from which they arise.

Suppose that we set as many of the priority-type specific cutoffs $\hat{C}^P$ to 0 as possible, i.e. we let $\hat{C}_i^P$ be 0 if all students have s_i in their affordable set, 1 if none do, and keep it the same otherwise. Then under this new definition, $C, \hat{C}^P$ satisfies the local order condition consistently with all other PLDAs.

Specifically, let

$$f_{p,i}^P(x) = |\{l : l \geq C_{p,i} \text{ or } P(l) \geq x\}|$$

be the proportion of students of priority type p who have school s_i in their affordable set if the first and second round cutoffs are $C_{p,i}$ and x respectively. Notice that f is decreasing in x . If $f_{p,i}^P(\hat{C}_{p,i}^P) = 1$ we set $\hat{C}_{p,i}^P = 0$, and otherwise we keep $\hat{C}_{p,i}^P$ unchanged.

Since $\mathcal{E}$ satisfies the global order condition, for all p there exists an ordering σ_p such that $C_{\sigma_p(1)} \geq C_{\sigma_p(2)} \geq \dots \geq C_{\sigma_p(n)}$ and $\hat{C}_{\sigma_p(1)}^{P'} \geq \hat{C}_{\sigma_p(2)}^{P'} \geq \dots \geq \hat{C}_{\sigma_p(n)}^{P'}$ for all permutations P' . We show that the global order condition implies that the newly defined cutoffs $\hat{C}^P$ satisfy $\hat{C}_{\sigma_p(1)}^P \geq \hat{C}_{\sigma_p(2)}^P \geq \dots \geq \hat{C}_{\sigma_p(n)}^P$. This is because the global order condition implies that f is increasing in i , i.e. for each $p, i < j$ and x it holds that $f_{p,\sigma_p(i)}^P(x) \leq f_{p,\sigma_p(j)}^P(x)$. Hence if we set $\hat{C}_{p,\sigma_p(i)}$ to be 0 then for all $j > i$

$$\begin{aligned} f_{p,\sigma_p(j)}^P(\hat{C}_{p,\sigma_p(j)}) &\geq f_{p,\sigma_p(j)}^P(\hat{C}_{p,\sigma_p(i)}) \quad (\text{since } f \text{ is decreasing}) \\ &\geq f_{p,\sigma_p(i)}^P(\hat{C}_{p,\sigma_p(i)}) \quad (\text{since } f \text{ is increasing in } i) \\ &= 1, \end{aligned}$$

so we set $\hat{C}_{p,\sigma_p(j)}$ to be 0. In other words, we have essentially taken a set of the smallest cutoffs and set them all to 0.

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

We now show that $\tilde{C}_{p,i}^P$ depends on p only via p_i , and does not depend on p_j for all $j \neq i$. We will use the fact that $\mathcal{E}_p$ satisfies the global order condition, and characterize $\tilde{C}_{p,i}^P$ using an affordable set argument.

Now since $\mathcal{E}_p$ satisfies the order condition, all PLDA mechanisms on $\mathcal{E}_p$ are type equivalent, and the proportion of students who have each school in their affordable set is the same across

all PLDA mechanisms. In other words, for all permutations P, P' , priority types p and schools s_i it holds that $f_{p,i}^P(\tilde{C}_{p,i}^P) = f_{p,i}^{P'}(\tilde{C}_{p,i}^{P'})$.

We write this in terms of RLDA as follows. Note that

$$f_{p,i}^R(x) = (1-x) + \min\{x, (1-C_{p,i})\}.$$

- Case 1: $f_{p,i}^R(\hat{C}_{p,i}^R) = 0$.

Then no students with priority type p have s_i in their second round affordable set, so $\tilde{C}_{p,i}^P = 1$.

- Case 2: $f_{p,i}^R(\hat{C}_{p,i}^R) \in (0, 1)$.

Then $\tilde{C}_{p,i}^P$ satisfies the following equation in terms of $\hat{C}_{p,i}^R$, $C_{p,i}$ and P :

$$f_{p,i}^P(\tilde{C}_{p,i}^P) = f_{p,i}^R(\hat{C}_{p,i}^R) = 2 - \hat{C}_{p,i}^R - C_{p,i}. \quad (1)$$

We note that an application of the intermediate value theorem shows that this equation always has a solution in $[0, 1]$, since $f_{p,i}^P(0) = 1 - C_{p,i}$, $f_{p,i}^P(1) = 1$, $f_{p,i}^P$ is continuous and decreasing on $[0, 1]$, and we are in the case when $1 - C_{p,i} < \hat{C}_{p,i}^R < 1$.

Hence $\tilde{C}_{p,i}^P$ is defined by $f_{p,i}^P$ and $f_{p,i}^R$.

- Case 3: $0 \leq \hat{C}_{p,i}^R \leq 1 - C_{p,i}$.

In this case, all students with priority type p have s_i in their second round affordable set and we set $\tilde{C}_{p,i}^P = 0$.

In all three cases, the value of $\tilde{C}_{p,i}^P$ depends on $f_{p,i}^P(\cdot)$, $f_{p,i}^R(\cdot)$ and $\hat{C}_{p,i}^R$, which depend only on $C_{p,i}$ and the permutations P or R . Since $C_{p,i}$ depends on p only through p_i , it follows $\tilde{C}_{p,i}^P$ depends on p only through p_i .

Hence if p, p' are two priority vectors such that $p_i = p'_i$, then $\tilde{C}_{p,i}^P = \tilde{C}_{p',i}^P$, so the $\tilde{C}_{p,i}^P$ are consistent across priority types. Hence the $\tilde{C}_{p,i}^P$ define cutoffs $\tilde{C}_i^P$ that are independent of priority type.

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

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

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

We are now ready to prove Theorem 1

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

Consider the subproblem with primitives $\mathcal{E}_p = (S, q_p, \Lambda_p, \eta_p)$. By Lemma 1, for all $\theta \in \Theta$ and s_i ,

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

Since η_p is the restriction of η to λ_p , it follows that all PLDA mechanisms are type equivalent. $\square$

B.3 Proof of Theorem 2

Proof of Theorem 2. Fix $\theta = (\succ^\theta, \hat{\succ}^\theta, p^\theta) \in \Theta$ and school $s_i \in S$. We will show that R minimizes the measure of reassigned students with type θ who were assigned to school s_i in the first round. The idea is that reversing the lottery shortcuts improvement chains within a particular type, moving one student many schools up their preference list instead of moving many students each a few schools up their preference list. We make this formal.

Let P be a permutation. Let the measure of students with type θ leaving school s_i in the second round under PLDA(P) be denoted by $\ell_P = \eta(\{\lambda \in \Lambda : \theta^\lambda = \theta, \mu(\lambda) = s_i, \hat{\mu}_P(\lambda) \neq s_i\})$. Let the measure of students with type θ entering school s_i under permutation P be denoted $e_P = \eta(\{\lambda \in \Lambda : \theta^\lambda = \theta, \mu(\lambda) \neq s_i, \hat{\mu}_P(\lambda) = s_i\})$. Due to type equivalence, there is a constant c , independent of P , such that $\ell_P = e_P - c$. If $s_{n+1} \hat{\succ}^\theta s_i$, then there are no reassigned students with type θ entering school s_i under either PLDA(P) or RLDA, since all students of type θ prefer to be unassigned. So assume $s_i \hat{\succ}^\theta s_{n+1}$.

We will show that $\ell_R \leq \ell_P$ for all permutations P . If $e_R = 0$ and no students of type θ enter s_i under RLDA, then $\ell_R = e_R - c \leq e_P - c = \ell_P$ for all permutations P , and fewer students of type θ leave s_i under RLDA. So assume $e_R > 0$. We claim that in this case $\ell_R = 0$, which will complete our proof. We will show that having both $e_R > 0$ and $\ell_R > 0$ means that students 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.

Since $e_R > 0$, there exists some student $\lambda \in \Lambda$ with type $\theta^\lambda = \theta$ for which $s_i = \hat{\mu}_R(\lambda) \hat{\succ}^\theta \mu(\lambda)$. By consistency, we have $s_i \succ^\theta \mu(\lambda)$, and therefore λ could not afford (meet the cutoff for) s_i in the first round. If $\ell_R = 0$ we are done. If $\ell_R > 0$, there exists some student $\lambda' \in \Lambda$ with type $(\succ^{\lambda'}, \hat{\succ}^{\lambda'}, p^{\lambda'}) = \theta^\lambda$ for which $s_j = \hat{\mu}_R(\lambda') \hat{\succ}^\theta \mu(\lambda') = s_i$.

By definition, λ' could afford s_i in the first round and λ could not, and hence $L(\lambda') > L(\lambda)$. Note that since $s_i \hat{\succ}^\theta s_{n+1}$, then $s_j \hat{\succ}^\theta s_{n+1}$, and thus consistency gives that $s_j \succ^\theta s_i$.

Now since λ' received a better second round assignment under RLDA than λ , $\hat{\mu}_R(\lambda') \hat{\succ}^\theta \hat{\mu}_R(\lambda)$, and both λ and λ' were reassigned under RLDA, it follows that $R(L(\lambda')) > R(L(\lambda))$.

This contradicts that $L(\lambda') > L(\lambda)$, and therefore $\ell_R = 0$, completing our proof. $\square$

B.4 Proof of Theorem 3

Proof of Theorem 3. 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.

(1) Every school has a single priority group

We first consider the case when $n_i = 1$ 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 $C_1 \geq C_2 \geq \dots \geq C_n \geq C_{n+1}$. 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 $\hat{C}_1^R \geq \hat{C}_2^R \geq \dots \geq \hat{C}_n^R \geq \hat{C}_{n+1}^R$.

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 $\tilde{C}_1^P \geq \tilde{C}_2^P \geq \dots \geq \tilde{C}_n^P \geq \tilde{C}_{n+1}^P$ 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 $\beta_{i,j} = \eta(\{\lambda \in \Lambda : \arg\max_{\lambda} X_j = s_i\})$ 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 $\gamma_i^R = \eta(\{\lambda \in \Lambda : \hat{E}_P^\lambda(\hat{\mathbf{C}}^R) = X_i\})$ be the fraction of students whose affordable set in the second round of RLDA is X_i , and let $\gamma_i^P = \eta(\{\lambda \in \Lambda : \hat{E}_P^\lambda(\hat{\mathbf{C}}^P) = X_i\})$ 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 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 their (second round) 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 $\hat{n}$ always exists, since every student has the outside option s_{n+1} in their 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

$$f_i^P(\tilde{C}_i^P) = f_i^R(\hat{C}_i^R) \quad (2)$$

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

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

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 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 $\tilde{C}_1^P \geq \tilde{C}_2^P \geq \dots \geq \tilde{C}_n^P \geq \tilde{C}_{n+1}^P$.

(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 $\hat{r}_i^\lambda = P(L(\lambda)) + n_i \mathbb{1}_{\{L(\lambda) \geq C_i\}} + p_i^\lambda \mathbb{1}_{\{L(\lambda) > C_i\}}$. 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 that we have defined these to be equal,

$f_i^P(\tilde{C}_i^P) = f_i^R(\hat{C}_i^R)$. It follows that $\tilde{C}^P$ are market-clearing and give the PLDA(P) cutoffs, 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 $\hat{C}_1^R \geq \hat{C}_2^R \geq \dots$, under RLDA this is given by

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

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

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

if $i < \hat{n}$ which is precisely γ_i^R , and 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 $\tilde{C}^P$ are market-clearing cutoffs when the second round scores are given by $\hat{r}_i^\lambda = P(L(\lambda)) + n_i \mathbb{1}_{\{\lambda \geq C_i\}} + p_i^\lambda \mathbb{1}_{\{\lambda \geq C_i\}}$, so PLDA(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 $\mathcal{E}_p$ on each priority type.

Fix a priority type p . Let q_p be a restricted capacity vector, where $q_{p,i}$ is the measure of students of priority type p assigned to school s_i under RLDA. Let Λ_p be the set of students λ such that $p^\lambda = p$, and let η_p be the restriction of the distribution η to Λ_p . Let $\mathcal{E}_p$ denote the primitives $(S, q_p, \Lambda_p, \eta_p)$.

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

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

Since RLDA satisfies the local order condition on $\mathcal{E}$, RLDA also satisfies the local order condition on $\mathcal{E}_p$ for all p . It follows from (1) that the global order condition holds on each of the little economies $\mathcal{E}_p$. Hence by Theorem 1 all PLDA mechanisms on $\mathcal{E}_p$ are type equivalent. Moreover, as in the proof of Theorem 6, the global order condition on $\mathcal{E}_p$ also allows us to determine aggregate student demand in $\mathcal{E}_p$ 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 p and school s_i the cutoff $\tilde{C}_{p,i}^P$ is the minimal real satisfying

$$f_{p,i}^P(\tilde{C}_{p,i}^P) = f_{p,i}^R(\hat{C}_{p,i}^R)$$

where for each permutation P

$$f_{p,i}^P(x) = |\{l : l \geq C_{p,i} \text{ or } P(l) \geq x\}|$$

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

It follows that $\tilde{C}_{p,i}^P$ depends on p only via p_i , and does not depend on p_j for all $j \neq i$. This is because the value of $\tilde{C}_{p,i}^P$ depends on $f_{p,i}^P(\cdot)$, $f_{p,i}^R(\cdot)$ and $\hat{C}_{p,i}^R$, which depend only on $C_{p,i}$ and the permutations P or R . Moreover, $C_{p,i}$ depends on p only through p_i . Hence if p, p' are two priority vectors such that $p_i = p'_i$, then $\tilde{C}_{p,i}^P = \tilde{C}_{p',i}^P$, so the $\tilde{C}_{p,i}^P$ are consistent across priority types. Hence the $\tilde{C}_{p,i}^P$ define cutoffs $\tilde{C}_i^P$ that are independent of priority type.

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

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

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

B.5 Proof of Theorem 4

Proof of Theorem 4. Let a *permuted lottery stable matching* mechanism (PLSM) be a reassignment mechanism that is non-atomic and outputs a matching that is stable under the same second round school preferences as in PLDA. In a slight abuse of notation, we will also let PLSM refer to the two-round mechanism defined by this reassignment mechanism. We need to show that PLSM satisfies the axioms and that any mechanism satisfying the axioms is a PLSM.

We first recall the cutoff characterization of the set of stable matchings for given student preferences and responsive school preferences (encoded by student scores r), as provided by Azevedo and Leshno in Azevedo and Leshno [2014]. Namely, if $\mathbf{C} \in \mathbb{R}_+^{n+1}$ 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, $\mu(\lambda) = \max_{\succ \lambda} (\{s_i \in S \cup \{s_{n+1}\} : r_i^\lambda \geq C_i\})$. We say that $\mathbf{C}$ is 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 $s_i \in S \cup \{s_{n+1}\}$, with equality if $C_i > 0$. Then the set of all stable matchings is precisely given by the set of assignments defined by market-clearing vectors Azevedo and Leshno [2014].

Recall also that under PLDA with permutation 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 $s_i \in S \cup \{s_{n+1}\}$. In a slight abuse of notation, we will sometimes let $\hat{\mathbf{C}}^P$ refer to the second round cutoffs from some fixed PLSM with permutation P (not necessarily corresponding to the student-optimal stable matching given by PLDA).

(1) Any PLSM satisfies the axioms: The proof that any PLSM satisfies the axioms essentially follows from Proposition 1. We provide a full proof here for completeness.

Fix a permutation P and some PLSM with permutation P . We first show that this particular PLSM satisfies all the axioms. 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 PLSM for this distribution of student types. PLSM respects guarantees because fewer students are guaranteed at each school than the capacity of the school. PLSM is non-wasteful because the second round finds a stable matching where all schools find all students acceptable, which is non-wasteful.

We now show that the PLSM is two-round strategyproof. Since students are non-atomic, no student can change the cutoffs $\hat{C}^P$ by changing their 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, $\hat{r}_i^\lambda$ increases by 1. This means that having a guarantee at a school s_i changes the student's second round assignment in the following way. She receives school s_i when she would have otherwise 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, so it is a dominant strategy for students with consistent preferences to report truthfully in the first round.

PLSM is Pareto efficient among reassigned students, since we use single tie-breaking and the output is stable with respect to the second round lotteries $\hat{r}$. This is easily seen via the cutoff characterization. Without guarantees, if a student gets assigned a seat at a school s_i when they prefer s_{i+1} , then $\hat{C}_i^P < \hat{C}_{i+1}^P$, so there cannot be a cycle in the preferences of such students.

We make the above intuition formal. Fix the first round allocation μ . Suppose there exists a Pareto-improving cycle among reassigned students, that is an ordered set of types $(\theta_1, \theta_2, \dots, \theta_m) \in \Theta^m$, sets of students $(\Lambda_1, \Lambda_2, \dots, \Lambda_m), \Lambda_i \subseteq \Lambda$ and schools $(\tilde{s}_1, \tilde{s}_2, \dots, \tilde{s}_m) \in S^m$ such that for all i it holds that $\eta(\Lambda_i) > 0$, $\tilde{s}_{i+1} \succ_{\theta_i} \tilde{s}_i$ (where we define $\tilde{s}_{m+1} = \tilde{s}_1$), and for all $\lambda \in \Lambda_i$ it is the case that $\theta^\lambda = \theta_i, \hat{\mu}(\lambda) = s_i, \mu(\lambda) \neq s_i$.

Let $\lambda_i \in \Lambda_i$. Then, if l_i is the first round lottery number of student λ_i at school s_i (which is a random variable), and $\hat{r}_i$ is the second round lottery number of student λ_i at school s_i , then $\hat{r}_i = P(l_i)$ for all i . Moreover, since $s_{i+1} \succ_{\lambda_i} s_i$ but $\eta(\Lambda_{i+1}) > 0$, it follows that the second round lottery number is in the interval $[\hat{C}_i^P, \hat{C}_{i+1}^P]$ (where we define $\hat{C}_{m+1}^P := \hat{C}_1^P$) with positive probability. Since P is fixed, we may also assume that the cutoffs $\hat{C}_i^P$ are fixed and so this interval is non-empty for all i . But this is clearly a contradiction.

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. Anonymity is easily checked.

(2) *Any mechanism satisfying the axioms is a PLSM:*

We now show that any mechanism M satisfying the axioms is a PLSM in a particular sense. We will show specifically that if we assume that each instantiation of M provides type equivalent output, then M is type equivalent to a PLSM. Moreover, if we assume that conditional on their reports, students' assignments under M are uncorrelated, we are able to explicitly construct a PLSM that provides the same output 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 deferred acceptance with uniform-at-random single tie-breaking 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 uniform-at-random single tie-breaking, and the mechanism satisfies the averaging axiom, we may assume that each pathwise realization of the mechanism gives type equivalent (two-round) allocations. Hence, for the majority of the proof we perform our constructions of aggregate cutoffs and measures of students pathwise, and assume that any realization 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 Pareto efficiency among reassigned students to construct a first round overdemand ordering $s_1, s_2, \dots, s_n, s_{n+1}$ and a permutation σ giving the second round overdemand ordering $s_{\sigma(1)}, s_{\sigma(2)}, \dots, s_{\sigma(n+1)}$, as in Ashlagi and Shi [2014a], 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 guarantee.) The existence of these orderings follow from the fact that the first round mechanism, DA with a single priority class and uniform-at-random single tie-breaking, is Pareto efficient, combined with the fact that the two-round mechanism is Pareto efficient among reassigned students. We let $X_i = \{s_i, s_{i+1}, \dots, s_{n+1}\}$ denote the set of schools after s_i in the first round overdemand ordering, and $\tilde{X}_i = \{s_{\sigma(i)}, s_{\sigma(i+1)}, \dots, s_{\sigma(n+1)}\}$ 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 their favorite school in their affordable set for that round. We use 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 (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. We first define a 'prefix property' and show that it holds (Lemma 2). We then use this prefix property to show that it is possible to construct the proportions $\gamma_{i,j}$ so that they are both the same for every student type θ , and consistent with the proportion of students in each type assigned to each pair of schools jointly over the first and second round assignments.

We then construct the lottery L and verify that if second round school preferences $\succeq^S$ 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 defining the first and second round affordable sets via cutoffs given by $\gamma_{i,j}$ is consistent with the first round assignments and gives a second round assignment that is stable with respect to second round school and student preferences.

Formal Proof:

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

Let the schools be numbered $s_1, s_2, \dots, s_n$ 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 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 the schools fill in the order $s_{\sigma(1)}, s_{\sigma(2)}, \dots, s_{\sigma(n)}$ in the second round, and which we will eventually show gives the same outcome as a PLSM with cutoffs $\hat{C}_{\sigma(1)} \geq \hat{C}_{\sigma(2)} \geq \dots$. 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 are not assigned to s' in the first round, but are assigned to s' in the second round, then $s = s_{\sigma(i)}$, $s' = s_{\sigma(j)}$ for some $i < j$. We assume that σ is the unique permutation satisfying this 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 exist, $\sigma(i) < \sigma(j)$ iff $i < j$. We also define $\sigma(n+1) = n+1$. An ordering σ with the required properties exists since the mechanism is Pareto efficient among reassigned students, so every trading cycle involves a student swapping a share in their first round guarantee.

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 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 $\rho^\theta(I, \{s'\})$.³⁸ In this section, for brevity, when defining preferences $\succ$ we will sometimes write $\succ: [s_{i_1}, s_{i_2}, \dots, s_{i_k}]$ instead of $s_{i_1} \succ s_{i_2} \succ \dots \succ s_{i_k}$.

(2a) 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 $j = 1, 2, \dots, n$ to be the unique solutions to the following n equations:

$$\begin{aligned} \gamma_{i,j} &= 0 \quad \text{for all } j \text{ such that } \sigma(j) \geq i \\ \gamma_{i,1} + \dots + \gamma_{i,j} &= \rho_{i,j} \quad \text{for all } j \text{ such that } \sigma(j) < i. \end{aligned}$$

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 $s_{\sigma(j-1)}$. If $\sigma(j) \geq i$ then this would include all students whose first round lottery is in $[C_i, C_{i-1}]$, and so we define $\gamma_{i,j} = 0$.

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.

³⁸Here we are assuming that this proportion is the same for every realization of the first round of M . This requires non-atomicity and anonymity.




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.

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 i ,

$$\tilde{I}_{n+1,n+1}, \tilde{I}_{n+1,n}, \dots, \tilde{I}_{1,2}, \tilde{I}_{1,1}.^{39}$$

The interval $\tilde{I}_{i,j}$ can be thought of as the lottery numbers of students whose first round lottery is in $[C_i, C_{i-1}]$ and whose second round affordable set contains $s_{\sigma(j)}$ but not $s_{\sigma(j-1)}$ (whose only available school in the second round comes from their first round guarantee, if $j = n + 1$).

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

$$P(\tilde{I}_{n+1,n+1}), P(\tilde{I}_{n,n+1}), \dots, P(\tilde{I}_{2,1}), P(\tilde{I}_{1,1}).^{40}$$

In Figure 5, we show an example for 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 show that there exists a PLSM mechanism with permutation P , where the students with first round scores in $\tilde{I}_{i,j}$ are precisely the students with first round affordable set X_i and second round affordable set $\tilde{X}_j$, and that this PLSM 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 so 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 a stable matching under the second round school preferences given by P .

(2b) 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 round affordable set X_i and second round affordable set $\tilde{X}_j$, 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 lemma about prefixes.

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

³⁹Specifically, let $\tilde{I}_{i,j} = [C_{i-1} - \sum_{j' \leq j} \gamma_{i,j'}, C_{i-1} - \sum_{j' < j} \gamma_{i,j'}]$.

⁴⁰Specifically, let $\hat{C}_{\sigma(j)} = 1 - \sum_{i', j': j' \leq j} \gamma_{i', j'}$, and let $P(\tilde{I}_{i,j}) = [\hat{C}_{\sigma(j-1)} - \sum_{i' \leq i} \gamma_{i', j}, \hat{C}_{\sigma(j-1)} - \sum_{i' < i} \gamma_{i', j}]$.

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 the proportion of such 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 such 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. Equivalently, $\rho^\theta([C_j, C_{i(S')}], S') = \rho^{\theta'}([C_j, C_{i(S')}], S')$.

Sketch of proof of Lemma 2. The idea of the proof is to use the full support assumption to identify students who are essentially indifferent between schools in S' , and then use two-round strategy-proofness to show that they are indifferent between reporting either θ or θ' . This shows that the conditional probabilities of being assigned to S' are the same for students of type θ or θ' (conditional on certain first round assignments). We then invoke anonymity to argue that proportions of types of students assigned to a certain school are given by the conditional probabilities for individual students of 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.

(2b.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 $\succ: [s_{i_1}, s_{i_2}, s_{n+1}]$ and $\hat{\succ}: [s_{i_1}, s_{n+1}]$, where all other schools are unacceptable. 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 s' in the second round, then the ordered pairs are (s_{i_1}, s_{i_1}) , (s_{i_2}, s_{i_1}) , (s_{i_2}, s_{n+1}) , (s_{n+1}, s_{i_1}) , 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_{i_2}, s_{i_1}) , and the same proportion of students assigned to (s_{n+1}, s_{i_1}) .

The proportion of students with preferences θ who are assigned to (s_{i_2}, s_{i_1}) and (s_{n+1}, s_{i_1}) under M are given by $\rho^\theta([C_{i_2}, C_{i_1}], s_{i_1})$ and $\rho^\theta([0, C_{\max\{i_1, i_2\}}], s_{i_1})$ respectively. We want to show that this is the same as the proportion of students with preferences θ who are assigned to (s_{i_2}, s_{i_1}) and (s_{n+1}, s_{i_1}) respectively when first and second round affordable sets are given by the affordable set distribution $\gamma_{i,j}$. We remark that when $i_1 > i_2$ 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 $i_1 < i_2$.

Let $\theta' = (\succ', \hat{\succ}')$ be the preferences given by $\succ': [s_{i_1}, s_{i_1+1}, \dots, s_{i_2-1}, s_{i_2}, s_{n+1}]$ and $\hat{\succ}': [s_{i_1}, s_{n+1}]$, where only the schools with indices between i_1 and i_2 are acceptable in the first round, only s_{i_1} is acceptable in the second round, and all other schools are unacceptable.

Recall that for all $i > i_1$, $\theta_{i, \sigma^{-1}(i_1)} = (\succ_{i, \sigma^{-1}(i_1)}, \hat{\succ}_{i, \sigma^{-1}(i_1)})$ are the student preferences such that $\succ_{i, \sigma^{-1}(i_1)}: [s_{i_1}, s_{i-1}, s_i, s_{n+1}]$ and $\hat{\succ}_{i, \sigma^{-1}(i_1)}: [s_{i_1}, s_{n+1}]$, with all other schools unacceptable, and that $\rho_{i, \sigma^{-1}(i_1)}$ is the proportion of students of type $\theta_{i, \sigma^{-1}(i_1)}$ whose first round assignment is

s_i and whose second round assignment is school s_{i_1} . (We note that in the case when $i = i_1 + 1$, we let the first two schools under the preference ordering $\succ_{i, \sigma^{-1}(i_1)}$ coincide.)

Then the proportion of students with preferences θ who are assigned to (s_{i_2}, s_{i_1}) under M is given by $\rho^\theta([C_{i_2}, C_{i_1}], s_{i_1})$, where

$$\begin{aligned}
\rho^\theta([C_{i_2}, C_{i_1}], s_{i_1}) &= \rho^{\theta'}([C_{i_2}, C_{i_1}], s_{i_1}) \text{ (by the prefix property (Lemma 2))} \\
&= \sum_{i_1 < i \leq i_2} \rho^{\theta'}([C_i, C_{i-1}], s_{i_1}) \\
&= \sum_{i_1 < i \leq i_2} \rho^{\theta_{i, \sigma^{-1}(i_1)}}([C_i, C_{i-1}], s_{i_1}) \\
&\quad \text{(since the second round assignment does not depend on the first round report)} \\
&= \sum_{i_1 < i \leq i_2} \rho_{i, \sigma^{-1}(i_1)} \text{ (by the definition of } \rho_{i, \sigma^{-1}(i_1)}) \\
&= \sum_{i_1 < i \leq i_2} \sum_{j \leq \sigma^{-1}(i_1)} \gamma_{i,j} \text{ (by the definition of } \gamma_{i,j}),
\end{aligned}$$

which is precisely the proportion of students with preferences θ who are assigned to (s_{i_2}, s_{i_1}) 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 $\succ'': [s_{i_1}, s_{i_2}, s_{i_2+1}, \dots, s_n, s_{n+1}]$ and $\hat{\succ}'' : [s_{i_1}, s_{n+1}]$, where only s_{i_1} and the schools with indices greater than i_2 are acceptable in the first round, only s_{i_1} 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_{i_1}) under M is given by $\rho^\theta([0, C_{\max\{i_1, i_2\}}], s_{i_1})$, where

$$\begin{aligned}
\rho^\theta([0, C_{\max\{i_1, i_2\}}], s_{i_1}) &= \rho^{\theta''}([0, C_{\max\{i_1, i_2\}}], s_{i_1}) \text{ (by the prefix property (Lemma 2))} \\
&= \sum_{\max\{i_1, i_2\} < i \leq n} \rho^{\theta''}([C_i, C_{i-1}], s_{i_1}) \\
&= \sum_{\max\{i_1, i_2\} < i \leq n} \rho^{\theta_{i, \sigma^{-1}(i_1)}}([C_i, C_{i-1}], s_{i_1}) \\
&\quad \text{(since the second round assignment does not depend on the first round report)} \\
&= \sum_{\max\{i_1, i_2\} < i \leq n} \rho_{i, \sigma^{-1}(i_1)} \text{ (by the definition of } \rho_{i, \sigma^{-1}(i_1)}) \\
&= \sum_{i,j: \max\{i_1, i_2\} < i \leq n, j \leq \sigma^{-1}(i_1)} \gamma_{i,j} \text{ (by the definition of } \gamma_{i,j}),
\end{aligned}$$

which is precisely the proportion of students with preferences θ who are assigned to (s_{n+1}, s_{i_1}) if the first and second round affordable sets are given by $\gamma_{i,j}$.

(2b.ii.) Students with general preferences:

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

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

for some $k > l$ and where all other schools are unacceptable. We wish to show that for every pair of schools $s, s' \in \{s_{i_1}, s_{i_2}, \dots, s_{i_k}, s_{n+1}\}$, 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 $s \in \{s_{i_2}, \dots, s_{i_k}, s_{n+1}\}$, 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 $\succ$ and a school $s = s_{i_j}$, $1 < j \leq k$. Let $m \leq k$ be such that $S' = \{s_{i_1}, s_{i_2}, \dots, s_{i_m}\}$. If $j \leq m$ then $s \in S'$, 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.

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 $j > m$ (ie. $s \notin S'$) 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

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

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

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

where the school $s_{i(\sigma(S'))}$ comes first and otherwise the schools in S' are ordered arbitrarily.

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 $i > i(\sigma(S'))$, $\theta_{i, \sigma^{-1}(i(\sigma(S'))))} = (\succ_{i, \sigma^{-1}(i(\sigma(S'))))}, \hat{\succ}_{i, \sigma^{-1}(i(\sigma(S'))))})$ are the student preferences such that

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

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

Let $\hat{S} = \{s_{i_1}, s_{i_2}, \dots, s_{i_{j-1}}\}$, 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

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

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 $\gamma_{i, j'}$.

(3) 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 $\bar{j}$ 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 $\{X_{i'}, X_{j'} \cup \{s_i\} : \underline{i} \leq i' \leq i, \underline{j} \leq j' \leq \bar{j}\}$. (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 $\cup_{i', j': \underline{i} \leq i' \leq i, \underline{j} \leq j' \leq \bar{j}} \tilde{I}_{i', j'}$,

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

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 $\text{Unif}[0, 1]$.

(4) Constructing the PLSM and verifying stability

Given the constructed lottery L , we construct the second round cutoffs $\hat{C}_i$ for the PLSM and verify that the allocation $\tilde{\mu}$ is feasible and stable with respect to the schools' second round

preferences, as defined by $P \circ L$ and the guarantee structure. Specifically, in PLSM, each student with first round score l and first round allocation s has a second round score $\hat{r}_i = P(l) + \mathbb{1}(s = s_i)$ at each school $s_i \in S \cup \{s_n + 1\}$, and students are assigned to their favorite school s_i at which their second round score exceeds the school's second round cutoff, $\hat{r}_i \geq \hat{C}_i$.

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

By the characterization of stable allocations given in Azevedo and Leshno [2014], it suffices to show that if each student with first round allocation s and second round lottery number in $[\hat{C}_{\sigma^{-1}(i)}, \hat{C}_{\sigma^{-1}(i-1)}]$ is assigned to their favorite school in $\{s\} \cup \tilde{X}_i$, where $\tilde{X}_i = \{s_{\sigma(i)}, s_{\sigma(i+1)}, \dots, s_{\sigma(n+1)}\}$, then the resulting assignment $\hat{\mu}$ is equal to the second round assignment $\tilde{\mu}$ of our mechanism M , and satisfies that $\eta(\hat{\mu}^{-1}(s_i)) \leq q_i$ 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 to $\gamma_{i, j}$.)

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 first round lottery number $L(\lambda)$ and reported preferences $\theta = (\succ, \hat{\succ})$. Let i, j be such that $L(\lambda) \in \cup_{i', j': \underline{i} \leq i' \leq i, \underline{j} \leq j' \leq \bar{j}} \tilde{I}_{i', j'}$, where $\underline{i}$ is the smallest index i' such that $\max_{\succ} X_{i'} = s_i$, let $\underline{j}$ is the smallest index j' such that $\max_{\hat{\succ}} \tilde{X}_{j'} \cup \{s_i\} = s_j$, and $\bar{j}$ 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

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

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 their favorite school in $\{s_i\} \cup \tilde{X}_{j'}$ for some $\underline{j} \leq j' \leq \bar{j}$, which is the school s_j . Hence $\tilde{\mu}(\lambda) = \hat{\mu}(\lambda) = s_j$.

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

Finally, let us check that the allocation 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 $\succ: [s_j, s_1, s_2, \dots, s_{j-1}, s_{j+1}, \dots, s_{n+1}]$. 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 Pareto efficient among reassigned students, and hence non-wasteful, this means that $\eta(\tilde{\mu}^{-1}(s_j)) = q_j$.

□

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 and by two-round strategy-proofness, any schools reported to be acceptable but ranked below s in the second round are inconsequential. We show this below.

Suppose there are two student preferences $\theta = (\succ, \hat{\succ})$ and $\theta' = (\succ, \hat{\succ}')$ that differ only in the second round preferences, and $\hat{\succ}'$ is given by truncating the preferences $\hat{\succ}$ at the school s and making all schools worse than s under $\hat{\succ}$ unacceptable. Let $\hat{\succ}' : [s_{i_1}, s_{i_2}, \dots, s_{i_k} = s, s_{n+1}]$ and let $S' = \{s_{i_1}, s_{i_2}, \dots, s_{i_k}\}$. Since the mechanism respects guarantees, students reporting θ and θ' are assigned only to schools in S' . Suppose for the sake of contradiction that the mechanism M treats students with reported preferences θ and θ' who are assigned to s in the first round differently, that is, the distribution of $(\tilde{\mu}(\lambda) \mid \mu(\lambda) = s, \lambda \text{ reports preferences } \theta)$ is different to the distribution of $(\tilde{\mu}(\lambda) \mid \mu(\lambda) = s, \lambda \text{ reports preferences } \theta')$. Let j be the smallest index for which the probability of assignment differs,

$$\mathbb{P}(\tilde{\mu}(\lambda) = s_{i_j} \mid \mu(\lambda) = s, \lambda \text{ reports } \theta) \neq \mathbb{P}(\tilde{\mu}(\lambda) = s_{i_j} \mid \mu(\lambda) = s, \lambda \text{ reports } \theta').$$

Suppose the left hand side is larger than the right hand side by some amount Δ . Consider a student with true preferences θ' whose cardinal utilities for schools in $\{s_{i_1}, s_{i_2}, \dots, s_{i_j}\}$ are $v, v - \varepsilon, \dots, v - (j - 1)\varepsilon$, for all other schools in S' are $\frac{\varepsilon}{j+1}, \dots, \frac{\varepsilon}{k}$, and for all other schools is 0. Then by misreporting $\hat{\succ}'$ in the second round, they increase their interim expected utility by at least $\Delta(v - (j - 1)\varepsilon) - \varepsilon$, which for sufficiently small ε is positive, contradicting two-round strategy-proofness. Hence the left hand side is smaller than the right hand side, and a similar argument for a student with true preferences θ shows that the probabilities must in fact be equal.

Hence it suffices to prove the lemma for first round preference orderings $\succ$ and $\succ'$ for which s is the last acceptable school. We note that the above argument proves the prefix property for any two sets of preferences θ, θ' and first round assignment s where the first round preferences are identical. When the first round preferences are different, because students cannot misreport their first round preferences after the fact, we will need to use two-round strategy-proofness for a number of additional student preferences in order to show that the prefix property holds. When the second round preferences are different, we will need to consider more fine-grained information that students have about the lottery number, and use two-round strategy-proofness for a number of different student preferences based on this fine-grained information. We provide the formal argument below.

Let $i_1, \dots, i_k$ 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 and vary only θ' . Therefore, we may, without loss of generality, fix the preferences θ to satisfy that

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

and all other schools 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, it follows from the structure of the preferences θ and θ' that the proportion of students with preferences θ who can deduce that their score is in $[C_j, C_{i(S')}]$ is precisely $C_{i(S')} - C_j$, and similarly for θ' . 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 $1 - C_{i(S')}$. (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.) We remark that the set of students with preferences θ and lottery number in $[C_{i(S')}, 1]$ is precisely the set of students with preferences θ who are assigned to a school in S' in the first round, and similarly for θ' .

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 $\theta'' = (\succ', \succ)$ 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 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 a student with preferences θ'' and 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 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 Bayesian incentive compatibility, 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 run through the full argument below.

Consider a student λ_0 with preferences θ whose second round cardinal utilities over schools in S' are $v, v - \varepsilon, v - 2\varepsilon, \dots, v - (|S'| - 1)\varepsilon$, for some small ε and some large v , and are 0 for all other schools and the outside option.

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.

By truthfully reporting preferences θ , student λ_0 has expected ex ante utility of at most

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

where the first term is an upper bound on their expected utility from having a lottery number in $[C_{i(S')}, 1]$ in the first round, the second term is an upper bound on their expected utility from having a lottery number in $[C_j, C_{i(S')}]$ in the first round, and the third term is an upper bound on their expected utility from having a lottery number in $[0, C_j]$ in the first round.

By misreporting preferences θ'' , they have expected ex ante utility of at least

$$(1 - C_{i(S')})(v - |S'|\varepsilon) + (C_{i(S')} - C_j)\rho''(v - |S'|\varepsilon) + C_j\pi(v - |S'|\varepsilon).$$

From two-round strategy-proofness for λ_0 and taking ε to zero, it follows that $\rho \geq \rho''$. A symmetric argument for a student with preferences θ'' with the same second round cardinal utilities gives the reverse inequality, and hence $\rho = \rho''$.

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 Bayesian incentive compatibility for students who have high value for schools in S' and low value outside of S' 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 $\{j_1 \leq \dots \leq j_m\}$ be the indices between $i(S')$ and j corresponding to schools that a student with preferences θ' and 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 $l = 1, \dots, m$, 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 students of 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 $[C_{j_l}, C_{j_{l-1}}]$. 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 with preferences θ' assigned to school s_{j_l} in the first round whose second round cardinal utilities are, for some small ε and some large v , $v, v - \varepsilon, v - 2\varepsilon, \dots, v - (|S'| - 1)\varepsilon$

in the appropriate order for schools in S , $\varepsilon, \frac{\varepsilon}{2}, \dots$ for the other acceptable schools, and 0 for all other schools and the outside option.

By truthfully reporting preferences θ' , student λ_l has expected interim utility of at most

$$\rho'_l v + (1 - \rho'_l) \varepsilon,$$

and by misreporting preferences θ'' , they have expected interim utility of at least

$$\rho''_l (v - |S'| \varepsilon),$$

so from second round strategy-proofness for λ_l and taking ε to 0, it follows that $\rho'_l \geq \rho''_l$.

Similarly, consider a student λ'_l with preferences θ'' with the same first round guarantee s_l , second round cardinal utilities over schools in S' also $v, v - \varepsilon, \dots, v - (|S'| - 1)\varepsilon$ in the order corresponding to preferences $\succ$, and second round cardinal utilities for other acceptable schools $\varepsilon, \frac{\varepsilon}{2}, \dots$. Second round strategy-proofness for these students gives the reverse inequality $\rho'_l \leq \rho''_l$, and hence $\rho'_l = \rho''_l$. This completes the proof of the lemma. $\square$

C Example of non-strategy-proofness of PLDA

In this section, we provide an example illustrating that when non-atomicity does not hold, PLDA mechanisms are not necessarily strategyproof.

Example 2. Consider a setting with $n = 2$ schools and $m = 4$ 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.

Assume student 2's utility is M for s_1 , ϵ for s_2 , and 0 for s_3 in both rounds, where $M \gg \epsilon > 0$. Consider the lottery which yields $1 \succ^B 2 \succ^B 3 \succ^B 4$. If the students report truthfully, the first round assignment is

$$\mu(A) = (\mu(1), \mu(2), \mu(3), \mu(4)) = (s_1, s_2, \emptyset, \emptyset),$$

and the reassignment is

$$\hat{\mu}(A) = (\hat{\mu}(1), \hat{\mu}(2), \hat{\mu}(3), \hat{\mu}(4)) = (\emptyset, s_2, s_1, \emptyset).$$

However, consider what happens if student 2 reports $s_1 \succ_2^r \emptyset \succ_2^r s_2$ in both rounds. Then, the first round assignment becomes $\mu(A) = (s_1, \emptyset, s_2, \emptyset)$, and the reassignment becomes

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

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

We remark that this reassignment was not stable in the second round when students reported truthfully, since, in that case school s_2 had second round preferences $2 \hat{\succ}_2^S 4 \hat{\succ}_2^S 3 \hat{\succ}_2^S 1$, and so school s_2 and student 4 form a blocking pair.

Consider now the expected utility of student 2 from reporting truthfully and her expected utility from misreporting, when all other students report truthfully and the expectation is over the first round lottery order. With probability $\frac{1}{4!}$, the lottery order is $1 \succ^B 2 \succ^B 3 \succ^B 4$, in which case student 2 can change her assignment from s_2 to s_1 by reporting s_2 as unacceptable. Moreover, one can verify that for any lottery order, if student 2 received s_1 in the first or second round under truthful reporting, then she also received s_1 in the same round by misreporting.⁴¹ Hence, by misreporting in this particular fashion, student 2 increases her probability of receiving s_1 by at least $\frac{1}{4!}$. Thus, for M sufficiently large with respect to ϵ , this violates strategy-proofness.

⁴¹This 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.