

Undergraduate Course Allocation through Competitive Markets*

Daniel Kornbluth Alexey Kushnir

August 14, 2024

Abstract

We consider the problem of allocating courses to students in post-secondary institutions. We propose a mechanism that assigns course seats based on student preferences and respects course priorities. This mechanism uses fake money and competitive equilibrium to allocate courses without transfers and has desirable theoretical properties in terms of stability, efficiency, fairness, and strategy-proofness. In simulations drawing from real-world university data, we demonstrate that its outcomes improve student satisfaction and allocation fairness over the outcomes of several celebrated mechanisms.

JEL classification: D47, D63, D82, C63, C78, I21

Keywords: market design, matching, many-to-many assignment, deterministic assignment, course allocation, approximate competitive equilibrium, random serial dictatorship, set-asides

**Kornbluth:* Harvard University, Department of Economics/Carnegie Mellon University; dkornbluth@g.harvard.edu; *Kushnir (corresponding author):* Tepper School of Business, Carnegie Mellon University, Pittsburgh, 15213, USA; +1 412 268 6079; akushnir@andrew.cmu.edu. We are especially thankful to Dennis Epple for his resolute support of this project. We are indebted to individuals who provided us with course enrollment data and consulted us on the process of undergraduate course allocation. They have requested to remain anonymous. We are also thankful to Maryam Saeedi for providing access to her computer cluster, which allowed us to run more simulations in a shorter period. We are also grateful to Martin Bichler, Eric Budish, Karam Kang, Rebecca Lessem, Anh Nguyen, Marek Pycia, and participants at various seminars and conferences for their numerous comments and suggestions.

1 Introduction

Every academic term, over 6,500 post-secondary institutions across the United States assign course schedules to a total of nearly twenty million students.¹ Based on timing and prerequisites, the number of schedules to which students can possibly be assigned in a given semester is limited. Further, students’ course preferences are heterogeneous. In the face of room size and teaching constraints, university registrars use factors such as student seniority and home department to complete the challenging task of deciding how to allocate seats in over-demanded courses.

In this paper, we approach the task of allocating course schedules to students in the context of a many-to-many matching problem with heterogeneous student preferences and course priorities. We propose a novel *deterministic* allocation mechanism, the *Pseudo-Market with Priorities* (PMP) mechanism, that uses fake money and competitive equilibrium to allocate objects to agents without transfers. This mechanism elicits student preferences for course schedules, assigns almost equal budgets to students, and computes an approximate competitive equilibrium allocation and *priority-specific prices*. The equilibrium prices respect the priority structure by setting a “cutoff” priority level for each course, where students at the cutoff pay a non-negative price, students above the cutoff pay zero price, and students below the cutoff cannot afford the course.

The PMP mechanism delivers a small market-clearing error and has several desirable theoretical properties. In particular, in an environment where each course’s capacity equals the number of assigned seats, the PMP assignment ensures that no student wants to drop a course or enroll in a subset of courses without violating course priorities or capacity constraints (*approximate stability*). Keeping the same number of assigned seats, it is not possible to reassign course seats to a group of students, benefiting some members and hurting none, while ensuring that courses respect course priorities to the same degree as the PMP outcome (*approximate priority-constrained efficiency*). In addition, the outcome limits envy among students; that is, if a student envies another student who has the same or a lower priority level for each course, by removing a single course from the other student’s schedule, envy is eliminated (*schedule envy bounded by a single course*). Finally, we show that students cannot manipulate the PMP mechanism in large populations (*strategy-proof in the large*).

To better guide practitioners in the course allocation process, we compare the PMP mech-

¹Based on [National Center for Education Statistics, 2017-2018](#).

anism with several celebrated mechanisms used in practice. First, we consider the Deferred Acceptance mechanisms with a single tie-breaking and multiple tie-breakings (denoted as DA and DA(m), respectively). These mechanisms allocate courses to students by extending the seminal Gale-Shapley algorithm (Gale and Shapley, 1962) to many-to-many matching with additive student utilities, where priority ties are resolved via either a single tie-breaking or multiple course-specific tie-breakings. The Deferred Acceptance mechanism is considered the gold standard in the matching literature, with many practical applications including matching medical residency to hospitals, school choice, and job matching (see Roth, 2018).

Second, we consider a mechanism closely resembling the one currently used in practice. Students select courses in order of seniority (with ties broken randomly), and priority structures based on student major/department/college are enforced by setting aside seats in each course. Top US universities that closely follow this process include Princeton, Johns Hopkins, Duke, Vanderbilt, Washington University in St. Louis, Columbia, Notre Dame, and Carnegie Mellon.² We call this mechanism the Random Serial Dictatorship with set-asides (RSD).³ In practice, the number of set-asides is set large at the beginning of the course allocation process. After course registration, university departments relax set-aside constraints to ensure full course enrollment. We do not model this period or the process of relaxing set-asides. Instead, we calculate the *optimal set-asides* that a university registrar should set at the beginning of the process if set-asides cannot be changed. Hence, one should interpret our simulation results (described below) as a comparison of the four theoretical mechanisms rather than an evaluation of the actual course allocation mechanism used in practice.

Using data on course allocation from a private institution in the mid-Atlantic region, we calibrate a student utility model and compare the performance of the four mechanisms

²The listed universities are in the top 30 of the [US News Best National University Ranking](#). In general, however, there is some heterogeneity in how university registrars assign course seats to students. Some universities use two passes (e.g., UCLA), allowing students to only register for a limited number of courses in the first pass. Seniority can be based on academic year (e.g., Carnegie Mellon), number of earned credits (e.g., University of North Carolina), or time to graduation (e.g., Washington University in St. Louis). Some universities use priorities (e.g., Dartmouth College) or course reserves (e.g., Vanderbilt) to determine how to enroll students, whereas others delegate the decision to departments (e.g., Cornell). At some universities, only a few classes are oversubscribed (e.g., Caltech), while at others, course allocation is rather congested (e.g., Berkeley).

³The term Random Serial Dictatorship has been used in the matching literature since [Shapley and Scarf \(1974\)](#). The mechanism has little to do with real “dictatorship,” though. The main idea behind the name is that the student with the highest priority is like a dictator who can choose any schedule that he/she wants. Then, the second-highest priority student becomes the dictator, chooses any schedule that he/she wants from the remaining course seats, and so on.

described above using simulations in terms of *student utilities*, *the standard deviation of student utilities* across students, and *the number of students who experience schedule envy*. The first metric measures the satisfaction of students from course allocation. The two other metrics measure the fairness of course allocation.

The simulations show that PMP delivers the highest mean student utility across all years of study. 1223 of 1565 first-year and 734 of 1611 second-year students received different schedules in the PMP and RSD mechanisms, with average improvements in student utility of 9% and 7%, respectively. Changes are smaller for more senior students. 307 of 1422 third-year and 62 of 1425 fourth-year students received a different schedule, with average improvements of 7% and 10%, respectively. The DA and DA(m) mechanisms only slightly improve the mean student utility for years 2-4 compared to the RSD benchmark. For first-year students, DA and DA(m) perform worse than the RSD benchmark. This results from the presence of many course seat reservations for first-year students in the RSD mechanism.

In addition to delivering high student utility, the PMP mechanism decreases the standard deviation of student utilities for almost all years of study compared to the RSD benchmark. The largest reduction occurs for first-year students (11.16%); the effects for second and third-year students are smaller (2.92% and 0.64%, respectively). For fourth-year students, the standard deviation slightly increases (0.05%). The standard deviation of student utilities for the DA mechanism is almost the same as in the benchmark for all years of study (the change is less than 1.5%). The relative values of the standard deviation of student utilities for the DA(m) mechanism are between the values of the DA and RSD mechanisms with the maximum decrease of 9.79% for first-year students. Among all four mechanisms, the PMP mechanism also results in outcomes in which individual students experience the least amount of envy towards students of the same or lower priority. Also, the share of such students is the smallest for PMP outcomes. Overall, our simulations show that the PMP mechanism delivers a more satisfactory and fair solution to the course allocation problem compared to the other three mechanisms.

Literature review. This paper contributes to the literature analyzing the assignment of courses to students. [Hylland and Zeckhauser \(1979\)](#) first proposed the use of fake money and competitive equilibrium to randomly allocate objects to agents without transfers through what are referred to as pseudo-market (competitive) mechanisms. An important advantage of pseudo-markets is that they elicit the participants' cardinal preferences, allowing them

to allocate objects more efficiently than is possible with most ordinal mechanisms (e.g., the deferred acceptance mechanism). [He, Miralles, Pycia, and Yan \(2018\)](#) incorporate a priority structure into pseudo-markets and analyze random mechanisms with an emphasis on unit-demand settings such as school choice.⁴ [Echenique, Miralles, and Zhang \(2021\)](#) and [Nguyen, Nguyen, and Teytelboym \(2021\)](#) also study a pseudo-market solution to random allocation problems with various complex constraints. [Pycia \(2021\)](#) provides an excellent survey of the use of pseudo-markets for random allocations in environments without transfers.⁵

In contrast to the studies cited above, the present paper considers only *deterministic* mechanisms and emphasize many-to-many matching problems such as undergraduate course allocation. For deterministic assignments, [Budish \(2011\)](#) introduces the idea of competitive markets that might not exactly satisfy the market clearing condition. Budish proposes the approximate competitive equilibrium from equal incomes mechanism, which is strategy-proof in the large, finds an allocation that bounds student envy by one course, and is approximately Pareto-efficient. This mechanism was successfully implemented at Wharton Business School and Columbia Business School (see [Budish, Cachon, Kessler, and Othman, 2017](#); [Budish and Kessler, 2022](#)).⁶ In this context, a key contribution of the present study is to extend Budish’s mechanism to many-to-many matching settings with course priorities, an important feature of the undergraduate course allocation problem.

Two recent papers also analyze deterministic assignments in many-to-many settings. In a recent contribution, [Nguyen and Vohra \(2022\)](#) establish the existence of a competitive equilibrium when all agent preferences satisfy a novel property of geometric substitutes, which is a strict generalization of gross-substitutes property (see [Kelso and Crawford, 1982](#)). When there is an upper limit on the number of goods k that can be acquired by an agent, by introducing a perturbation in agent utilities, they also establish that there always exists an

⁴[He, Miralles, Pycia, and Yan \(2018\)](#) also explain how their results can be extended to random allocations in many-to-many settings with additive utility.

⁵[Budish, Che, Kojima, and Milgrom \(2013\)](#) and [Akbarpour and Nikzad \(2020\)](#) investigate the implementation of a random allocation mechanism by randomizing over feasible integer allocations. [Nguyen, Peivandi, and Vohra \(2016\)](#) relax the requirement of strategy-proofness and design a random allocation mechanism that is ordinally efficient, envy-free, and weakly strategy-proof. This mechanism might also violate each course’s capacity constraint, but by no more than the size of one student’s course schedule. These appealing theoretical properties led to the implementation of the mechanism at the Technical University of Munich (see [Bichler and Merting, 2021](#)).

⁶[Budish and Cantillon \(2012\)](#) draw on theory and field data to argue that the properties of a non-strategy proof course allocation mechanism used at Harvard Business School are superior to those of the serial dictatorship. Also, [Rusznák, Biró, and Fleiner \(2021\)](#) analyze the course allocation mechanism at the largest university in Hungary.

approximate competitive equilibrium with good-by-good clearing bound equal to $k - 1$. The good-by-good bound dominates the aggregate market-clearing bound of Budish (2011) when preferences are close substitutes, whereas the aggregate market-clearing bound dominates when student preferences are close complements. Though our main theoretical results follow the aggregate market-clearing bound approach, we ensure in our simulations to find an approximate competitive equilibrium allocation in which no courses are oversubscribed by more than $k - 1$ seats.

As in the previous paper, Lin, Nguyen, Nguyen, and Altinkemer (2022) analyze deterministic allocation mechanisms through approximate competitive equilibrium with the good-by-good bound. They establish the existence of an approximate competitive equilibrium that satisfies stability and fairness properties similar to the one analyzed in this paper. We additionally establish that our mechanism is approximately priority-constrained efficient and strategy-proof in the large (see Azevedo and Budish, 2019), while Lin, Nguyen, Nguyen, and Altinkemer (2022) contributes to the analysis of matchings under complex feasibility constraints (Echenique, Miralles, and Zhang, 2021; Kamada and Kojima, 2023). We also differ in our applications. While they provide an interesting analysis of how to reassign season tickets to families using artificially generated data, we address a larger-scale problem of allocating course seats to students based on real-world university data.

Our paper also compares the performance of our pseudo-market mechanism with the performance of the deferred-acceptance algorithm with single- and multiple tie-breaking in a many-to-many matching setting. Both versions of the mechanism were previously studied in the matching literature in the context of schools choice (Abdulkadiroğlu, Pathak, and Roth, 2009; Abdulkadiroğlu, Che, and Yasuda, 2015; Erdil and Ergin, 2008). These papers combine theory and simulations to show that single tie-breaking has better welfare properties.

The remainder of the paper is organized as follows. We introduce the model in Section 2 and investigate the properties of the PMP mechanism in Section 3. Simulation results are presented in Section 4. We offer concluding remarks in Section 5. Appendices A and B present proofs, Appendix C describes how we calibrate a simple student utility model, Appendix D provides additional simulations. Finally, Supplementary Materials provides more details on calibration and presents additional results for a common priorities setting.⁷

⁷Supplementary Materials can be accessed online through the link: <https://www.dropbox.com/s/vnkug7eqontpj7p/Supplementary-Materials.pdf?dl=0>.

2 Environment

Course allocation is a many-to-many matching problem described by the tuple $(\mathcal{S}, \mathcal{C}, Q, V, \mathfrak{R})$.

- $\mathcal{S} = \{1, \dots, S\}$ is a set of students; in reference to students, we use she/her pronouns.
- $\mathcal{C} = \{1, \dots, M\}$ is a set of courses.
- $Q = (q_1, \dots, q_M)$ is a vector of course capacities; each c can enroll at most q_c students.
- $V = (\succsim_1, \dots, \succsim_S)$ is a vector of student preferences over course schedules. Students are typically restricted to a set of permissible course schedules due to factors such as course meeting times and prerequisites. Each student can take at most one seat in any course and k total courses. We assume that these restrictions are incorporated into student preferences, and, for simplicity of exposition, that $1 \leq k \leq M/2$. Our specification permits general substitutability and complementarity of preferences over courses. We also assume that student preferences over permissible course schedules are *strict*.
- $\mathfrak{R} = \{r_{s,c}\}_{s \in \mathcal{S}, c \in \mathcal{C}}$ is a course priority structure. Each $r_{s,c} \in \mathcal{R} \equiv \{1, \dots, R\}$, $R \geq 1$, specifies the level of priority for each student s and course c , with *a larger number meaning a higher level of priority*. The priority levels need not be distinct and could be the same for multiple students.

We consider deterministic allocations of courses to students. An allocation $x = (x_s)_{s \in \mathcal{S}}$ assigns a course schedule to each student, where $x_s \subseteq \mathcal{C}$ for each $s \in \mathcal{S}$. For ease of notation, we view the schedule x_s as both a set of courses assigned to student s and a vector from the set $\{0, 1\}^M$. Allocation x is *feasible* if $\sum_{s \in \mathcal{S}} x_{s,c} \leq q_c$ for each $c \in \mathcal{C}$; that is, no course is assigned a greater number of students than its capacity if a feasible allocation. Allocation x is *individually rational* if, for any $s \in \mathcal{S}$, $x_s = \max_{\succeq_s} \{x'_s : x'_s \subseteq x_s\}$; that is, student s does not want to drop any of the assigned courses. A pair (s, C) of student $s \in \mathcal{S}$ and a subset of courses $C \subset \mathcal{C}$ is a *block* of x if $C = \max_{\succeq_s} \{x'_s : x'_s \subseteq x_s \cup C\}$, $C \neq x_s$, and for each $c \in C$ such that $c \notin x_s$, we have that either there is a student $s' \in \mathcal{S}$ with $c \in x_{s'}$ and $r_{s',c} < r_{s,c}$ or $\sum_{s \in \mathcal{S}} x_{s,c} < q_c$; that is, there is a block if a student is willing to replace the current schedule with a new one, and for each course in the new schedule that was not previously assigned, there is a lower priority student who is assigned a seat in the course or the course has available seats.

We evaluate allocations based on stability, efficiency, and fairness. An allocation x is *stable* if it is *feasible*, *individually rational*, and admits no *blocks* (see Roth and Sotomayor, 1990; Echenique and Oviedo, 2006). As we explain in the next section, we analyze approximate market equilibria, for which the market-clearing condition is satisfied with a small error. To account for that, we use an approximate version of stability.

Definition 1. An allocation x is **approximately stable** if it is stable in the environment with a capacity of $q'_c = \sum_{s \in S} x_{s,c}$ for each course $c \in \mathcal{C}$.⁸

The relationship between approximate stability and stability is similar to the relationship between Pareto efficiency and approximate Pareto efficiency introduced by Budish (2011). An allocation y *Pareto dominates* an allocation x if there is at least one student who strictly prefers her course schedule in y and all other students weakly prefer their course schedules in y . If y Pareto dominates x , we say that x has a *Pareto improvement*. An allocation is *Pareto efficient* if there are no Pareto improvements. An allocation is *approximately Pareto efficient* if there is no Pareto improvements with the number of assigned seats for each course equal to the number of assigned seats in the current allocation.

In our setting, a meaningful notion of efficiency should also account for course priorities. Following Schlegel and Mamageishvili (2020), we extend course priorities over individual students to course priorities over subsets of students using the first-order stochastic dominance relation. We say that y dominates x for course c , and write $y_c \succeq_c x_c$, if for all $r \in \mathcal{R}$ we have that $\sum_{s:r_{s,c} \geq r} y_{s,c} \geq \sum_{s:r_{s,c} \geq r} x_{s,c}$. Using this definition, we consider the following criterion of constrained efficiency.

Definition 2. An allocation x is **approximately priority-constrained efficient** if for each assignment y that Pareto dominates x and has the same total number of assigned seats— $\sum_{s \in S} y_{s,c} = \sum_{s \in S} x_{s,c}$ for all $c \in \mathcal{C}$ —there is a course $c' \in \mathcal{C}$ for which $y_{c'} \not\preceq_{c'} x_{c'}$.

In environments without priorities, Definition 2 reduces to the approximate Pareto efficiency used by Budish (2011). The requirement that any Pareto improvement must have the same total number of seats assigned accounts for the possibility that the market-clearing condition

⁸A related concept is justified course envy. Allocation x *prevents justified course envy* if there are no students $s, s' \in S$ and course $c \in \mathcal{C}$ such that $r_{s,c} > r_{s',c}$, $c \notin x_s$, $c \in x_{s'}$ and $c \in \max_{\succeq_s} \{x'_s : x'_s \subseteq x_s \cup c\}$. The absence of justified course envy prevents envy towards students of a lower priority, but does not account for individual rationality and the possibility of envy towards several students assigned to several courses.

is satisfied with a small error. If this requirement is dropped, our definition coincides with the priority-constrained efficiency introduced by [Schlegel and Mamageishvili \(2020\)](#).

The seminal measure of fairness is envy-freeness, as introduced by [Foley \(1967\)](#): an allocation x prevents *schedule envy* if there are no students $s, s' \in S$ such that $x_{s'} \succ_s x_s$. However, without using lotteries, we cannot hope to allocate courses in a completely envy-free way among students (e.g., if two students have the same priorities and the same preferences over all courses). Extending the notion originally proposed by [Budish \(2011\)](#), we consider a more permissive concept of *schedule envy bounded by a single course* among students of the same or lower levels of priority (see also [Lin, Nguyen, Nguyen, and Altinkemer, 2022](#)).

Definition 3. *An allocation x has **schedule envy bounded by a single course** among students of the same or lower levels of priority if, for any $s, s' \in \mathcal{S}$ such that $r_{s,c} \geq r_{s',c}$ for all $c \in \mathcal{C}$, either $x_s \succeq_s x_{s'}$ or there exists some course c^* such that $x_s \succeq_s (x_{s'} \setminus \{c^*\})$.*

This definition provides a criterion of fairness for students weakly ordered in the same way across all course priority orders and bounds envy among such students in a minimal way.

Allocations are found through mechanisms that systematically elicit student preferences $(\succeq_s)_{s \in \mathcal{S}}$ over course schedules. Evidence from business schools demonstrates that mechanisms requiring strategic play on behalf of students can lead to large complications with efficiency (see [Budish and Cantillon, 2012](#); [Budish and Kessler, 2022](#); [Sönmez and Ünver, 2010](#)).

Definition 4. *A course allocation mechanism is **strategy-proof** if there is no student s , who, by reporting manipulated preferences \succeq'_s , receives an allocation she strictly prefers to the course schedule she would get by reporting \succeq_s .*

We are interested mainly in mechanisms that are *strategy-proof in the large*. To avoid unnecessary heavy notation early in the paper, we introduce this concept formally in the proof of [Theorem 5](#). Here, we only provide its interpretation. A mechanism that is strategy-proof in the large is strategy-proof in a limit market in which each student regards the “prices” in pseudo-market mechanisms as exogenous to her report (see [Azevedo and Budish, 2019](#); [Budish, 2011](#)).⁹ We also assume that strategic play is a concern only for the students, as course priorities are typically set based on commonly observable factors such as student seniority or department.

⁹See [He, Miralles, Pycia, and Yan \(2018\)](#) for asymptotic incentive compatibility in random assignment matching models.

3 Pseudo-Market with Priorities

In this section, we present our novel mechanism, which allocates courses to students by extending the concept of approximate competitive equilibrium from equal incomes of [Budish \(2011\)](#) to settings with course priorities. For this purpose, we allocate to each student $s \in \mathcal{S}$ a budget of fake money b_s^* . To obtain good fairness properties, we assume that budgets are almost the same, with $1 \leq \min_s b_s^* \leq \max_s b_s^* < \bar{b} \equiv 1 + \beta$ for some small $\beta > 0$. The parameter β is interpreted as the maximum allowable budget inequality across students. We also allow for slack in the market-clearing condition, which is bounded by $\alpha \geq 0$. In addition, to respect the course priority structure, we require a special structure on priority-specific prices, as defined below.

Definition 5. *An allocation x^* , prices p^* , and budgets b^* constitute an (α, β) -Pseudo-Market Equilibrium with Priorities if*

1. $x_s^* = \max_{\succsim_s} \left\{ x'_s \subseteq \mathcal{C} : \sum_{c \in \mathcal{C}} p_{c,r_{s,c}}^* x'_c \leq b_s^* \right\}$ for each student $s \in \mathcal{S}$.

2. For each $c \in \mathcal{C}$, there is a cutoff priority level r_c^* such that $\sum_{\{s \in \mathcal{S} : r_{s,c} > r_c^*\}} x_{s,c}^* < q_c$ and

$$p_{c,r}^* \in \begin{cases} \{0\} & r > r_c^* \\ [0, \bar{b}) & r = r_c^* \\ [\bar{b}, +\infty) & r < r_c^* \end{cases} \quad (1)$$

3. $\|z^*\|_2 \leq \alpha$, where $z^* = (z_1^*, \dots, z_M^*)$ and

- (a) $z_c^* = \sum_s x_{s,c}^* - q_c$ if $p_{c,1}^* > 0$,

- (b) $z_c^* = \max(\sum_s x_{s,c}^* - q_c, 0)$ if $p_{c,1}^* = 0$.

4. $1 \leq \min_s b_s^* \leq \max_s b_s^* < \bar{b} \equiv 1 + \beta$.

Unlike [Budish \(2011\)](#), the above definition of competitive market equilibrium allows for course prices to depend on priority levels, so that the vector of prices is $p^* = \{p_{c,r}^*\}_{c \in \mathcal{C}, r \in \mathcal{R}} \in \mathbb{R}^{MR}$. When $R = 2$, there are only two levels of priority: prioritized students and all others. In undergraduate course allocation, department-specific priorities are an example of such a priority structure. Condition (1) ensures that the equilibrium allocation satisfies the priority

structure; that is, there exists a cutoff level of priority such that higher priority students can obtain the course seats for free, students at the cutoff level of priority face non-negative prices, and lower priority students cannot afford seats in the course. A similar condition on prices appeared in [He, Miralles, Pycia, and Yan \(2018\)](#) in the context of random allocation with priorities. Also, condition (1) ensures that the equilibrium allocation *prevents justified course envy*, that is there is no student allocated a seat in a course and another student who is not allocated a seat in the course, has a higher priority for the course, and wants to add the course to her schedule.

The market-clearing error for a course depends on its price for the lowest priority $r = 1$, guaranteeing that under-demand is only counted as an error if the price for the lowest priority group is positive. This requirement is a non-trivial extension of [Budish \(2011\)](#)'s definition to settings with course priorities. Alternatively, if we try to define the market-clearing error per priority level, one could use the cutoff level to artificially lower the market-clearing error. For example, consider a situation with the price of zero at and above the cutoff level of priority and a price of \bar{b} below the cutoff level. In this case, the market-clearing error is zero according to the alternative definition. At the same time, the unfilled course seats should be counted towards the market-clearing error, as they could be potentially eliminated if the price below the cutoff level were decreased.

The condition $\sum_{\{s \in \mathcal{S}: r_{s,c} > r_c^*\}} x_{s,c}^* < q_c$ also did not appear in the previous literature. Note that if we were to allow r_c^* to satisfy $\sum_{\{s \in \mathcal{S}: r_{s,c} > r_c^*\}} x_{s,c}^* = q_c$, a price $p_{c,r_c^*}^*$ can be selected such that no student at the cutoff level of priority is able to take a seat in course c . Then, raising the cutoff level of priority to $r_c^* + 1$ does not lead any student to lose their seat in the course. In turn, requiring that $\sum_{\{s \in \mathcal{S}: r_{s,c} > r_c^*\}} x_{s,c}^* < q_c$ guarantees the cutoff cannot be raised without causing under-demand for the course.

3.1 Pseudo-Market with Priorities Mechanism

[Budish \(2011\)](#) establishes the existence of an approximate competitive equilibrium from equal incomes under the upper bound on the market-clearing error. The upper bound on the market-clearing error is proportional to the square root of the dimension of the price space. In the presence of R course priorities for each of M courses, the dimension of a price space is MR . This could potentially require a large market-clearing error for a Pseudo-Market Equilibrium with Priorities to exist. The main idea of the proof of [Theorem 1](#) is

that we can reduce the effective dimension of the price space from MR to M , where we can establish the existence of a Pseudo-Market Equilibrium with Priorities. In particular, we consider the following parameterization. For each course c and priority r , we consider only price vectors $p \in \mathbb{R}^{MR}$ that satisfy

$$p_{c,r}(t) = \max(t_c - (r - 1)\bar{b}, 0),$$

for some $t \in \mathbb{R}^M$.¹⁰ We then establish the existence of an equilibrium in the smaller dimensional space. It is clear that for equilibrium vector t^* one could find cutoff levels such that the corresponding price vector p^* satisfies condition (1).

Theorem 1 (Existence). *For any $\beta > 0$, there exists a $(\sqrt{kM/2}, \beta)$ -Pseudo-Market Equilibrium with Priorities.*¹¹

Theorem 1 establishes the existence of the Pseudo-Market Equilibrium with Priorities, which is an extension of Theorem 1 in Budish (2011) to environments with course priorities. The equilibrium has the same upper bound on the market-clearing error as an approximate competitive equilibrium from equal incomes in the same environment without priorities.

Now, using the Pseudo-Market Equilibrium with Priorities, we introduce the Pseudo-Market with Priorities (PMP) mechanism and investigate its properties in terms of stability, efficiency, fairness, and strategy-proofness.

Definition 6. *The Pseudo-Market with Priorities mechanism with market-clearing error α and budget inequality β is defined through the following steps:*

1. *Each student s reports preferences \succsim_s over permissible course schedules.*
2. *Each student s is assigned a distinct budget b_s^* in $[1, 1 + \beta]$.*
3. *Compute an (α, β) -Pseudo-Market Equilibrium with Priorities (x^*, p^*, b^*) . Allocate courses to students according to x^* .*

In general, the PMP mechanism depends on the level of allowable market-clearing error α and the level of budget inequality β . For results that hold for all non-negative $\alpha \geq 0$ and

¹⁰We can think of this pricing scheme as if there is only one price per course, but each priority level $r \in \mathcal{R}$ is entitled to a rebate of $(r - 1)\bar{b}$. We thank Ran Shorrer for suggesting this intuitive interpretation.

¹¹Budish (2011) shows the existence of an approximate market equilibrium from equal incomes with the upper bound $\sqrt{\min\{2k, M\}M/2}$. As we assume $k \leq M/2$, the upper bound reduces to $\sqrt{kM/2}$.

$\beta \geq 0$, we avoid this dependence in our exposition. We will be specific when any restrictions are necessary.

We first establish that the PMP mechanism results in an *approximately stable* course allocation. The individual maximization condition in Definition 6 ensures that the course allocation is individually rational. Also, the condition on priority-specific prices prevents the possibility of schedule blocking if courses are regarded as being at their full capacity.

Theorem 2 (Stability). *The outcome of every Pseudo-Market with Priorities mechanism is approximately stable.*

Proof. Let (x^*, p^*, b^*) be an outcome of the Pseudo-Market with Priorities mechanism. The individual maximization condition in Definition 6 implies that x_s^* is individually rational for each student $s \in \mathcal{S}$. Consider now an environment with course capacities $q'_c = \sum_{s \in \mathcal{S}} x_{s,c}^*$ for each $c \in \mathcal{C}$. Hence, all courses are at their full capacities and x^* is feasible in the new environment. In addition, the only possibility for the existence of a block (s, C) is that $C = \max_{\succeq_s} \{x'_s : x'_s \subseteq x_s \cup C\}$, $C \neq x_s^*$, and for each $c \in C$ with $c \notin x_s^*$, there is a student $s' \in \mathcal{S}$ with $c \in x_{s'}^*$ and $r_{s',c} < r_{s,c}$. The individual maximization in Definition 6 implies that there must exist at least one course $c \in C$ with $c \notin x_s^*$ such that $p_{c,r_{s,c}}^* > 0$; otherwise x_s^* should be replaced with C in the maximization problem of student s as we assume that student preferences over course schedules are strict. Condition (1) on equilibrium prices then implies that for any student s' with a lower priority for course c —that is $r_{s',c} < r_{s,c}$ —we must have $p_{c,r_{s',c}}^* \geq \bar{b}$. Therefore, we obtain a contradiction: we must have $c \notin x_{s'}^*$, as an allocation x^* cannot possibly assign a student s' a seat in course c . Hence, there are no blocks, and the outcome of every Pseudo-Market with Priorities mechanism is approximately stable. \square

The main reason that the outcome of the PMP mechanism is approximately stable, but not stable, is the possibility of undersubscribed and oversubscribed courses. In particular, there may be students who want to take courses with unfilled seats and oversubscribed courses may want to drop some students from their assignments. However, as Theorem 1 shows, such instances are rare.

The presence of course priorities prevents the PMP mechanism from being Pareto efficient. This is similar to how stable allocations in one-sided matching markets might not be Pareto efficient (see Roth and Sotomayor, 1990). We illustrate this in the following example.

Example 1. Let us consider an economy with two students $\mathcal{S} = \{1, 2\}$ and two courses $\mathcal{C} = \{A, B\}$, each with a capacity of one. Student preferences are $1 : A \succ B$ and $2 : B \succ A$. Student budgets are $b_1^* = 1$ and $b_2^* = 1 + \beta$ for some $0 < \beta < 1$. We assume that course priorities are the opposite of student preferences, with $r_{1,B} = r_{2,A} = 2$ and $r_{1,A} = r_{2,B} = 1$. Both price vectors $p_A^* = p_B^* = (2, 1)$ with priority cutoffs $r_A^* = r_B^* = 2$ and allocation

Student	$x_{s,A}^*$	$x_{s,B}^*$
1:	0	1
2:	1	0

constitute an exact market equilibrium with course priorities. However, the allocation is not Pareto efficient, as both students would be made better off by exchanging their assigned seats.

The next result shows that the outcome of the PMP mechanism satisfies a form of constrained Pareto efficiency. The PMP mechanism's allocation cannot be Pareto dominated by another allocation in which all courses have the same total number of seats assigned and priorities are respected to at least the same degree; that is, for each course, the distribution of priorities for the allocation cannot first-order stochastically dominate the distribution of priorities for the PMP outcome (see Definition 2).

Theorem 3 (Efficiency). *The outcome of the Pseudo-Market with Priorities mechanism is approximate priority-constrained efficient.*

The proof of the above result resembles the proof of [Schlegel and Mamageishvili \(2020\)](#) for single-unit settings with random allocations and is postponed to Appendix A. Note that we assume each agent has strict preferences over course schedules. Hence, for the above result, we do not require the condition that a student choose the cheapest course schedule when multiple course schedules are optimal as in [Miralles and Pycia \(2021\)](#) and [Schlegel and Mamageishvili \(2020\)](#).

We also want to comment on the relationship between approximate priority-constrained efficiency and approximate stability. Though these concepts are related, they do not imply each other. If we assume in Example 1 that both students have the same priority for both courses, then there are two (approximately) stable allocations: $x_1 = \{A\}$, $x_2 = \{B\}$ and $x_1 = \{B\}$, $x_2 = \{A\}$. However, only the first allocation is approximately priority-constrained efficient. The concept of approximately priority-constrained efficiency selects among approximately stable allocations in the presence of weak priorities, favoring more Pareto optimal

outcomes. In turn, consider an assignment of two courses to one student, in which the student wants to take only one of the courses and drop the other course. The assignment is approximately priority-constrained efficient, but not approximately stable. Example S1 in [Supplementary Materials](#) presents an approximately priority-constrained efficient allocation that is individually rational, but not approximately stable.¹²

The fairness of a course allocation should also account for the course priority structure. Theorem 2 precludes the possibility that the PMP outcome has *justified course envy*, guaranteeing that there is no student left wanting to get a course for which she has a higher priority than some student who is assigned a seat in the course. However, the priorities in undergraduate course allocation are weak, with hundreds of students enjoying the same level of priority for many courses. The next result addresses the question of fairness among students who could be weakly ordered in terms of their course priorities for all courses.

Theorem 4 (Fairness). *If $\beta \leq \frac{1}{k-1}$, the Pseudo-Market with Priorities mechanism results in an allocation that has schedule envy bounded by a single course among students of the same or lower levels of priority.*

The result of Theorem 4 extends the one established in [Budish \(2011\)](#) for settings where all students are on the same priority level. In the setting with priorities based on seniority and student’s department, the above result implies that the PMP mechanism with small budget inequality produces an allocation in which any two students from the same department and the same year of study have schedule envy bounded by a single course.

We also establish that students have almost no incentive to manipulate the PMP mechanism in large populations, which is the case for undergraduate course allocation, where student bodies can consist of thousands of students.

Theorem 5 (Strategy-Proofness). *The Pseudo-Market with Priorities mechanism is strategy-proof in the large.*

To alleviate the burden of extra notation, we define the concept of a mechanism that is strategy-proof in the large in the proof of Theorem 5 in Appendix A. To prove this result, we show that the PMP mechanism is a *semi-anonymous* mechanism that is *envy-free but for tie-breaking* (see [Azevedo and Budish, 2019](#); [Kalai, 2004](#)). Intuitively, students can be

¹²To be precise, Example S1 in [Supplementary Materials](#) presents an allocation that is Pareto efficient, but not stable.

partitioned into groups based on their course priorities, where students in the same group are at the same level of priority in all courses. The PMP mechanism treats the group members in the same way; that is, they face the same course prices and the same random lottery over budgets. Hence, the PMP is a semi-anonymous mechanism. Moreover, the PMP mechanism is envy-free but for tie-breaking, as a student with a larger budget cannot envy an allocation of any other student within the same group with a lower budget. The statement the theorem then follows from Appendix C of [Azevedo and Budish \(2019\)](#), which proves that any semi-anonymous mechanism that is envy-free but for tie-breaking is strategy-proof in the large.¹³

Last, we relate the PMP mechanism to some mechanisms discussed in the literature. In the case of no priorities, our properties reduce to those derived in [Budish \(2011\)](#). With strict priorities, it is of interested to compare the PMP mechanism to the deferred acceptance mechanism. If students have additive preferences, the student-proposing deferred acceptance mechanism results in the student-optimal stable matching (see [Roth, 1984](#)), which can be supported with a Pseudo-Market Equilibrium with Priorities where the market clears exactly, the cutoff level of priority is at the lowest-priority student to receive a seat, and the price of each course is zero at the cutoff. In fact, any stable allocation can be supported by a Pseudo-Market Equilibrium with Priorities.¹⁴ This could be problematic, as one could inquire how to select among several possible equilibria. In large markets, however, we expect the set of stable matching to be small (see [Kojima and Pathak, 2009](#)) and, hence, the PMP mechanism outcomes should not differ much from each other.

When each student wants to enroll only in one course, our setting reduces to the unit-demand variation of the problem, which is known in the literature as school choice. The standard mechanisms in school choice are ordinal ([Abdulkadiroğlu and Sönmez, 2003](#)) or elicit only restricted information about cardinal preferences from agents ([Abdulkadiroğlu, Che, and Yasuda, 2015](#)). The set of ordinal mechanisms is restrictive, and one may be implementing Pareto-dominated assignments using these mechanisms. [He, Miralles, Pycia, and Yan \(2018\)](#) show how one can extend the pseudo-market approach, which elicits the cardinal preferences of students, to school choice settings with *random allocations*. The PMP mechanism studied in this paper complements their results in school choice settings by analyzing

¹³Alternatively, one can consider a continuum replication of an economy with course priorities. By leveraging on the price structure as in Theorem 1, the steps of Theorem 4 of [Budish \(2011\)](#) can be adapted to obtain the result. The alternative proof is available upon request.

¹⁴See [Miralles and Pycia \(2021\)](#) for a related result showing that every efficient assignment can be decentralized through prices in random allocation settings.

the pseudo-market mechanisms with *deterministic allocations*. While deterministic allocations are preferable from a practical market design point of view, they are also associated with additional complications regarding the existence of an equilibrium. The existence of equilibrium requires market clearing conditions to be satisfied approximately, leading many theoretical results to hold only in an approximate sense.

4 Simulations

In this section, we analyze the performance of the Pseudo-Market with Priorities (PMP) mechanism using simulations based on course allocation data from a private institution in the mid-Atlantic region. For this purpose, we compare the PMP mechanism with the Deferred Acceptance with single and multiple tie-breakings and the Random Serial Dictatorship with set-asides. The two variants of the Deferred Acceptance mechanism were previously used in the allocation of students to public schools and doctors to hospitals, while variants of the Random Serial Dictatorship are used in undergraduate course allocation in many U.S. universities. We begin with the description of the course allocation data that we use for our simulations.

4.1 Course and Student Data

Our simulations utilize undergraduate course allocation data for one semester (Spring 2018). The undergraduate population consists of 6023 students from 7 colleges, 41 departments, and 5 classes of students based on year of study.¹⁵ We merge the data for the fourth and fifth-year students into one year of study, as some colleges have only a few fifth-year

College	A	B	C	D	E	F	G
# of students	853	1642	259	1274	745	741	509
# of courses	180	84	12	269	88	84	39

TABLE 1: The number of students and courses in various colleges.

¹⁵We consider only undergraduate students who were enrolled full-time. There are also some graduate students and exchange students who take undergraduate courses. We exclude these students from the data.

students. The data also contains information about 756 courses across 42 departments.¹⁶ Table 1 presents information on the number of students and courses offered by different colleges. We use capital letters to denote colleges. Colleges differ in the size of their student population and the number and type of courses they offer. For example, some colleges serve the whole university, offering many general education courses in topics such as math and science, whereas some colleges offer courses mainly for their students.

In our simulations, we use the data on maximum enrollment, actual enrollment, and the number of course seat reservations. Table 2 presents several quantiles of the distribution of these characteristics. The table shows that a typical median quantile course has a capacity for 25 students, the median actual enrollment is 15, and the median number of reserved seats is 3. Overall, 33 455 course seats were available in all courses during the considered term. Among them, 23 369 seats got occupied by students in the actual enrollment. 13 922 total seats were reserved as set-asides at the start of the course enrollment process.

Course reservation (or set-asides). Course seat reservations can be done at college, department, major, or year of study levels. There are no major-level course reservations in the data. We also convert the college-level reservations to department-level reservations; instead of a college, we list all college departments for a course reservation. Then, we unite all department reservations corresponding to the same year of study into one reservation. This significantly simplifies the course reservation data.¹⁷ For each year of study, we have almost a binary reservation structure: a set of departments and several course seat reservations.

Quantile	10%	25%	50%	75%	90%
Max Enrollment	8	15	25	50	98
Act Enrollment	3	7	15	35	72
# of Course Reserves	0	0	3	20	53

TABLE 2: Quantiles of maximum enrollment, actual enrollment, and number of reserved seats.

¹⁶There are different numbers of departments for courses and students. Some departments are used only for course classification (i.e., no students are assigned to these departments) and some departments are used only for student classification (i.e., no courses are assigned to these departments). We also restrict our sample to semester-long courses.

¹⁷This simplification ignores some redistribution concerns of course reservations. Several department-level reservations could aim to obtain a balanced class enrollment. Our simulations do not account for this aspect of the actual course allocation process.

Course	Departments	Year of Study	#Reserved Seats
Economics	Dept 1, Dept 2	Year 1	25
Economics	All	Year 4	2
Economics	Dept 2, Dept 3	All	30

TABLE 3: Some examples of course reservations.

The only exception is reservations done for all years of study. Table 3 presents examples of course reservation types observed in the data. The criteria for course reservations are not mutually exclusive and could potentially intersect. If there is a student who is eligible for both a year-specific reservation and a reservation for all years of study (labeled “All” in Table 3), the student is assigned to the year-specific reservation without deducting the number of reserved seats for the reservation for all years of study. If no reserved seats are left in the year-specific reservation, the student is assigned a seat in the reservation for all years of study.

Binding maximum enrollment capacity. The percentage of courses that are at or above the maximum enrollment capacity is an important indicator of how much student interests overlap and the need for a design of the course allocation system. Table 4 shows that 11.2% of all courses are at the maximum enrollment capacity. This number does not reflect, however, a large heterogeneity across the colleges. The percentage of courses that are at or above the binding capacity within each college ranges between 5.6% and 33%.

Table 4 also shows that the enrollment capacity is not a strictly binding constraint for more than 7% of the courses. The enrollment capacity in such courses is intended to ensure proper class dynamics and balanced course sections. For example, it is difficult to teach business communication or drama classes of a large size. Also, it might be difficult for one instructor to teach one section with more than sixty students and another with fewer than

# seats above max capacity	≥ 0	≥ 1	≥ 2	≥ 3	≥ 4	≥ 5
% of courses	11.2%	7.3%	4.1%	3.3%	2.5%	1.9%

TABLE 4: The percentage of courses at or above the maximum enrollment capacity.

twenty students of the same course in the same semester. In these situations, the maximum enrollment constraint is flexible and such classes could accommodate students above their capacity. This fits our treatment of the enrollment constraints in the PMP mechanism, which allows some enrollment above course capacity. In some courses, though, the maximum enrollment capacity is strictly binding. These situations are typically associated with the physical limitations of a classroom assigned to the course. In our simulations, we set each course capacity equal to the maximum enrollment capacity observed in the data and the number of actually assigned students.

Student utility. The course allocation process that is currently used by the university closely follows the Random Serial Dictatorship with set-asides. However, the data does not contain information on students’ utilities over courses. Hence, we need to recover this information. We accomplish this by exploiting the information on actual student-course enrollment, the order in which students register for courses, and the information on course reservations.¹⁸

One major obstacle in recovering student utilities is that course reservations do not remain constant during the course allocation process. After the first week of course registration, the demand for courses from students becomes more or less clear. To ensure full course enrollment, departments responsible for course allocation start relaxing course reservations by admitting students from waiting lists. While there are best practices for handling waiting lists, the actual enrollment is done solely at each department’s discretion.¹⁹ As a result, the final enrollment could violate individual course reserves set at the beginning of the process. To avoid this problem, we adjust the actual course reserves to ensure that the course reserves are consistent with the student-course allocation observed in the data. We use the celebrated *Hall’s Marriage Theorem* to identify course reserve violations. We explain how we apply the theorem and minimally adjust course violations in Appendix C. Using *the adjusted course reserves* and the final student-course allocation, we calibrate a simple student utility model.

Our student utility model assumes that a student’s utility from taking a course depends

¹⁸The order we observe reflects the earliest time for a student to register for a course. This time does not always coincide with the earliest time each student could have registered for the course (the true lottery number in the random serial dictatorship). For example, it is possible that the course registration window for a student opened earlier in the morning, but the student decided to register only in the evening. Some students may have their registration windows delayed if they have outstanding financial obligations in front of the university.

¹⁹The best practices include prioritizing senior students and students who have majors or minors in that department.

only on her college, her year of study, and the course’s college plus student- and course-specific idiosyncratic component; that is, for a student from college a and year y and course c from college a' ,

$$u_{sc} = \theta_{aya'} + \varepsilon_{sc}, \tag{2}$$

where $\theta_{aya'}$ is a fixed utility component and ε_{sc} is a random utility component. We assume that $\varepsilon_{sc} \sim N(0, \sigma)$ are independently and normally distributed random variables with zero mean and variance σ . Student utilities are additive across courses. As a normalization, we consider the value of the outside option from not taking a course to be zero and the standard deviation of the noise parameter to be $\sigma = 1$.²⁰

In addition, we assume that each student’s choice set is limited to 80 courses. This is done for computational purposes and is further motivated by the practical observation that any given student typically only considers taking a subset of all offered courses during the semester (see also [Budish and Cantillon, 2012](#); [Diebold and Bichler, 2017](#)). We draw these 80 courses randomly among those taken by at least one student from the same college and year. Moreover, the probability that a course from college a' is drawn equals the share of students from the same college-year pair enrolled in courses in college a' in the actual data.²¹ We emphasize that students might not have all 80 utility entries positive. Instead, some utility draws may be below the value of the outside option.

The actual course allocation process (that resembles Random Serial Dictatorship with set-asides) does not necessarily result in a stable allocation. This creates a substantial complication, as most of the empirical matching literature relies on stability to identify the preferences of agents (see, e.g., [Agarwal and Budish, 2021](#); [Fack, Grenet, and He, 2019](#)). To overcome this complication, we use a simulated method of moments. Using the order in which students choose courses, adjusted course reserves, and the final course allocation in the university data, we calibrate θ to match the number of courses in each college taken by students from each college-year pair ($7 \times 4 \times 7$ moments in total; the same number as the number of calibrated parameters). Appendix C describes the calibration procedure in detail. In the appendix, we also discuss some interesting patterns of the calibrated values of

²⁰In Appendix C, we also analyze an alternative normalization with a variable outside option that can depend on the student’s college and the year of study. We also present the welfare comparison of the four mechanisms using this normalization in Appendix D.1.

²¹If students from a given college-year pair took fewer than ten courses in a college, the student choice sets are enlarged to include courses of the same or next level in the same department.

θ presented in Table C2.

We want to acknowledge some limitations of our student utility model. The model does not allow systematic differences in popularity across courses within a college. We also do not account for scheduling or other constraints that are present parallel to a student’s willingness to take a course. The additive model specification could also be restrictive. Hence, one should be careful in interpreting the calibrated parameter values. Though we believe a more detailed student preference calibration might be possible, this is not the main focus of our paper. Rather, we want to obtain a sensible student utility model to evaluate the performance of the four course allocation mechanisms described in Section 4.2.

4.2 Course Allocation Mechanisms

Using the simple student utility model described above, we calculate the outcomes of four course allocation mechanisms across 100 runs. Within each run, we consider a fixed draw of students’ utilities. We assume that students report their utilities truthfully, as all considered mechanisms are either strategy-proof or strategy-proof in the large (see Theorem 5).

The Pseudo-Market with Priorities (PMP) mechanism. To find the outcome of the PMP mechanism, we assign budgets evenly spaced between 1 and $1 + \beta$ with $\beta = 1/(k - 1) = 0.25$. The theoretical market-clearing error bound equals $\alpha = \sqrt{kM/2} \approx 43.5$. We also assign students to $R = 8$ priority groups for each course. Priority is on four years of study and whether a student is within a department with course reserves. We consider two cases: one is with year-specific priority taking precedence over department-specific priority and the other one is with precedence being reversed.

The algorithm that searches for a Pseudo-Market Equilibrium with priorities departs from the one used in Budish (2011) to accommodate the presence of priority-specific prices and to obtain a tighter bound on the number of oversubscribed courses. The algorithm has two phases. Phase I searches for priority-specific prices that ensures the market-clearing error is below theoretical bound α . Using the idea of the proof of Theorem 1, we conveniently parameterize prices with $t \in [0, R\bar{b}]^M$, where $\bar{b} = 1.251$. For each $t \in [0, R\bar{b}]^M$, course $c \in \mathcal{C}$, and priority $r \in \mathcal{R}$, we define priority-specific prices as $p_{c,r}(t) = \max(t_c - (r - 1)\bar{b}, 0)$. Then, we look for a pseudo-market market equilibrium in this lower-dimensional space. This idea significantly speeds up the search for a Pseudo-Market Equilibrium.

The search in Phase I starts with an educated guess for equilibrium prices based on the student demand at zero prices.²² Then, the algorithm finds each student’s utility-maximizing schedule and adjusts prices proportional to the number of over and undersubscribed seats. The algorithm searches in the lower dimensional space of t -parameters for student demand with the market-clearing error smaller than α . If the algorithm fails to improve the error in a given iteration, the current price vector is retained and the size of the price adjustment is reduced in the next iteration. The prices and the allocation of the smallest market clearing error are also retained. Once theoretical bound α is reached, the algorithm continues to improve course allocation until the market-clearing error fails to improve by more than 1% within six consecutive iterations.²³

Phase II takes the allocation obtained in Phase I and gradually increases the prices only for oversubscribed courses to ensure that no courses are oversubscribed by more than $k-1 = 4$ seats and the percentage of oversubscribed courses is smaller than the one observed in the data (see Table 4). As a result of Phase II, an allocation with a market-clearing error greater than the theoretical bound α might be obtained. In this case, the algorithm returns to Phase I. Simulations show that there are typically no more than 1-3 iterations between Phases I and II. However, if there are more than 6 iterations between the phases, the algorithm restarts with a new set of initial prices.

Deferred Acceptance mechanism with single and multiple tie-breakings. Both of these mechanisms use the same priority groups as the PMP mechanism. The main difference between the two variants is how they break ties among students with the same priority. The first variant uses a single tie-breaking lottery across all courses, which determines the strict priority order for each course, then executes the student-optimal deferred acceptance algorithm. We refer to this version of the mechanism as DA. The second variant uses multiple course-specific tie-breakings. We refer to this version of the mechanism as DA(m). Both versions have received a thorough analysis in the many-to-one matching literature (see, e.g., Erdil and Ergin, 2008, 2017).

²²We found that starting with an educated guess improves the speed of convergence over starting with random prices as in Budish (2011).

²³The allocation search process is adjustable. The current settings balance the trade-off between a small market clearing error and the algorithm’s run time. We use the Walras Tâtonnement procedure to find an approximate equilibrium, a simpler algorithm than the one used in Budish (2011). At the same time, we do not consider course complementarities, which are a part of Budish’s paper.

The Random Serial Dictatorship (RSD) with set-asides. We finally describe how we determine the outcome of the RSD mechanism with set-asides. We consider the strict priority order of students as in the DA mechanism with the single tie-breaking lottery. Using this order, the program assigns available course seats according to student utility. The course reservations are treated as described in Section 4.1: for each year of study (including the “All” option), there is a set of departments and amount of course reserved seats for each course (which could be zero). If a student satisfies one of the course reservation criteria, the student can be enrolled in one of the reserved or regular seats; otherwise, the student can be enrolled only in a regular seat. If a student is eligible for two course reservation criteria, the set-asides with year-specific reservations are assigned first and the set-asides with all years of study reservations are assigned last.

During numerous rounds of simulations, we encountered the question of how to pick the number of reserved seats. We first set the number of reserved seats equal to the number of *adjusted course reserves* as described in Section 4.1. This led to significant under-performance of RSD with set-asides.²⁴ The mechanism treated the number of reserved seats as fixed, which resulted in numerous inefficiencies, especially when too many seats were reserved. As described in Section 4.1, the course enrollment process in practice is more flexible: the university departments relax course reservation constraints during the course allocation process. Doing so releases pressure on the department heads to set course reservations precisely at the beginning of the process. The PMP, DA, and DA(m) compute “effective set-asides” during the course allocation process. Hence, comparing the performance of the other three mechanisms to the performance of RSD with a fixed number of set-asides would not be a fair exercise.

Instead, we decided to take a different route and estimate *the optimal set-asides* in the RSD mechanism. For this purpose, we utilize the DA mechanism. We generate 100 environments with random student utilities according to the student utility model and a random tie-breaking and run the DA mechanism. To determine the optimal set-asides for a given set of departments, we take the average across all environments (rounded to the nearest integer).

In addition to the inefficiencies associated with accurately determining set-asides, the RSD with set-asides mixes two priority structures: year-specific priority and department-

²⁴The RSD with set-asides from the actual data leads to even worse performance.

specific priority. Hence, the RSD mechanism does not respect any priority structure.²⁵ This leads to the question of what priority structure we should choose when running the PMP, DA, and DA(m) mechanisms and whether we should prioritize the year- or department-specific priority. Our main simulations compare the performance of the four mechanisms with a course priority structure in which year-specific priority takes precedence over department-specific priority. Appendix D contains simulations in which department-specific priorities take precedence over year-specific priority.

4.3 Simulation results

We consider the Random Serial Dictatorship (RSD) mechanism with the optimal set-asides as our benchmark. We then evaluate the performance of the Pseudo-Market with Priorities (PMP) and Deferred Acceptance algorithm with single and multiple tie-breakings (DA and DA(m), respectively) relative to this benchmark using *student utilities*, *the standard deviation of student utilities* across students, and *the number of students who experience schedule envy*. The first metric measures students' satisfaction in their respective schedules. The two other metrics measure the fairness of course allocation among students. We conduct 100 simulation runs, where each run corresponds to a random utility draw and a random tie-breaking, which we keep the same across RSD, PMP, and DA.²⁶ DA(m) has multiple course-specific tie-breakings. We highlight that the standard deviation of student utilities measures how much student utilities are dispersed within each run of the simulation, not across the runs.

Table 5 presents the results of our simulations averaged across all runs. We observe that the PMP mechanism delivers the highest mean student utility for all years of study among all four mechanisms. The improvements compared to the benchmark for first and second-year students are 6.92% and 3.23%, respectively. The percentage numbers correspond to 1223 out of 1565 frosh and 734 out of 1611 sophomores to receive a schedule in the PMP mechanism that differs from the RSD outcome with the average improvements in utility of 9% and 7%,

²⁵If we assume one takes precedence over the other, it is always possible to come up with a counter-example leading to a priority violation for the final RSD outcome (a stability violation). For instance, assume that year-specific priority takes precedence over department-specific priority. Then, some senior students may remain unassigned even though set-aside seats remain unfilled.

²⁶The program code using Mathematica is available. The majority of the simulations were run on a server with 12 Cores of CPU (2.6 GHz of an Intel Xeon Gold 6126 CPU), 64 GBs of RAM, 70 GBs of SSD hard drive space, and Windows Server 2019 Standard operating system. On average, it takes about one hour to calculate the allocation outcomes of all four mechanisms at one core.

Mean Utility	Year 1	Year 2	Year 3	Years 4
Pseudo-Market with Priorities	6.92%	3.23%	1.43%	0.43%
Deferred Acceptance with single tie-breaking	-0.34%	0.67%	0.68%	0.39%
Deferred Acceptance with multiple tie-breakings	-3.29%	0.11%	0.62%	0.41%
St. Dev. of Utility	Year 1	Year 2	Year 3	Years 4
Pseudo-Market with Priorities	-11.16%	-2.92%	-0.64%	0.05%
Deferred Acceptance with single tie-breaking	-1.36%	-0.40%	-0.02%	0.09%
Deferred Acceptance with multiple tie-breakings	-9.79%	-2.84%	-0.38%	0.07%

TABLE 5: The performance of the Pseudo-Market with Priorities, Deferred Acceptance with single tie-breaking, and Deferred Acceptance with multiple tie-breakings mechanisms compared to the benchmark of Random Serial Dictatorship with *optimal set-asides*, with year-specific priorities taking precedence over department-specific priorities. The results are based on 100 runs with different random utility draws.

respectively. The improvements for third and fourth-year students, where the allocation of course seats is barely constrained, are 1.43% and 0.43%, respectively. These small relative improvements are actually significant for many students: 307 out of 1422 juniors and 62 out of 1425 seniors receive a different schedule in the PMP mechanism with the average improvement in utility of 7% and 10%, respectively. The improvements in student mean utilities compared to the RSD benchmark are statistically significant for all years of study (see Appendix D.1 for details). Both the DA and DA(m) mechanisms deliver only a small improvement (less than 1%) relative to the benchmark for years 2-4. For first-years, both the DA and DA(m) underperform relative to the benchmark. This occurs because the RSD mechanism reserves some course seats to Year 1 students leading to better outcomes for them. The performance of the DA(m) mechanism is slightly inferior to that of DA for years 1-3, which is in line with the comparison of the DA and DA(m) mechanisms in the previous literature (e.g., [Abdulkadiroğlu, Pathak, and Roth, 2009](#); [Abdulkadiroğlu, Che, and Yasuda, 2015](#)). However, this difference is not statistically significant for most years of study. Overall, the result that the PMP delivers the highest student utility arises because the PMP mechanism is a cardinal mechanism, accounting for the cardinal utilities of students. In contrast, the other three are all ordinal mechanisms, based only on the ordinal preferences

# seats above max capacity	≥ 1	≥ 2	≥ 3	≥ 4	≥ 5
% of courses	4.23%	0.86%	0.24%	0.16%	0%

TABLE 6: The percentage of oversubscribed courses in the PMP mechanism.

of students.

In principle, the higher mean utility in the outcomes of the PMP mechanism might be due to the presence of the market-clearing error, which allows some courses to be oversubscribed.²⁷ On average, the market-clearing error of the PMP outcomes is about 14 seats, which is much smaller than theoretical error $\alpha \approx 43.5$. Table 6 also reports the distribution of the number of oversubscribed courses. No courses are oversubscribed by more than $k - 1 = 4$ seats, in line with [Lin, Nguyen, Nguyen, and Altinkemer \(2022\)](#), which provides a theoretical justification for the existence of an equilibrium with a good-by-good clearing bound equal to $k - 1$. The number of oversubscribed courses is also smaller than the ones observed in the data presented in Table 4.

In addition, Table 5 presents evidence that the PMP mechanism delivers a reduced standard deviation in student utilities for years 1-3 compared to the three other mechanisms. The standard deviation for year 4 is slightly higher than the RSD benchmark (0.05%). This increase is primarily associated with an increase in the utilities of seniors compared to the RSD benchmark. The change is also small for second and third-year students, where the impact of a good allocation mechanism is limited (2.92% and 0.64%, respectively). The largest reduction occurs for first-year students (11.16%). The drop in standard deviation for the DA mechanism is less than 1.5% for all years of study. However, we observe a more pronounced decrease in the standard deviation for years 1-3 for the DA(m) mechanism (between 0.38% and 9.79%). F-tests for the PMP and DA(m) mechanisms in each simulation run reject the null hypothesis that the standard deviation of first-year student utilities is larger than that of the RSD mechanism on a 5% significance level (see Appendix D.1 for details)

To provide additional information on allocation fairness, we also report in Table 7 data on students who experience envy towards students of the same or lower priority level. The first column shows the percentage of students who experience no such envy. To obtain the

²⁷Alternatively, we could have compared the performance of the four mechanisms with quotas equal to the number of seats assigned in the PMP mechanism for each run. We thank Olivier Tercieux for this suggestion.

	0 courses	1 course	2 courses	3 courses	4 courses	5 courses
Pseudo-Market with Priorities	98.74%	1.26%	0%	0%	0%	0%
Random Serial Dictatorship with optimal set-asides	92.63%	6.30%	0.91%	0.14%	0.02%	0.001%
Deferred Acceptance with single tie-breaking	93.38%	5.68%	0.82%	0.11%	0.01%	0.0005%
Deferred Acceptance with multiple tie-breakings	93.96%	5.74%	0.29%	0.01%	0.0003%	0%

TABLE 7: The percentage of students who experience envy towards students of the same or lower priority for the four mechanisms. The first column provides information about students who experience no envy. The other columns show the percentage of students who experience envy bounded by 1, ..., 5 courses.

numbers in the other columns, we remove courses one by one from envied schedules and check whether envy vanishes. Our simulations support Theorem 4 that shows that the PMP mechanism satisfies schedule envy bounded by a single course among students of the same or lower levels of priority. None of the other three mechanisms satisfy this property. Still, our results show that envy is generally by no more than by a single course in the RSD, DA, and DA(m) mechanisms. However, the percentages of students who experience any envy are relatively large (7.37%, 6.62%, and 6.04%, respectively). In contrast, just 1.26% of students in the PMP mechanism prefer the course schedule assigned to another student of the same or lower levels of priority in each course. Overall, our simulations support that the PMP mechanism limits envy and provides a fairer course allocation among students.

Last, we discuss the course prices in the equilibria of the PMP mechanism. Table 8 presents information about price cutoff levels and the percentages of students allocated different numbers of courses with positive prices (when student budgets are between 1 and 1.25). The table shows that 62.8% of courses have the lowest possible price cutoff, reflecting the non-binding capacity of most courses in the data. The average prices are more or less stable at most cutoffs except the lowest one.²⁸ The overwhelming majority of senior students obtain all their courses for zero price, with only around 12% paying a positive price for at least one course. A similar picture is observed for second and third-year students, with many also obtaining all their seats for free. The picture is completely different for first-year

²⁸A positive price at a cutoff level means that only students of the same or higher priority can obtain a seat in this course in the equilibrium.

Cutoff level of priority and their average prices								
Cutoff level of priority	1	2	3	4	5	6	7	8
% of courses	62.8%	4.8%	13.8%	3.2%	10.7%	1.7%	2.9%	0.1%
average cutoff price	0.18	0.62	0.81	0.85	0.83	0.84	0.77	0.86

Courses with positive prices (% is the percentage of students)				
# courses with positive prices	Year 1	Year 2	Year 3	Years 4
0 courses	2.4%	37.3%	62.6%	88.0 %
1 course	33.3%	49.4%	34.3%	11.6%
2 courses	41.1%	12.2%	2.9%	0.4%
3 courses	18.8 %	1.1%	0.2%	0.01%
4 courses	4.1%	0.04%	0.01%	0%
5 courses	0.3%	0.01%	0%	0%

TABLE 8: The prices and course allocation in the PMP mechanism with students budget in $[1, 1 + \beta] = [1, 1.25]$. The top section presents the information about cutoff levels and their average prices. The bottom section shows the percentage of students allocated different numbers of courses with positive prices by year of study.

students. Only 2.4% of first-year students consume all courses for free. A typical first-year student pays a positive price for 1-3 courses. To afford this, course prices at the lowest cutoff level should be low, with an average price of 0.18 at the lowest cutoff level.

Overall, the Pseudo-Market with Priorities mechanism delivers a higher mean student utility, a lower standard deviation of student utilities, and a smaller number of students who experience envy compared to the Random Serial Dictatorship with set-asides and the Deferred Acceptance algorithms with single and multiple tie-breakings. Hence, the Pseudo-Market with Priorities mechanism leads to higher average satisfaction among students and produces fairer outcomes than the mechanisms commonly used in practice.

5 Conclusion

In this paper, we explore a many-to-many matching problem that arises in undergraduate course allocation. We allow courses to prioritize students based on factors such as their year of study and department and design a deterministic allocation mechanism, the *Pseudo-Market with Priorities* (PMP) mechanism, that respects this priority structure. The PMP

mechanism is an extension of the approximate competitive equilibrium from equal incomes mechanism (Budish, 2011) to settings with course priorities. This mechanism maintains a small market-clearing error and has desirable properties in terms of stability, efficiency, fairness, and strategy-proofness.

To employ the PMP mechanism, university registrars need to adopt a preference reporting language that allows students to express how they compare different course schedules and account for substitutabilities and complementarities that may exist among courses. Software developed by Budish, Cachon, Kessler, and Othman (2017) to implement the approximate competitive equilibrium from equal incomes at Wharton Business School and Columbia Business School exhibits this in a manner that is easily accessible to students (see also Bichler and Merting, 2021). Moreover, Soumalias, Zamanlooy, Weissteiner, and Seuken (2023) recently developed powerful machine learning-based techniques to help with student preference elicitation.²⁹

Finding a market equilibrium with a small market-clearing error allowing for general student preferences could be also a difficult computational problem (see Othman, Papadimitriou, and Rubinstein, 2016; Vazirani and Yannakakis, 2021). However, when complementarities in preferences are limited, heuristic algorithms find a market equilibrium with much tighter approximations than theoretical bound (see Budish, Cachon, Kessler, and Othman, 2017; Othman, Sandholm, and Budish, 2010). In this paper, we also show that when preferences are additive an algorithm based on Walrasian tâtonnement locates a market equilibrium fast and leads to tight approximation bounds. Hence, for practical applications, it is important to understand what is the best input language of preferences that is expressive enough and also guarantees the calculation of a market equilibrium with a market-clearing error below theoretical bounds for almost every instance (see Boutilier and Hoos, 2001; Sandholm and Boutilier, 2006).

In closing, we highlight the ways in which university registrars can improve undergraduate course allocation by adopting the PMP mechanism. In comparison to the RSD mechanism, PMP has several noteworthy advantages. A first advantage is that the PMP mechanism endogenously computes course set-asides. The PMP mechanism prevents students of lower priority from taking course seats that higher priority students could benefit from, thereby

²⁹See also Brero, Lubin, and Seuken (2018) for the use of these techniques for the preference elicitation in combinatorial clock auctions.

relieving university registrars of attempting to correctly estimate the number of set-asides. Thus, the mechanism ensures that all the students who require a course in their department are assigned these courses and limits the number of remaining places that must be assigned manually. The second advantage of the PMP mechanism over the RSD mechanism (as well as the DA and DA(m) mechanisms) is that it uses student cardinal preferences to allocate course seats. Hence, it is more flexible than the other three mechanisms, given its potential to accommodate preference complementarities and deliver higher utility to students. The third advantage of the PMP mechanism is that it allocates courses more fairly among students—an outcome that none of the other three mechanisms can deliver. Our simulations based on actual university data support these advantages of the PMP mechanism.

Finally, adopting the Pseudo-Market with Priorities mechanism would provide valuable data to university registrars on student demand. As students have almost no incentive to misrepresent their preferences, the market-clearing prices of these mechanisms can serve as indicators of student demand and help universities distinguish popular courses from unpopular ones. With this information, universities can adjust class sizes, timing, and sections to further increase student satisfaction.

A Appendix: Proofs

Proof of Theorem 1. Consider an economy $(\mathcal{S}, \mathcal{C}, Q, \mathcal{V}, \mathfrak{R})$ and a budget vector $b = (b_1, \dots, b_M)$ that satisfies $1 \leq \min_s(b_s) \leq \max_s(b_s) < \bar{b} \equiv 1 + \beta$ for some $\beta > 0$.

We consider the M -dimensional set $\mathcal{T} = [0, R\bar{b}]^M$, which allows us to conveniently parameterize priority-specific prices and look for a competitive market equilibrium in a lower dimensional space. In particular, for each $t \in \mathcal{T}$, course $c \in \mathcal{C}$, and the level of priority $r \in \mathcal{R}$, we define priority-specific prices as

$$p_{c,r}(t) = \max(t_c - (r - 1)\bar{b}, 0). \quad (\text{A.1})$$

For each $t \in \mathcal{T}$ and each $c \in \mathcal{C}$ there is a unique cutoff level of priority $r_c^*(t) \in \mathcal{R}$ such that for any $r \in \{1, \dots, R\}$, $p_{c,r}(t)$ satisfies

$$p_{c,r}(t) \in \begin{cases} \{0\} & r > r_c^*(t) \\ [0, \bar{b}] & r = r_c^*(t) \\ [\bar{b}, +\infty) & r < r_c^*(t) \end{cases}. \quad (\text{A.2})$$

We will also consider an auxiliary enlargement of this set, $\tilde{\mathcal{T}} = [-1, R\bar{b} + 1]^M$, and similarly define $p_{c,r}(\tilde{t})$ for $\tilde{t} \in \tilde{\mathcal{T}}$. We define demand function $d_s : \tilde{\mathcal{T}} \rightarrow 2^{\mathcal{C}}$ as

$$d_s(\tilde{t}) = \max_{\tilde{x}_s} \left\{ x'_s \subseteq \mathcal{C} : \sum_{c \in \mathcal{C}} x'_{s,c} \max(\tilde{t}_c - (r_{s,c} - 1)\bar{b}, 0) \leq b_s + \tau_{s,x'_s} \right\},$$

where the τ_{s,x_s} are student- and schedule-specific taxes chosen to ensure that the demand is single-valued (similarly to [Budish, 2011](#)). For each course $c \in \mathcal{C}$, excess demand $z_c : \tilde{\mathcal{T}} \rightarrow \mathbb{Z}$ is defined by

$$z_c(\tilde{t}) = \sum_{s \in \mathcal{S}} x_{s,c}^* - q_c,$$

where $x_s^* = d_s(\tilde{t}^*)$ for all $s \in \mathcal{S}$. The excess demand is bounded because $-S \leq z_c \leq S - 1$ for all $c \in \mathcal{C}$. We also define a budget surface for each student $s \in \mathcal{S}$ and schedule $x_s \subseteq \mathcal{C}$ as

$$H(s, x_s) = \left\{ \tilde{t} \in \tilde{\mathcal{T}} : \sum_{c \in \mathcal{C}} x_{s,c} \max(\tilde{t}_c - (r_{s,c} - 1)\bar{b}, 0) = b_s + \tau_{s,x_s} \right\}.$$

Note that the budget surface $H(s, x_s)$ may not be a hyperplane as in the case without priorities (see [Budish, 2011](#)). Lemma [B1](#) in Appendix [B](#) shows that it is still possible to choose b_s and τ_{s, x_s} such that at most M budget constraints intersect for any $\tilde{t} \in \tilde{\mathcal{T}}$.

Next, we define a truncation function $trunc: \tilde{\mathcal{T}} \rightarrow \mathcal{T}$, where for each $c \in \mathcal{C}$

$$(trunc(\tilde{t}))_c = \min\{R\bar{b}, \max\{0, \tilde{t}_c\}\}.$$

Also, we introduce a tâtonnement price adjustment function $f: \tilde{\mathcal{T}} \rightarrow \tilde{\mathcal{T}}$ by

$$f(\tilde{t}) = trunc(\tilde{t}) + \gamma z(trunc(\tilde{t})),$$

where $\gamma \in (0, 1/S)$. Suppose that f has a fixed point $\tilde{t}^* = f(\tilde{t}^*)$, and denote its truncation by $t^* = trunc(\tilde{t}^*)$. We show that prices $\{p_{c,r}(t^*)\}_{c \in \mathcal{C}, r \in \mathcal{R}}$ defined by equation [\(A.1\)](#), allocation $x_s^* = d_s(t^*)$, and budgets $b_s^* = b_s + \tau_{s, x_s^*}$ for all $s \in \mathcal{S}$ constitute an exact competitive equilibrium (or $(0, \beta)$ -Pseudo-Market Equilibrium with Priorities as in [Definition 5](#)). We might only need to slightly adjust t^* to find another fixed point that satisfies condition

$$\sum_{\{s \in \mathcal{S}: r_{s,c} > r_c^*(t^*)\}} x_{s,c}^* < q_c.$$

- Prices $\{p_{c,r}(t^*)\}_{c \in \mathcal{C}, r \in \mathcal{R}}$ and cutoffs defined by [\(A.1\)](#) and [\(A.2\)](#) ensure that condition [\(1\)](#) is satisfied.
- The definition of demand function implies that any course schedule that student s prefers to $x_s^* = d_s(t^*)$ must cost strictly more than $b_s^* = b_s + \tau_{s, x_s^*}$.
- $p_{c,1}(t^*) > 0$ implies $z_c(t^*) = 0$. To see this, note that equation [\(A.1\)](#) implies $\tilde{t}_c^* > 0$. In addition, we must have $\tilde{t}_c^* < R\bar{b}$; otherwise $(trunc(\tilde{t}^*))_c = R\bar{b}$ and $z_c(\tilde{t}^*) < 0$, which contradicts the fixed point equation. Hence, $\tilde{t}_c^* \in (0, R\bar{b})$ and the fixed point equation ensures $z_c(t^*) = 0$.
- $p_{c,1}(t^*) = 0$ implies $z_c(t^*) \leq 0$. To see this, consider two cases. If $\tilde{t}_c^* \in (0, R\bar{b})$, the fixed point equation implies $z_c(t^*) = 0$. If $\tilde{t}_c^* \in [-1, 0]$, we have $t_c^* \equiv trunc(\tilde{t}_c^*) = 0$, $p_{c,1}(t^*) = 0$, and the fixed point ensures $z_c(t^*) \leq 0$.
- To make sure condition $\sum_{\{s \in \mathcal{S}: r_{s,c} > r_c^*(t^*)\}} x_{s,c}^* < q_c$ is satisfied (see condition [\(1\)](#)), assume from the contrary that the demand for a course c across priority levels lower than $r_c^*(t^*)$ is greater than its number of seats; that is, $\sum_{\{s \in \mathcal{S}: r_{s,c} > r_c^*(t^*)\}} x_{s,c}^* \geq q_c$. The

previous two bullet points establish that $z_c(t^*) \leq 0$. Hence, our assumption implies $\sum_{\{s \in \mathcal{S}: r_{s,c} > r_c^*(t^*)\}} x_{s,c}^* - q_c = \sum_{s \in \mathcal{S}} x_{s,c}^* - q_c = 0$. In other words, there is no demand for course c from students at level of priority $r_c^*(t^*)$. So, we can consider $\hat{t}_c^* = t_c^* + \bar{b} - p_{c,r_c^*(t^*)}$, where the cutoff priority group faces price \bar{b} and the prices of lower priority levels do not change. With $\hat{t}^* = (\hat{t}_c^*, t_{-c}^*)$, we obtain that $p_{c,r_c^*(\hat{t}^*)} = 0$, $p_{c,r_c^*(\hat{t}^*)} = \bar{b}$, and the adjustment does not change the demand and excess demand for course c for all students, but it increases the cutoff level $r_c^*(\hat{t}^*) = r_c^*(t^*) + 1$. If we still have $\sum_{\{s \in \mathcal{S}: r_{s,c} > r_c^*(\hat{t}^*)\}} x_{s,c}^* \geq q_c$, we repeat the price adjustment until the condition is satisfied. Finally, $z_c(\hat{t}^*) = z_c(t^*) = 0$ implies $\hat{t}_c^* < R\bar{b}$ and $\hat{t}_c^* > t_c^* \geq 0$. Therefore, $\hat{t}_c^* = \text{trunc}(\hat{t}_c^*) + z_c(\hat{t}^*)$, and, hence, vector \hat{t}^* satisfies the fixed point equation and is such that $\sum_{\{s \in \mathcal{S}: r_{s,c} > r_c^*(\hat{t}^*)\}} x_{s,c}^* < q_c$.

Overall, if f has a fixed point $\tilde{t}^* = f(\tilde{t}^*)$, its truncation $t^* = \text{trunc}(\tilde{t}^*)$ (or its adjustment \hat{t}^*) is an exact competitive equilibrium price vector for allocation $x_s^* = d_s(t^*)$ and budgets b_s^* for all $s \in \mathcal{S}$.

Though, the fixed point of operator f might fail to exist. Following [Budish \(2011\)](#), we define a ‘‘convexification’’ of f , $F : \tilde{\mathcal{T}} \rightarrow \tilde{\mathcal{T}}$, by

$$F(\tilde{t}) = \text{co}\{y : \exists \text{ a sequence } \tilde{t}^w \rightarrow \tilde{t}, \tilde{t} \neq \tilde{t}^w \in \tilde{\mathcal{T}} \text{ such that } f(\tilde{t}^w) \rightarrow y\},$$

where co denotes the convex hull of the set. F is nonempty, $\tilde{\mathcal{T}}$ is compact and convex, and $F(t)$ is convex. From Lemma 2.4 of [Cromme and Diener \(1991\)](#), F is an upper hemicontinuous correspondence and hence has a fixed point by Kakutani’s fixed point theorem. We denote a fixed point by $\tilde{t}^* \in F(\tilde{t}^*)$, and let again $t^* = \text{trunc}(\tilde{t}^*)$ be its truncation.

Note that the reduction of the MR -dimensional space of prices $\{p_{c,r}\}_{c \in \mathcal{C}, r \in \mathcal{R}}$ to M -dimensional space $\tilde{\mathcal{T}} = [-1, R\bar{b} + 1]^M$ allows us to use the steps of [Budish \(2011\)](#) to establish the existence of a $(\sqrt{kM/2}, \beta)$ -Pseudo-Market Equilibrium with Priorities for any $\beta > 0$. For completeness, we provide adapted steps in Appendix B. \square

Proof of Theorem 3. Let (x^*, p^*, b^*) be a (α, β) -Pseudo-Market Equilibrium with Priorities. Suppose that allocation y Pareto dominates allocation x^* and has the same number of course seats assigned for each course. Then, there is a student s' who strictly prefer

allocation y to x^* . In addition,

$$\sum_{c \in \mathcal{C}} p_{c,r_{s',c}}^* y_{s',c} > b_{s'}^* \geq \sum_{c \in \mathcal{C}} p_{c,r_{s',c}}^* x_{s',c}^*.$$

For the other students we must also have

$$\sum_{c \in \mathcal{C}} p_{c,r_{s,c}}^* y_{s,c} \geq b_s^* \geq \sum_{c \in \mathcal{C}} p_{c,r_{s,c}}^* x_{s,c}^*$$

as we assume that students have strict preferences among course schedules. Summing the above inequalities over all agents, we obtain

$$\sum_{s \in \mathcal{S}} \sum_{c \in \mathcal{C}} p_{c,r_{s,c}}^* y_{s,c} > \sum_{s \in \mathcal{S}} \sum_{c \in \mathcal{C}} p_{c,r_{s,c}}^* x_{s,c}^*.$$

We can further rearrange

$$\sum_{c \in \mathcal{C}} \sum_{r=1}^R p_{c,r}^* \sum_{\{s:r_{s,c}=r\}} y_{s,c} > \sum_{c \in \mathcal{C}} \sum_{r=1}^R p_{c,r}^* \sum_{\{s:r_{s,c}=r\}} x_{s,c}^*.$$

Hence, we obtain

$$0 < \sum_{c \in \mathcal{C}} \sum_{r=1}^R p_{c,r}^* \sum_{\{s:r_{s,c}=r\}} (y_{s,c} - x_{s,c}^*) = \sum_{c \in \mathcal{C}} \sum_{r=2}^R (p_{c,r}^* - p_{c,r-1}^*) \sum_{\{s:r_{s,c} \geq r\}} (y_{s,c} - x_{s,c}^*) + p_{c,1}^* \sum_{\{s:r_{s,c} \geq 1\}} (y_{s,c} - x_{s,c}^*).$$

As the number of course seats allocated to students for allocations y and x^* is the same, the last term equals zero. Also, $p_{c,r}^* - p_{c,r-1}^* \leq 0$, and we must have $\sum_{\{s:r_{s,c} \geq r\}} (y_{s,c} - x_{s,c}^*) < 0$ for at least some course c and rank r . This implies $y_c \not\preceq_c x_c^*$ and that allocation x^* is approximately priority-constrained efficient. \square

Proof of Theorem 4. We assume that $\beta \leq \frac{1}{k-1}$ and show that the Pseudo-Market with Priorities mechanism has schedule envy bounded by a single course among students of the same or lower levels of priority. Let us consider two students $s, s' \in \mathcal{S}$ such that $r_{s,c} \geq r_{s',c}$ for all $c \in \mathcal{C}$. We denote the prices faced by s and s' in market equilibrium as $p_s^* = \{p_{c,r_{s,c}}^*\}_{c \in \mathcal{C}}$ and $p_{s'}^* = \{p_{c,r_{s',c}}^*\}_{c \in \mathcal{C}}$, and the course schedule assigned in market equilibrium to student s' as $x_{s'}^* = (c_{j_1}, \dots, c_{j_{k'}})$ for $j_1, \dots, j_{k'} \in \{1, \dots, M\}$ and $k' \leq k$. Suppose that s envies student s'

and that this envy is not bounded by one course. Therefore, student s is not able to afford the course schedule of student s' even if one course from schedule $x_{s'}^*$ is dropped. That is,

$$p_s^* \cdot x_{s'}^* \setminus \{c_{j_\ell}\} > b_s^*,$$

for $\ell = 1, \dots, k'$. Since s' has the same or lower priority, we have $p_{s'}^* \geq p_s^*$ and

$$p_{s'}^* \cdot x_{s'}^* \setminus \{c_{j_\ell}\} > b_s^*,$$

for $\ell = 1, \dots, k'$. Summing these inequalities over ℓ and using $b_{s'}^* \geq p_{s'}^* \cdot x_{s'}^*$, we obtain

$$(k' - 1)b_{s'}^* \geq (k' - 1)p_{s'}^* \cdot x_{s'}^* > k'b_s^*.$$

The latter implies $\frac{b_{s'}^*}{b_s^*} > \frac{k'}{k'-1} \geq \frac{k}{k-1} \geq 1 + \beta$, which contradicts how budgets are allocated. \square

Proof of Theorem 5. To prove the statement of the theorem, we first formally define a direct mechanism, a semi-anonymous direct mechanism, and the property of being *strategy-proof in the large*. Then, we show that the PMP mechanism is a semi-anonymous direct mechanism that is *envy-free but for tie-breaking*. The result then follows from Appendix C in [Azevedo and Budish \(2019\)](#).

Consider a sequence of markets labeled by the number of students $|\mathcal{S}^N| \equiv N$.³⁰ We assume that the set of courses \mathcal{C} is fixed with the number of courses equal to M , but the capacity of each course q_c^N can vary. An allocation $x \in X_0^N \equiv \{0, 1\}^{MN}$ specifies schedule x_s for each student $s \in \mathcal{S}^N$. Also, $X = \Delta X_0$ denotes the set of random course allocations. Each student s has a type $v_s = (\succeq_s, r_s)$ that describes her preferences over course \succeq_s and her priorities r_s . The set of all possible types is denoted as \mathcal{V} . We assume that each student can misrepresent her preferences, but not her priorities. Each student has also a von Neumann-Morgenstern utility function $u_v : X \rightarrow [0, 1]$ that is consistent with type $v \in \mathcal{V}$.

We consider a sequence of *direct mechanisms* $\{\Phi^N\}_{N \in \mathbb{N}}$ that for each report of student preferences assigns a distribution over course allocations; that is, $\Phi^N : \mathcal{V}^N \rightarrow \Delta(X_0^N)$. The PMP mechanism is a typical example of a *semi-anonymous direct mechanism* where agents are divided into groups, agents within each group are treated the same way, but agents across the groups can be treated differently. Formally, we partition \mathcal{S}^N into groups according to their course priorities $S_{r_1, r_2, \dots, r_M}^N$, where $r_j \in \{1, \dots, R\}$, $j = 1, \dots, M$. This partition collects

³⁰We slightly depart from the notation of the main text where we labeled the number of students as S .

students who have the same level of priority for all courses in one group. The students in the same group face the same prices in the PMP mechanism and the same budget distribution.

To define the property of a direct mechanism being strategy-proof in the large, we consider function $\phi^N : \mathcal{V} \times \Delta\mathcal{V} \rightarrow X$ for each N according to

$$\phi^N(v_s, m) = \sum_{v_{-s} \in \mathcal{V}^{N-1}} \Phi_s^N(v_s, v_{-s}) Pr(v_{-s} | v_{-s} \sim iid(m)),$$

where $\Phi_s^N(v_s, v_{-s})$ denotes the course schedule obtained by student s when she reports type v_s and all other students report v_{-s} . $Pr(v_{-s} | v_{-s} \sim iid(m))$ denotes the probability that profile v_{-s} is realized when other students' types v_{-s} are independent and identically distributed according to $m \in \Delta\mathcal{V}$. In other words, $\phi^N(v_s, m)$ describes the random outcome that student s expects to receive when she reports v_s and the other students' types are distributed independently and identically according to m —and they report their types truthfully. Using this notation, we have the following definition.

Definition A1. *The direct semi-anonymous mechanism $\{\Phi^N\}_{N \in \mathbb{N}}$ is strategy-proof in the large if, for any random distribution of reports by other students $m \in \Delta\mathcal{V}$ that has full support and $\varepsilon > 0$, there exists n_0 such that, for all $n \geq n_0$ and all $v'_s, v_s \in \mathcal{V}$ where students can misrepresent only their preferences, but not priorities, we have $u_{v_s}[\phi^N(v_s, m)] \geq u_{v_s}[\phi^N(v'_s, m)]$.*

In other words, in a large enough market, reporting preferences truthfully is approximately optimal, for any independent and identical distribution of the other students' types that has full support. Note that students can misrepresent only their preferences, but not priorities.

The PMP mechanism also assigns budgets b_s that can be represented through lottery number ℓ_s , which is uniformly distributed from $[0, 1]$ as $b_s = 1 + \ell_s \beta$. Because budgets are random, the outcome of the PMP mechanism for each $v \in \mathcal{V}^N$ can be represented as $\Phi^N(v) = \int_{\ell \in [0, 1]^N} x^N(v, \ell) d\ell$, where $x^N(v, \ell) \in \{0, 1\}^{MN}$ assigns a course schedule to each student in \mathcal{S}^N . Definition 5 guarantees that any student who has a larger budget b_s (a large lottery number ℓ_s) prefers her course schedule to any course schedule assigned to a student in the same group. As a result, the PMP mechanism is a semi-anonymous mechanism that is *envy-free but for tie-breaking*, which is defined as follows (see [Azevedo and Budish, 2019](#)).

Definition A2. A direct semi-anonymous mechanism $\{\Phi^N\}_{N \in \mathcal{N}}$ is envy-free but for tie-breaking if for each N there exists $x^N : \mathcal{V}^N \times [0, 1] \rightarrow \Delta(X_0^N)$, symmetric over its coordinates, such that $\Phi^N(v) = \int_{l \in [0, 1]^N} x^N(v, l) dl$, and, for all s, s', N, v , and ℓ , if $\ell_s \geq \ell_{s'}$, and v_s and $v_{s'}$ belong to the same group, then $u_{v_s}(x_s^N(v, \ell)) \geq u_{v_{s'}}(x_{s'}^N(v, \ell))$.

Appendix C in [Azevedo and Budish \(2019\)](#) provides an argument explaining why any semi-anonymous mechanism that is envy-free except for tie-breaking is strategy-proof in the large. This completes the proof of the theorem. \square

References

- ABDULKADIROĞLU, A., Y.-K. CHE, AND Y. YASUDA (2015): “Expanding choice in school choice,” *American Economic Journal: Microeconomics*, 7(1), 1–42.
- ABDULKADIROĞLU, A., P. A. PATHAK, AND A. E. ROTH (2009): “Strategy-proofness versus efficiency in matching with indifferences: Redesigning the NYC high school match,” *American Economic Review*, 99(5), 1954–78.
- ABDULKADIROĞLU, A., AND T. SÖNMEZ (2003): “School choice: A mechanism design approach,” *American economic review*, 93(3), 729–747.
- AGARWAL, N., AND E. BUDISH (2021): “Market design,” in *Handbook of Industrial Organization*, vol. 5, pp. 1–79. Elsevier.
- AKBARPOUR, M., AND A. NIKZAD (2020): “Approximate random allocation mechanisms,” *The Review of Economic Studies*, 87(6), 2473–2510.
- AZEVEDO, E., AND E. BUDISH (2019): “Strategy-proofness in the large,” *The Review of Economic Studies*, 86, 81–116.
- BICHLER, M., AND S. MERTING (2021): “Randomized scheduling mechanisms: Assigning course seats in a fair and efficient way,” *Production and Operations Management*.
- BOUTILIER, C., AND H. H. HOOS (2001): “Bidding languages for combinatorial auctions,” in *International Joint Conference on Artificial Intelligence*, vol. 17, pp. 1211–1217.
- BRERO, G., B. LUBIN, AND S. SEUKEN (2018): “Combinatorial Auctions via Machine Learning-based Preference Elicitation,” in *International Joint Conf. on Artificial Intelligence*, pp. 128–136.
- BUDISH, E. (2011): “The combinatorial assignment Problem: Approximate competitive equilibrium from equal incomes,” *Journal of Political Economy*, 119, 1061–1061.
- BUDISH, E., G. CACHON, J. KESSLER, AND A. OTHMAN (2017): “Course match: A large-scale implementation of approximate competitive equilibrium from equal incomes for combinatorial allocation,” *Operations Research*, 65(2), 314–336.
- BUDISH, E., AND E. CANTILLON (2012): “The multi-unit assignment problem: Theory and evidence from course allocation at Harvard,” *American Economic Review*, 102(5), 2237–71.

- BUDISH, E., Y.-K. CHE, F. KOJIMA, AND P. MILGROM (2013): “Designing random allocation mechanisms: Theory and applications,” *American Economic Review*, 103(2), 585–623.
- BUDISH, E., AND J. B. KESSLER (2022): “Can market participants report their preferences accurately (enough)?,” *Management Science*, 68(2), 1107–1130.
- CRAWFORD, G. S., R. GRIFFITH, AND A. IARIA (2021): “A survey of preference estimation with unobserved choice set heterogeneity,” *Journal of Econometrics*, 222(1), 4–43.
- CROMME, L. J., AND I. DIENER (1991): “Fixed point theorems for discontinuous mapping,” *Mathematical Programming*, 51(1), 257–267.
- DIEBOLD, F., AND M. BICHLER (2017): “Matching with indifferences: A comparison of algorithms in the context of course allocation,” *European Journal of Operations Research*, 260(1), 268–282.
- ECHENIQUE, F., A. MIRALLES, AND J. ZHANG (2021): “Constrained pseudo-market equilibrium,” *American Economic Review*, 111(11), 3699–3732.
- ECHENIQUE, F., AND J. OVIEDO (2006): “A theory of stability in many-to-many matching markets,” *Theoretical Economics*, 1(2), 233–273.
- ERDIL, A., AND H. ERGIN (2008): “What’s the matter with tie-breaking? Improving efficiency in school choice,” *American Economic Review*, 98(3), 669–89.
- (2017): “Two-sided matching with indifferences,” *Journal of Econ. Theory*, 171, 268–292.
- FACK, G., J. GRENET, AND Y. HE (2019): “Beyond truth-telling: Preference estimation with centralized school choice and college admissions,” *Amer. Econ. Review*, 109(4), 1486–1529.
- FOLEY, D. K. (1967): “Resource allocation and the public sector,” *Yale Economics Essays*.
- GALE, D., AND L. S. SHAPLEY (1962): “College admissions and the stability of marriage,” *The American Mathematical Monthly*, 69(1), 9–15.
- HE, Y., A. MIRALLES, M. PYCIA, AND J. YAN (2018): “A pseudo-market approach to allocation with priorities,” *American Economic Journal: Microeconomics*, 10(3), 272–314.
- HYLLAND, A., AND R. ZECKHAUSER (1979): “The efficient allocation of individuals to positions,” *Journal of Political economy*, 87(2), 293–314.
- KALAI, E. (2004): “Large Robust Games,” *Econometrica*, 72(6), 1631–1665.
- KAMADA, Y., AND F. KOJIMA (2023): “Fair Matching under Constraints: Theory and Applications,” *The Review of Economic Studies*, p. rdad046.
- KELSO, A. S., AND V. P. CRAWFORD (1982): “Job matching, coalition formation, and gross substitutes,” *Econometrica*, pp. 1483–1504.
- KOJIMA, F., AND P. A. PATHAK (2009): “Incentives and stability in large two-sided matching markets,” *American Economic Review*, 99(3), 608–627.
- LIN, Y.-S., H. NGUYEN, T. NGUYEN, AND K. ALTINKEMER (2022): “Allocation with weak priorities and general constraints,” *Operations Research*, 70(5), 2804–2819.
- MIRALLES, A., AND M. PYCIA (2021): “Foundations of pseudomarkets: Walrasian equilibria for discrete resources,” *Journal of Economic Theory*, 196, 105303.
- NGUYEN, H., T. NGUYEN, AND A. TEYTELBOYM (2021): “Stability in matching markets with

- complex constraints,” *Management Science*, 67(12), 7438–7454.
- NGUYEN, T., A. PEIVANDI, AND R. VOHRA (2016): “Assignment problems with complementarities,” *Journal of Economic Theory*, 165, 209–241.
- NGUYEN, T., AND R. VOHRA (2022): “(Near)-substitute preferences and equilibria with indivisibilities,” Discussion paper, University of Pennsylvania.
- OTHMAN, A., C. PAPADIMITRIOU, AND A. RUBINSTEIN (2016): “The complexity of fairness through equilibrium,” *ACM Transactions on Economics and Computation*, 4(4), 1–19.
- OTHMAN, A., T. SANDHOLM, AND E. BUDISH (2010): “Finding approximate competitive equilibria: efficient and fair course allocation,” in *AAMAS*, vol. 10, pp. 873–880.
- PYCIA, M. (2021): “Pseudomarkets,” in *Online and Matching-Based Market Design*, ed. by F. Echenique. Cambridge University Press.
- ROTH, A. E. (1984): “Stability and polarization of interests in job matching,” *Econometrica*, pp. 47–57.
- (2018): “Marketplaces, Markets, and Market Design,” *American Economic Review*, 108(7), 1609–1658.
- ROTH, A. E., AND M. A. O. SOTOMAYOR (1990): *Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis*. Cambridge University Press.
- RUSZNÁK, A., P. BIRÓ, AND R. FLEINER (2021): “Seat transfers in the course allocation mechanism of Eötvös Loránd University,” in *2021 IEEE 15th International Symposium on Applied Computational Intelligence and Informatics (SACI)*, pp. 503–508. IEEE.
- SANDHOLM, T., AND C. BOUTILIER (2006): “Preference elicitation in combinatorial auctions,” *Combinatorial Auctions*, 10.
- SCHLEGEL, J. C., AND A. MAMAGEISHVILI (2020): “Welfare theorems for random assignments with priorities,” *Games and Economic Behavior*, 124, 62–81.
- SHAPLEY, L., AND H. SCARF (1974): “On cores and indivisibility,” *Journal of Mathematical Economics*, 1(1), 23–37.
- SÖNMEZ, T., AND M. U. ÜNVER (2010): “Course bidding at business schools,” *International Economic Review*, 51(1), 99–123.
- SOUMALIAS, E., B. ZAMANLOOY, J. WEISSTEINER, AND S. SEUKEN (2023): “Machine learning-powered course allocation,” *arXiv preprint arXiv:2210.00954*.
- VAZIRANI, V., AND M. YANNAKAKIS (2021): “Computational complexity of the Hylland-Zeckhauserscheme for one-sided matching markets,” in *Innovations in Theoretical Computer Science*.

(For Online Publication)

B Proof of Theorem 1 (Omitted Details)

We have shown in the main text how the arguments of Steps 1-3 of Theorem 1 of [Budish \(2011\)](#) need to be modified for the setting with priority-specific course prices. The adaptation of Steps 4-9 closely follows the original proof. However, we first establish that it is possible to choose budgets b_s and τ_{s,x_s} such that at most M budget constraints intersect for any $\tilde{t} \in \tilde{\mathcal{T}}$.

Lemma B1. *One can choose taxes $\{\tau_{s,x_s}\}_{s \in \mathcal{S}, x_s \subseteq \mathcal{C}}$ that satisfy the following conditions:*

- (i) *Taxes are small ($-\varepsilon < \tau_{s,x_s} < \varepsilon$);*
- (ii) *Taxes favor more preferred bundles ($\tau_{s,x_s} > \tau_{s,x'_s}$ for $x'_s \succ_s x_s$);*
- (iii) *The inequality bounds are preserved ($-1 \leq \min_{s,x_s}(b_s + \tau_{s,x_s}) \leq \max_{s,x_s}(b_s + \tau_{s,x_s}) \leq 1 + \beta$);*
- (iv) *No perturbed budgets are equal ($b_s + \tau_{s,x_s} \neq b_{s'} + \tau_{s',x_{s'}}$);*
- (v) *There is no auxiliary price vector $\tilde{t} \in \tilde{\mathcal{T}}$ at which more than M budget constraints $H(s, x_s)$ intersect.*

Proof. Let us fix $x = \{x_s\}_{s \in \mathcal{S}}$. [Budish \(2011\)](#) establishes the possibility to choose $\{\tau_{s,x_s}\}_{s \in \mathcal{S}, x_s \subseteq \mathcal{C}}$ that satisfy the first four conditions. We now establish that it is always possible to slightly change taxes such that condition (v) is also satisfied.

For this purpose, let us assume that more than M budget constraints $H(s, x_s)$ intersect and denote

$$\mathcal{I} = \{s \in \mathcal{S} : \cap_s H(s, x_s) \neq \emptyset\},$$

with $|\mathcal{I}| > M$. For each $\tilde{t} \in \tilde{\mathcal{T}}$, consider prices $\{p_{c,r}(\tilde{t})\}_{c \in \mathcal{C}, r \in \mathcal{R}}$ defined by equation (A.1) and the cutoffs defined by (A.2). The definition of cutoffs $r_c^*(\tilde{t})$ implies that

$$\forall s \in \mathcal{I}, c \in \mathcal{C} : r_{s,c} > r_c^*(\tilde{t}), \quad x_{s,c} \cdot p_{c,r_{s,c}}(\tilde{t}) = 0.$$

In addition, the definition of cutoffs implies that $p_{c,r_{s,c}} \geq \bar{b}$ for $1 \leq r_{s,c} < r_c^*(\tilde{t})$. Therefore, a seat in course c is not allocated to agent s for $r_{s,c} < r_c^*(\tilde{t})$; that is, $x_{s,c} = 0$. Hence, we have

also

$$\forall s \in \mathcal{I}, c \in \mathcal{C} : r_{s,c} < r_c^*(\tilde{t}), \quad x_{s,c} \cdot p_{c,r_{s,c}}(\tilde{t}) = 0.$$

Therefore, we obtain that $x_{s,c} \cdot p_{c,r_{s,c}}(\tilde{t})$ might be non-zero only if $r_{s,c} = r_c^*(\tilde{t})$. That is, entries in agent s 's budget constraint, $s \in \mathcal{I}$, are non-zero only if $r_{s,c} = r_c^*(\tilde{t})$. Denote $p_c(\tilde{t}) \equiv p_{r_c^*(\tilde{t})}$. Hence, we obtain that the set of equations $\{s \in \mathcal{I} : H(s, x_s) = 0\}$ is the set of linear equations with coefficients $x_{s,c} \in \{0, 1\}$:

$$\begin{cases} x_{s,1} \cdot p_1(\tilde{t}) + x_{s,2} \cdot p_2(\tilde{t}) + \dots + x_{s,M} \cdot p_M(\tilde{t}) & = b_s + \tau_{s,x_s} \\ \dots & \dots \\ x_{s',1} \cdot p_1(\tilde{t}) + x_{s',2} \cdot p_2(\tilde{t}) + \dots + x_{s',M} \cdot p_M(\tilde{t}) & = b_{s'} + \tau_{s',x_{s'}} \end{cases}$$

for $s, s' \in \mathcal{I}$. Since $|\mathcal{I}| > M$ for any $\tilde{t} \in \mathcal{T}$, there are at most M independent linear equations of prices. So, the Rouché–Capelli theorem implies that we can choose τ_{s,x_s} such that only at most M equations are satisfied for any \tilde{t} .³¹ \square

Step 4. Similar to Theorem 1 of Budish (2011), if the price vector t^* is not on any budget constraint, then t^* is an exact competitive equilibrium price vector. Suppose that t^* is not on any budget constraint. Then, there is a neighborhood around t^* where each agent's demand is unchanging in price. At price t^* , f is continuous, and as a result, $F(t^*) = f(t^*)$.

- If $t^* = \tilde{t}^*$, then $F(\tilde{t}^*) = F(t^*) = f(t^*)$ and, thus, $t^* = \tilde{t}^* \in F(\tilde{t}^*) = f(\tilde{t}^*)$. Therefore, t^* is a fixed point. Hence, as we established earlier, it is an exact competitive equilibrium price vector.
- If $t^* \neq \tilde{t}^*$, we establish the following lemma.

Lemma B2. For any $\tilde{t} \in \tilde{\mathcal{T}} \setminus \mathcal{T}$, (i) $f(\tilde{t}) = f(\text{trunc}(\tilde{t}))$ and (ii) $F(\tilde{t}) \subseteq F(\text{trunc}(\tilde{t}))$.

Proof. Statement (i) follows from the definition of f . For statement (ii), consider a point y for which there exists a sequence $\tilde{t}^w \rightarrow \tilde{t}$, $\tilde{t}^w \neq \tilde{t}$ such that $f(\tilde{t}^w) \rightarrow y$. Consider $\text{trunc}(\tilde{t}^w)$. As $\text{trunc}(\cdot)$ is continuous, this sequence will converge to $\text{trunc}(\tilde{t})$. Statement (i) implies $f(\text{trunc}(\tilde{t}^w))$ converges to y . As a result, $y \in F(\tilde{t})$ implies that $y \in F(\text{trunc}(\tilde{t}))$ and, thus, $F(\tilde{t}) \subseteq F(\text{trunc}(\tilde{t}))$. \square

³¹Note that when we change τ_{s,x_s} , some other budget constraints might start intersecting. Since the set of possible intersecting budget constraints is finite and τ_{s,x_s} varies continuously, we can always choose τ_{s,x_s} without influencing the intersection property of the other budget constraints.

As a result, since $F(t^*) = f(t^*)$ and $\tilde{t}^* \in F(\tilde{t}^*)$, $\tilde{t}^* \in F(t^*) = f(t^*) = f(\tilde{t}^*)$ and, thus, $\tilde{t}^* = f(\tilde{t}^*)$. Therefore, t^* is an exact competitive equilibrium price vector.

Step 5. Next, suppose that t^* is on $1 \leq L \leq M$ budget constraints. We denote $\Phi = \{0, 1\}^L$ and construct a set of 2^L price vectors $\{t^\phi\}_{\phi \in \Phi}$ that satisfy the following conditions:

1. Each t^ϕ is close enough to t^* such that there is a path from t^ϕ to t^* that does not cross any budget constraint.
2. Each t^ϕ is on the “affordable” side of the ℓ th budget constraint if $\phi_\ell = 0$ and is on the “unaffordable” side if $\phi_\ell = 1$.

To construct vectors $\{t^\phi\}_{\phi \in \Phi}$, note that each of the L intersecting budget constraints defines two sets:

$$H_\ell^0 = \left\{ \tilde{t} \in \tilde{\mathcal{T}} : \sum_{c \in \mathcal{C}} x_{s_\ell c} \max(\tilde{t}_c - (r_{s_\ell, c} - 1)\bar{b}, 0) \leq b_{s_\ell} + \tau_{s_\ell, x_{s_\ell}} \right\}$$

$$H_\ell^1 = \left\{ \tilde{t} \in \tilde{\mathcal{T}} : \sum_{c \in \mathcal{C}} x_{s_\ell c} \max(\tilde{t}_c - (r_{s_\ell, c} - 1)\bar{b}, 0) > b_{s_\ell} + \tau_{s_\ell, x_{s_\ell}} \right\}$$

The first set delineates the set of prices for which agent s_ℓ can afford schedule x_{s_ℓ} , whereas the second set delineates the set of prices for which agent s_ℓ can't afford x_{s_ℓ} . Let $\phi = (\phi_1, \dots, \phi_L) \in \Phi$ be an L -dimensional vector of zeros and ones, and the polytope $\pi(\phi) := \bigcap_{\ell=1}^L H_\ell^{\phi_\ell}$ be the set of points in $\tilde{\mathcal{T}}$ that belongs to the intersection of sets indexed by ϕ . Let $H = \{H(s, x_s)\}_{s \in \mathcal{I}, x_s \subseteq \mathcal{C}}$ be the finite set of all budget constraints formed by any student-schedule pair (s, x_s) . We then define

$$\delta < \inf_{\tilde{t}'' \in \tilde{\mathcal{T}}, H \in \mathcal{H}} \{ \|t^* - \tilde{t}''\|_2 : \tilde{t}'' \in H, t^* \notin H \},$$

which denotes the distance such that any budget constraint that t^* does not belong to is further than δ away from t^* . Let $B_\delta(t^*)$ be a δ -ball of t^* . Now, for each $\phi \in \Phi$ we define \tilde{t}^ϕ to an arbitrary element of $\pi(\phi) \cap B_\delta(t^*)$. Such price vectors satisfy the two requirements outlined above.

Step 6. We now show that a perfect market-clearing excess demand vector lies in the convex hull of $\{z(t^\phi)\}_{\phi \in \Phi}$. For this purpose, we first show that for any $y \in F(t^*)$, we must

have $y = t^* + \sum_{\phi \in \Phi} \lambda^\phi z(t^\phi)$, where $\sum_{\phi \in \Phi} \lambda^\phi = 1$. Take some $y \in F(t^*)$. Consider a sequence $t^w \rightarrow t^*$, $t^* \neq t^w \in \tilde{\mathcal{T}}$ such that $f(t^w) \rightarrow y'$. Note that the sequence t^w consists of a finite number of subsequences $t^{w,\phi} \in \pi(\phi) \cap B_\delta(t^*)$ for some $\phi \in \Phi$.³² Since all elements of the subsequence $t^{w,\phi}$ are on the same side of set H_ℓ^ϕ , every agent has the same choice at every point. Hence, if the subsequence has an infinite number of elements, we must have $t^{w,\phi} \rightarrow t^*$ and

$$f(t^{w,\phi}) \rightarrow t^* + \gamma z(t^\phi).$$

Therefore, the limit of $f(t^w)$ for the original sequence t^w must be also $t^* + \gamma z(t^{\phi'})$ for some $\phi' \in \Phi$. So, $y' = t^* + \gamma z(t^{\phi'})$. Since, by definition, y is a convex combinations of such y' , we must have $y = t^* + \sum_{\phi \in \Phi} \lambda^\phi z(t^\phi)$, where $\sum_{\phi \in \Phi} \lambda^\phi = 1$.

Lemma B2 implies that $\tilde{t}^* \in F(\tilde{t}^*) \subseteq F(t^*)$. Hence, we must have

$$\tilde{t}^* = t^* + \sum_{\phi \in \Phi} \lambda^\phi z(t^\phi).$$

for some $\{\lambda_\phi\}_{\phi \in \Phi}$ with $\sum_{\phi \in \Phi} \lambda^\phi = 1$. We also denote

$$\zeta = \sum_{\phi \in \Phi} \lambda^\phi z(t^\phi) = \frac{\tilde{t}^* - t^*}{\gamma}.$$

We note that $\zeta \leq 0$ and $\zeta_c < 0$ imply $t_c^* = 0$. Hence, vector ζ is a perfect market-clearing excess demand vector, and it lies in the convex hull of $\{z(t^\phi)\}_{\phi \in \Phi}$.

Steps 7-9. The structure of excess demand has the same geometric structure as in Budish (2011). In particular, denote L' as the number of agents whose budget constraints intersect at price t^* . We rename agents such that $s = 1, \dots, L'$. We denote the number of budgets of student s that intersect at t^* as w_s . Since at most M budgets constraints can intersect, we must have $L \equiv \sum_{s=1}^{L'} w_s \leq M$. We also denote the bundles pertaining to s 's budget constraints as $x_s^1 \succ \dots \succ x_s^{w_s}$.

Similarly to Budish (2011), we consider bundles that s demands at prices near t^* . In the set $H^0(s, x_s^1)$, agent s can purchase her favorite bundle x_s^1 . Hence, one does not need to know whether prices belong to sets $H^0(s, x_s^2), \dots, H^0(s, x_s^{w_s})$. Let us denote the demand for prices at $H^0(s, x_s^1) \cap B_\delta(t^*) \cap \tilde{\mathcal{T}}$ as d_s^0 . Similarly, we consider prices in $H^1(s, x_s^m) \cap H^0(s, x_s^{m+1})$

³²Some subsequences can have only a finite number of elements.

and denote the corresponding demands as d_s^m for $m = 1, \dots, w_s$. Overall, agent $s = 1, \dots, L'$ purchase $w_s + 1$ distinct bundles at prices near to t^* .

Let us denote the excess demand of the remaining agents as

$$z_{S \setminus \{1, \dots, L'\}}(t^*) = \sum_{s=L'+1}^S d_s(t^*) - q.$$

Hence, a perfect market-clearing excess demand vector lies in the convex hull of $\{z(t^\phi)\}_{\phi \in \Phi}$ with the elements

$$z_{S \setminus \{1, \dots, L'\}}(t^*) + \sum_{s=1}^{L'} \sum_{f=1}^{w_s} a_s^f d_s^f$$

where $0 \leq a_s^f \leq 1$, $s = 1, \dots, L'$, $f = 1, \dots, w_s$ and $\sum_{f=1}^{w_s} a_s^f = 1$, $s = 1, \dots, L'$. Budish (2011) shows in Step 8 of Theorem 1 that there exists a vertex of the above geometric structure that is within $\sqrt{kM/2}$ distance from the perfect market-clearing excess demand vector. He also explains how to adjust agent budgets to find an approximate competitive equilibrium. These purely mathematical arguments remain unchanged from the original paper. \square

C Appendix: A Simple Student Utility Model

This section describes the calibration of the simple student utility model introduced in Section 4.1 based on the real-world university data. The data contains information about students from seven colleges $a = A, \dots, G$ and four years of study $y = 1, \dots, 4$.³³ The model assumes that the utility of student s from college a and year of study y for taking course c from college a' equals

$$u_{sc} = \theta_{aya'} + \varepsilon_{sc},$$

where $\theta_{aya'}$ is a fixed component and ε_{sc} is a random component with $\varepsilon_{sc} \sim N(0, \sigma)$ being independently and normally distributed random variables with zero mean and variance σ . Student's utilities are additive across courses. The outside option is normalized to 0. The standard deviation of the noise is normalized to one as well $\sigma = 1$.

Each student's choice set is limited up to 80 courses. The courses in student's choice set are drawn randomly such that the probability that a course from college a' is drawn equals

³³There are five years of study observed in the actual university data. We unite students in year four and year five cohorts as there are only a few students in year five cohort for some colleges.

the share of students from the same college-year who are enrolled to courses in college a' in the actual data. In addition, if the number of courses taken by college-year students in a given college is fewer than ten, the student choice sets are additionally enlarged by courses of the same or the next level in the same department.

Before we describe the process of how we calibrate the utility parameters, some discussion is necessary. The standard discrete choice identification techniques cannot be used to calibrate the student utility model, as the student choice sets are unobserved. We do not observe individual student's past enrollment and, hence, the set of courses for which the student satisfies course prerequisites. In addition, the student's choice set during the course allocation process is influenced by the courses chosen by the other students who are more senior or have earlier time-slots and, hence, students' choice sets are determined endogenously. There are also complications associated with the changing course reserves during the course allocation process described in detail in Section 4.1.

This forces us to take a stand on how student's choice sets are formed. We assume that courses are drawn in proportion to the number of courses assigned in the actual data in order to mimic closely unobserved choice sets.³⁴ However, any changes in the way we draw the course lists could potentially change the calibrated parameter values. Crawford, Griffith, and Iaria (2021) provide an insightful discussion of preference estimation techniques with unobserved choice sets. We also want to mention that it is not clear whether increasing the size of the course lists is relevant as the previous studies report that students typically consider only a limited subset of courses in a given semester (see Budish and Cantillon, 2012; Diebold and Bichler, 2017). We explore how the performance of the four allocation mechanisms changes for various sizes of choice sets (and the size of the noise parameter) using simulations in Appendix D.3.

Course reserves adjustments. We use simulations to calibrate student utility parameters. However, we need first to perform adjustments to course reserves. Course reserves are typically set large at the beginning of the course allocation process and relaxed at later stages. Hence, the final course allocation could violate the initial course reserves. In fact,

³⁴We also considered an alternative model on how students' choice sets are formed. We assumed that each student has a fixed choice set that coincides with the set of courses taken by at least one students from the same department-year in the data. For some department-years, the choice set might be as large as 150 courses. The calibration of this modelled to the results in which many students enroll in the maximum possible number of courses $k = 5$ and some students enroll in a few or even zero courses, which significantly differs from the distribution of the number of courses taken by student in the data.

170 out of 756 courses violate initial course reservations in the data.

We perform the minimum course reserve adjustment using the following procedure. Consider a course with multiple reserves based on the year and department of students (see Table 3 in Section 4.1) and the final course allocation. Determining if the set of students assigned to the course can satisfy course reservation constraints is difficult, as a student can be attributed to several reservations. For example, a course can have one reservation for students of any year of study in a certain department and another reservation for students in a certain year of study and a given set of departments. It might not be clear to which reservation one should attribute each student. We use the *Hall's Marriage Theorem*, which provides the necessary and sufficient condition for the existence of a matching where each student is assigned exactly one seat in the course and satisfies course reserves. We explain the details in [Supplementary Materials](#). Once we identify course reserve violations, we reduce the number of reserves one at a time, starting with reservations for larger years of study, until the course allocation satisfies course reserve constraints. In addition, we decrease the number of reserved seats for the courses in colleges A and B by 15%. The departments in these colleges typically assign course reserves only to their students and the number of reserved seats is close to the full course capacity. This additional adjustment appears to be necessary for a stable behavior of the calibration process.

Calibration. To calibrate parameters $\theta_{aya'}$, $a, a' = A, \dots, G$, $y = 1, \dots, 4$, we use the Random Serial Dictatorship with set-asides, the adjusted course reserves, the order in which students choose the courses in the university data, and the final course allocation. The calibration proceeds as follows.

- Before the start of the process, we draw 100 times 80-long course lists with non-zero utilities for each student. We also draw 100 vectors of random utility components. If we denote the number of students as $S = 6023$, there are in total $S \times 100 \times 80$ draws of ε_{sc} . We fix these draws throughout the calibration process.
- For given parameters θ , we calculate student utilities u_{sc} using the random components.
- For each $j = 1, \dots, 100$ we calculate the outcome of the Random Serial Dictatorship with set-asides taking the order of student choices provided in the data and the adjusted course reserves. We also calculate the aggregate matching on the college-year-college level to obtain $m_{aya'}^j$ - the number of students from college-year (a, y) assigned to

courses at college a' . There are $7 \times 4 \times 7$ matching student numbers.

- We average the matching matrices across random draws $\bar{m} = \frac{1}{100} \sum_j m^j$ to obtain the average aggregate matching on college-year-college level.
- We compare the aggregate matching in the data m^{DATA} (see Table C1) with \bar{m} and adjust each $\theta_{aya'}$ proportionally to the difference in the number of students assigned

$$\lambda \frac{m_{aya'}^{DATA} - \bar{m}_{aya'}}{m_{aya'}^{DATA}},$$

where $\lambda > 0$ is the size of the adjustment step.

- For adjusted parameters θ , we recalculate matching m . Then, we evaluate the utility parameters based on the objective $\|m^{DATA} - m\|^2$. If the adjustment leads to a smaller value of the objective, we keep the new parameters and proceed to the next step. If the value of the objective becomes larger or the change in the objective is small, we decrease the size of the adjustment step λ , and recalculate matching m .³⁵ We terminate the algorithm if the objective is less than 1 or improves less than 1% during the six consecutive iterations.

Table C2 presents the parameters of the simple utility model calibrated using the above procedure. Each table shows parameters θ for students from one college. Each cell in each table presents the average utility of students for each year of study (column) taking courses from each college (row). When no students from the college-year (a, y) are assigned to courses at college a' in the data, we leave the corresponding parameter as NA. For the calibrated parameters, the student utility model fits well the observed aggregate matching with $\|m^{DATA} - m\|^2 < 1$.

Several points need to be highlighted about the calibrated values. First, the parameters corresponding to courses from the same college are typically the largest among seven colleges. The revealed preferences argument also implies that students should prefer to take courses from their own colleges. The only exception is college C , which does not offer first-year courses (all of its courses correspond to interdisciplinary programs).

Second, many parameters are smaller than 0, implying that the value of the outside option could be an important factor in a student's choice. To explore this point in more detail,

³⁵The threshold for the change in the objective is 5% for Years 3 and 4 and 1% for Years 1 and 2.

we consider an alternative normalization, in which the average student utility from taking the course from the same college is normalized to 1 and the value of the outside option can vary with respect to the year of study and student college. The value of the outside option corresponds to the value of taking non-semester-long courses, doing a part-time job (which is relevant for more senior students), or engaging in extracurricular activities. The alternative normalization results in a parallel shift in utilities, which does not change the calibration as it does not change the outcomes of the RSD mechanism. At the same time, the alternative normalization does have implications for the welfare comparison across mechanisms we explore in Appendix D. The calibrated parameters for an alternative normalization are presented in Table C3.

Finally, we want to note that one should be careful interpreting the above utility parameters. First, we consider only the calibration of the student utility model. We do not have an identification result. Second, the model specifies student utility only at the college-year level. This is not very precise, as there are many departments within each college and each department typically has several majors. Third, the model cannot disentangle student preferences from physical course enrollment constraints such as course prerequisites, schedule time, and location. To truly identify student preferences, one should consider a more fine-tuned model that can separate student preferences from physical constraints. We leave this exciting, but difficult exercise for future research.

College A

	Year 1	Year 2	Year 3	Year 4
A	565	514	448	439
B	2	3	4	4
C	5	17	11	19
D	189	98	68	109
E	22	13	7	5
G	22	30	34	51
F	8	21	20	24

College B

	Year 1	Year 2	Year 3	Year 4
A	13	19	23	62
B	431	931	770	596
C	2	10	5	21
D	293	311	242	303
E	657	331	111	88
G	244	169	192	153
F	25	64	63	104

College C

	Year 1	Year 2	Year 3	Year 4
A	18	31	28	33
B	3	2	3	6
C	0	1	22	25
D	86	113	95	105
E	99	123	58	34
G	41	41	43	38
F	0	7	12	7

College D

	Year 1	Year 2	Year 3	Year 4
A	11	22	35	60
B	4	11	25	11
C	3	11	9	20
D	901	949	927	550
E	293	178	90	61
G	161	168	191	156
F	21	141	118	76

College E

	Year 1	Year 2	Year 3	Year 4
A	6	6	6	18
B	27	27	11	12
C	0	1	1	2
D	167	172	189	184
E	492	529	361	203
G	142	130	122	80
F	3	14	18	27

College F

	Year 1	Year 2	Year 3	Year 4
A	1	11	15	15
B	19	12	11	10
C	1	5	14	4
D	232	173	140	118
E	167	153	133	57
G	447	504	369	173
F	8	28	20	10

College G

	Year 1	Year 2	Year 3	Year 4
A	7	8	4	12
B	1	4	1	3
C	0	1	2	3
D	313	132	138	138
E	125	30	21	2
G	50	32	21	37
F	192	425	320	285

TABLE C1: The aggregate matching observed in the data. Each table presents the matching for students from one college. Each cell presents the number of seats occupied by students from year of study (column) in courses from given college (row).

	Year 1	Year 2	Year 3	Year 4
A	-0.20	-0.95	-1.30	-1.53
B	-1.58	-1.73	-2.00	-2.07
C	-2.05	-1.33	-1.60	-1.35
D	-1.20	-1.74	-1.77	-1.82
E	-1.74	-1.84	-1.83	-2.09
G	-1.69	-1.77	-1.67	-1.80
F	-1.96	-1.26	-1.73	-1.84

	Year 1	Year 2	Year 3	Year 4
A	-1.92	-1.63	-1.78	-1.83
B	-0.70	-0.59	-1.19	-1.47
C	-2.51	-1.85	-2.16	-1.68
D	-1.58	-1.59	-1.66	-1.83
E	-1.09	-1.31	-1.65	-1.84
G	-0.98	-1.61	-1.68	-1.85
F	-1.96	-1.64	-1.71	-1.82

	Year 1	Year 2	Year 3	Year 4
A	-0.78	-1.22	-1.19	-1.67
B	-1.14	-1.24	-1.60	-1.43
C	-1.81	-2.14	-0.18	-0.50
D	-1.05	-1.42	-1.47	-1.68
E	-0.90	-1.32	-1.43	-1.68
G	-0.40	-1.36	-1.49	-1.71
F	-0.99	-1.25	-1.49	-1.39

	Year 1	Year 2	Year 3	Year 4
A	-1.62	-1.40	-1.52	-1.63
B	-1.76	-1.36	-1.59	-1.62
C	-2.31	-1.59	-1.77	-1.30
D	-0.91	-1.29	-1.38	-1.60
E	-1.24	-1.40	-1.56	-1.67
G	-1.07	-1.40	-1.51	-1.61
F	-1.58	-1.37	-1.50	-1.61

	Year 1	Year 2	Year 3	Year 4
A	-1.57	-1.80	-1.79	-1.72
B	-1.31	-0.95	-1.51	-1.44
C	-2.52	-2.53	-2.31	-2.12
D	-1.42	-1.60	-1.60	-1.67
E	-1.02	-1.21	-1.44	-1.61
G	-1.05	-1.37	-1.57	-1.71
F	-2.22	-1.51	-1.71	-1.70

	Year 1	Year 2	Year 3	Year 4
A	-1.91	-1.65	-1.82	-1.40
B	-1.26	-1.52	-1.71	-1.61
C	-2.42	-1.96	-1.52	-1.96
D	-1.25	-1.71	-1.79	-1.83
E	-1.38	-1.58	-1.47	-1.70
G	-0.48	-1.16	-1.39	-1.67
F	-2.05	-1.77	-1.82	-1.61

	Year 1	Year 2	Year 3	Year 4
A	-1.96	-1.73	-1.83	-1.99
B	-2.22	-1.91	-2.21	-2.32
C	-2.43	-2.03	-2.01	-2.19
D	-0.62	-1.49	-1.47	-1.62
E	-1.07	-1.62	-1.55	-2.44
G	-1.39	-1.64	-1.66	-1.82
F	-0.82	-0.68	-0.93	-1.21

TABLE C2: The calibrated parameters of the simple student utility model with zero outside option. Each table presents the utility parameters for students from one college. Each cell presents the utility of students from one year of study (column) taking courses from one college (row).

	Year 1	Year 2	Year 3	Year 4
A	1.00	1.00	1.00	1.00
B	-0.38	0.22	0.30	0.46
C	-0.86	0.62	0.70	1.18
D	0.00	0.22	0.53	0.71
E	-0.54	0.12	0.47	0.44
G	-0.49	0.18	0.63	0.73
F	-0.76	0.70	0.57	0.69
O	1.20	1.95	2.30	2.53

	Year 1	Year 2	Year 3	Year 4
A	-0.22	-0.04	0.40	0.64
B	1.00	1.00	1.00	1.00
C	-0.81	-0.26	0.02	0.79
D	0.11	-0.01	0.53	0.64
E	0.60	0.27	0.54	0.64
G	0.72	-0.03	0.51	0.63
F	-0.26	-0.05	0.47	0.65
O	1.70	1.59	2.19	2.47

	Year 1	Year 2	Year 3	Year 4
A	2.03	1.92	-0.01	-0.18
B	1.67	1.90	-0.42	0.07
C	NA	1.00	1.00	1.00
D	1.77	1.72	-0.29	-0.18
E	1.91	1.82	-0.25	-0.19
G	2.42	1.78	-0.31	-0.21
F	NA	1.89	-0.31	0.11
O	2.81	3.14	1.18	1.50

	Year 1	Year 2	Year 3	Year 4
A	0.30	0.89	0.86	0.97
B	0.16	0.93	0.80	0.98
C	-0.40	0.70	0.61	1.30
D	1.00	1.00	1.00	1.00
E	0.68	0.89	0.83	0.93
G	0.85	0.89	0.88	0.99
F	0.33	0.92	0.88	0.99
O	1.91	2.29	2.38	2.60

	Year 1	Year 2	Year 3	Year 4
A	0.44	0.41	0.65	0.89
B	0.71	1.26	0.93	1.17
C	NA	-0.32	0.13	0.49
D	0.60	0.61	0.84	0.94
E	1.00	1.00	1.00	1.00
G	0.97	0.84	0.87	0.90
F	-0.20	0.70	0.73	0.91
O	2.02	2.21	2.44	2.61

	Year 1	Year 2	Year 3	Year 4
A	-0.43	0.51	0.58	1.27
B	0.22	0.64	0.68	1.06
C	-0.94	0.20	0.88	0.71
D	0.23	0.46	0.61	0.84
E	0.10	0.59	0.92	0.97
G	1.00	1.00	1.00	1.00
F	-0.57	0.39	0.58	1.06
O	1.48	2.16	2.39	2.67

	Year 1	Year 2	Year 3	Year 4
A	-0.14	-0.05	0.10	0.22
B	-0.40	-0.23	-0.29	-0.11
C	NA	-0.36	-0.09	0.01
D	1.20	0.19	0.46	0.59
E	0.75	0.05	0.37	-0.23
G	0.43	0.04	0.27	0.38
F	1.00	1.00	1.00	1.00
O	1.82	1.68	1.93	2.21

TABLE C3: The calibrated parameters of the simple student utility model with variable. Each table presents the utility parameters for students from one college. Each cell presents the utility of students from one year of study (column) taking courses from one college (row). The last row shows the values of the outside option. The utility parameters corresponding to courses from the same college are normalized to one.

D Appendix: Additional Simulations

This section presents simulation results not present in the main text. Section D.1 supplies additional tables for the results in Section 4.3. In addition, we provide simulation results for the simple utility model with variable outside options (see Table C3). Both simulations assume that year-specific priority takes precedence over department-specific priority. Section D.2 supplements the main analysis by presenting the simulation results when department-specific priority takes precedence over year-specific priority. Section D.3 provides additional robustness checks by providing simulation results for several sizes of student choice sets and the standard deviation of the noise utility parameter.

D.1 Year-First Priorities: Additional Results

Table D1 below presents the actual levels of the mean student utility and the standard deviations of student utility for the analysis of Section 4.3. The simulated errors across runs are reported in parentheses. Using the standard T-test, for each mechanism $mech = \text{PMP}, \text{DA}, \text{DA(m)}$, we test null hypothesis $H_0 : \mu_{mech} = \mu_{RSD}$ against the alternative $H_a : \mu_{mech} \neq \mu_{RSD}$. We supply the corresponding mean student utility values in Table D1 with stars if the null hypothesis is rejected on 5% significance level. In addition, we perform F-tests for zero hypothesis $H_0 : \frac{\sigma_{mech}}{\sigma_{RSD}} \geq 1$ against the alternative $H_a : \frac{\sigma_{mech}}{\sigma_{RSD}} < 1$ for each of 100 runs. We supply the standard deviation of student utility with a star if the *maximum p-value across all 100 runs* is less than 5%.

In addition, we provide simulation results for the simple utility model with variable outside options. The university data contains several students who enroll in fewer than $k = 5$ courses. This is most common among students in their later years of study. To account for that, we consider an alternative normalization, in which the average student utility from taking the course from the same college is normalized to $\theta_{aya} = 1$, $a = A, \dots, G$, $y = 1, \dots, 4$ and the outside option for each student can vary with respect to the year of study and student college (parameters θ presented in Table C3 in Appendix C). Our interpretation is that the student's outside option corresponds to the value of enrolling in a non-semester-long course, participating in extra-curricular activities, doing a part-time job for more senior students, etc. If a student is enrolled in $k' < 5$ courses, the student's utility is supplemented with $5 - k'$ values of the outside option.

Mean Utility	Year 1	Year 2	Year 3	Years 4
Pseudo-Market with Priorities	3.39*(0.04)	2.98*(0.04)	2.32*(0.03)	1.62*(0.03)
Random Serial Dictatorship with optimal set-asides	3.17(0.04)	2.89(0.04)	2.29(0.03)	1.61(0.02)
Deferred Acceptance with single tie-breaking	3.16(0.04)	2.91*(0.04)	2.31*(0.03)	1.62(0.02)
Deferred Acceptance with multiple tie-breakings	3.06*(0.04)	2.89(0.04)	2.31*(0.03)	1.62(0.02)
St. Dev. of Utility	Year 1	Year 2	Year 3	Years 4
Pseudo-Market with Priorities	1.482*	1.368	1.220	1.051
Random Serial Dictatorship with optimal set-asides	1.669	1.409	1.227	1.051
Deferred Acceptance with single tie-breaking	1.646	1.403	1.227	1.052
Deferred Acceptance with multiple tie-breakings	1.505*	1.369	1.223	1.051

TABLE D1: The comparison of the four mechanisms using the simple utility model with *zero outside option* and with year-specific priorities taking precedence over department-specific priorities. The standard simulated deviations across runs are presented in parentheses.

The comparison of the four mechanisms using the alternative normalization is presented in Table D2. The major difference between Table D1 and Table D2 is the level of student utilities. The numbers in Table D2 account for the student utility from activities outside taking the courses. Inevitably, the student utility levels become higher. However, the absolute differences between mean student utilities across the allocations of different mechanisms remain the same. The alternative normalization also influences the standard deviation of student utility as the shift in student utilities differs by student college and year of study (see Table C3). Note that we do not need to recalculate the results about student envy (see Table 7) as they are not influenced by the alternative normalization.

Mean Utility	Year 1	Year 2	Year 3	Years 4
Pseudo-Market with Priorities	12.03*(0.04)	13.10*(0.04)	13.55*(0.03)	13.97*(0.03)
Random Serial Dictatorship with optimal set-asides	11.81(0.04)	13.01(0.04)	13.51(0.03)	13.96(0.02)
Deferred Acceptance with single tie-breaking	11.80(0.04)	13.03*(0.04)	13.53*(0.03)	13.97(0.03)
Deferred Acceptance with multiple tie-breakings	11.70*(0.04)	13.01(0.04)	13.53*(0.03)	13.97(0.02)
St. Dev. of Utility	Year 1	Year 2	Year 3	Years 4
Pseudo-Market with Priorities	2.022	1.984	1.671	1.571
Random Serial Dictatorship with optimal set-asides	2.175	2.049	1.680	1.570
Deferred Acceptance with single tie-breaking	2.155	2.042	1.677	1.570
Deferred Acceptance with multiple tie-breakings	2.080	2.027	1.677	1.571

TABLE D2: The comparison of the four mechanisms using the simple utility model with *variable outside option* and with year-specific priorities taking precedence over department-specific priorities. The standard simulated deviations across runs are presented in parentheses.

D.2 Department-First Priorities

Some U.S. universities consider department-specific priorities as more important than those based on year of study.³⁶ Here, we present additional simulation results comparing the performance of the four mechanisms assuming that department-specific priority takes precedence over year-specific priority. We assume each course has eight priority levels, where levels 5 through 8 correspond to students eligible for course reserves, and levels 1 through 4 correspond to students not eligible for course reserves. Note that the priority structure directly influences the performance of the PMP, DA, and DA(m) mechanisms, but not the RSD mechanism. However, the RSD mechanism relies on the calculation of *the optimal set-asides*, which are also determined by the priority structure.

Simulation results. As in Section 4.3, we consider the Random Serial Dictatorship (RSD) with optimal set-asides as a benchmark. Table D3 presents the performance of Pseudo-

³⁶For example, Dartmouth College explicitly states these priorities on its website (see https://www.dartmouth.edu/reg/registration/course_priorities.html).

Mean Utility	Year 1	Year 2	Year 3	Years 4
Pseudo-Market with Priorities	6.90%	3.66%	1.20%	-0.25%
Deferred Acceptance with single tie-breaking	1.49%	1.25%	0.43%	-0.33%
Deferred Acceptance with multiple tie-breakings	-0.33%	0.66%	0.29%	-0.41%
St. Dev. of Utility	Year 1	Year 2	Year 3	Years 4
Pseudo-Market with Priorities	-7.14%	-2.45%	-0.69%	-0.32%
Deferred Acceptance with single tie-breaking	-0.58%	0.07%	0.05%	-0.2%
Deferred Acceptance with multiple tie-breakings	-6.35%	-2.29%	-0.34%	-0.25%

TABLE D3: The performance of the Pseudo-Market with Priorities, Deferred Acceptance with single tie-breaking, and Deferred Acceptance with multiple tie-breakings mechanisms compared to the benchmark of Random Serial Dictatorship with optimal set-asides for each cohort of students and with *department-specific priorities taking precedence over year-specific priorities*.

Market with Priorities (PMP), Deferred Acceptance mechanism with single (DA) and multiple tie-breakings (DA(m)) compared to the benchmark over 100 simulation runs for different random utility draws.

The results are similar to the ones obtained in Section 4.3, with a slight change in the magnitude. The main differences are for Years 1 and 4. The relative importance of department-specific priorities pushes some senior students to a lower priority in PMP, DA, and DA(m) mechanisms. The same does not happen for the RSD mechanism, as the order in which students choose courses stays the same. As a result, there is an inferior allocation and smaller average utility for these students (-0.25% , -0.33% , and -0.41% , respectively). We also observe that first-year students receive slightly higher utility for DA and DA(m) mechanisms for the same reason. The performance of the PMP mechanism stays the same for these students.

The change in the standard deviation of student utilities stays almost the same as in Section 4.3. The major difference is for Year 4, where all three mechanisms deliver a reduced standard deviation versus the benchmark (-0.23% , -0.2% , and -0.25% , respectively). Unfortunately, these changes come with a decrease in the mean student utility as well. This is associated with the fact that the three mechanisms operate under a different set of constraints

	0 courses	1 course	2 courses	3 courses	4 courses	5 courses
Pseudo-Market with Priorities	98.69%	1.31%	0%	0%	0%	0%
Random Serial Dictatorship with optimal set-asides	93.61%	5.62%	0.68%	0.09%	0.008%	0.0005%
Deferred Acceptance with single tie-breaking	94.39%	5.02%	0.54%	0.05%	0.003%	0%
Deferred Acceptance with multiple tie-breakings	95.05%	4.78%	0.17%	0.003%	0%	0%

TABLE D4: The percentage of students who experience envy towards students of the same or lower priority for the four mechanisms with *department-specific priorities taking precedence over year-specific priorities*. The first column shows the information about students who experience no envy. The other columns show the percentage of students who experience envy bounded by 1, ..., 5 courses.

than the RSD benchmark.

The results regarding the percentage of students who experience envy are similar to those in the main part of the paper (see Table D4). The PMP mechanism leads to outcomes in which no student envies the schedule of any other student of the same or lower priority by more than a single course. No other mechanism satisfies such a property. Additionally, the PMP mechanism has the smallest percentage of students who experience envy bounded by a single course among the four mechanisms.

D.3 Additional Simulations: Comparative Statics

This section performs some additional robustness exercises. We take the student utility model calibrated based on the actual university data. Using this model, we consider the changes of two parameters a) the size of student choice sets 60, 70, 80, 90, and b) the standard deviation of the random component of student utility $\sigma = 0.75, 1, 1.25, 1.5$. We again report the performance of the PMP, DA, and DA(m) mechanisms relative to the RSD benchmark. In contrast to the main text, the reported simulation results are based on 50 random utility draws.

Figure D1 shows the comparative statics with respect to the size of student choice sets. The figure shows that the changes in the mean student utility compared to the RSD benchmark are small for third- and fourth-year students for all three mechanisms. For first- and second-year students, the ranking among the four mechanisms remains the same for all the

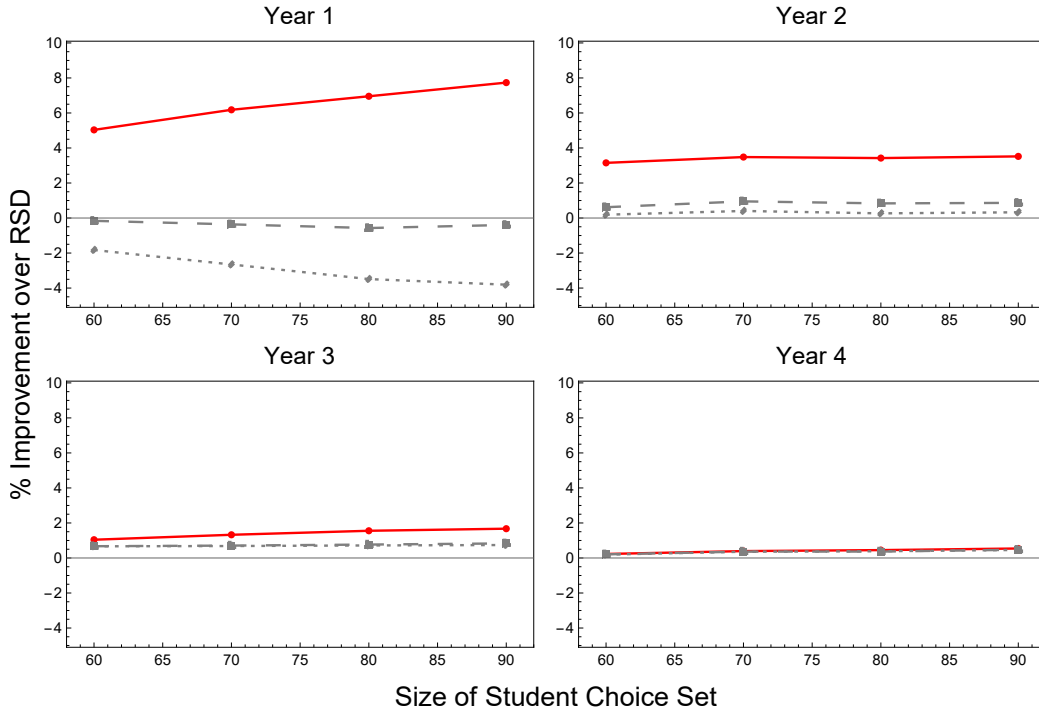
sizes of student choice sets. For first-year students, the change in the student mean utility compared to the RSD becomes larger as more courses are included in student choice sets. Such the change for the PMP mechanism increases from 5.05% for 60 courses to 7.73% for 90 courses.

The percentage changes in the standard deviation of student utility compared to the RSD benchmark are also small for third- and fourth-year students for all three mechanisms. For second-year students, the change remains approximately the same for all sizes of student choice sets. For first-year students, the drop in the standard deviation also grows as more courses are included in student choice sets. Such a drop in the standard deviation for the PMP mechanism increases from -6.43% for 60 courses to -13.17% for 90 courses. The relative ranking of the PMP and DA(m) mechanism reverses when student choice sets have 60 courses, but the magnitude of the reversal is small.

Figure D2 shows the comparative statics with respect to the standard deviation of the random component σ . As for the comparative statics for the size of student choice sets, the changes in the performance among four mechanisms are small for third- and fourth-year students. For second-year students, the change in the mean student utility remains stable, whereas the absolute change in the standard deviation of student utility slightly increases with the value of the noise parameter σ . For first-year students, the absolute changes in both the mean student utility and the standard deviation of student utility increase in σ . Mean student utility for the PMP outcomes increases from 3.37% for $\sigma = 0.75$ to 9.25% for $\sigma = 1.5$ and the standard deviation of student utility decreases from -0.7% for $\sigma = 0.75$ to -22.67% for $\sigma = 1.5$. We also observe a reversal in rankings for the standard deviation of student utility for the PMP and DA(m) mechanisms for first- and second-year students when $\sigma = 0.75$. However, the magnitude of the reversal is small.

Overall, the relative performance of the four mechanisms is relatively stable for a range of parameters of the simple student utility model. For the PMP mechanism, we observe an increase in the mean student utility and a decrease in the standard deviation of student utility compared to the RSD benchmark when student choice sets are larger and when student preferences over courses are more dispersed.

Mean Student Utility



Standard Deviation in Student Utility

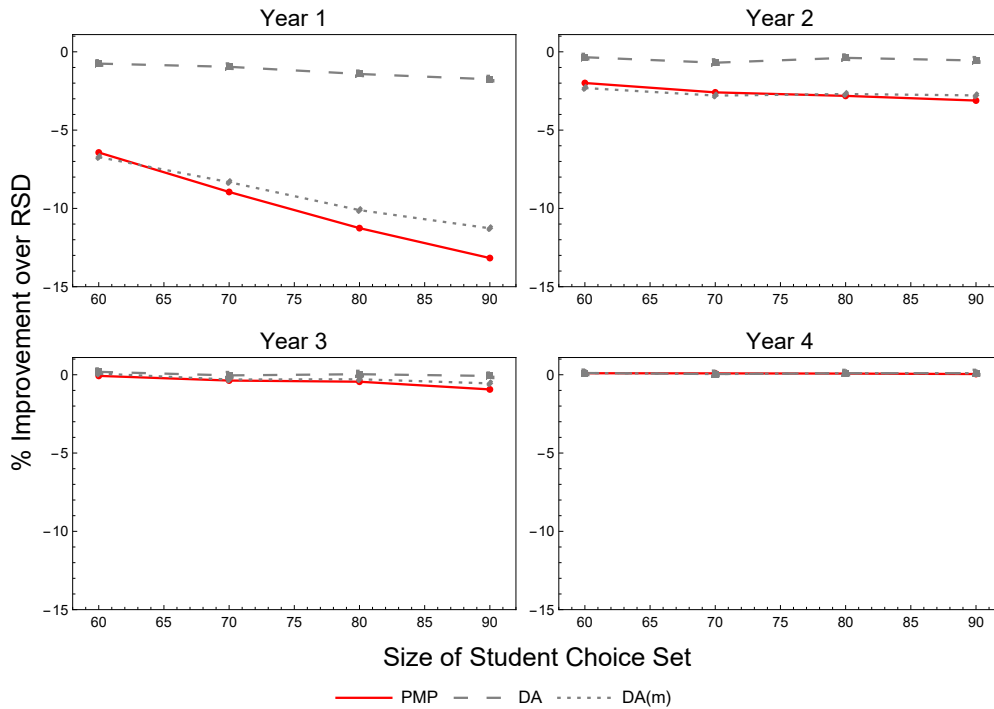


FIGURE D1: The performance of Pseudo-Market with Priorities (PMP), Deferred Acceptance with single tie-breaking (DA), and Deferred Acceptance with multiple tie-breaking DA(m) compared to the Random Serial Dictatorship (RSD) with optimal set-asides when students choice set consists of 60, 70, 80, and 90 courses. The results are based on 50 runs with different random utility draws for each parameter value.

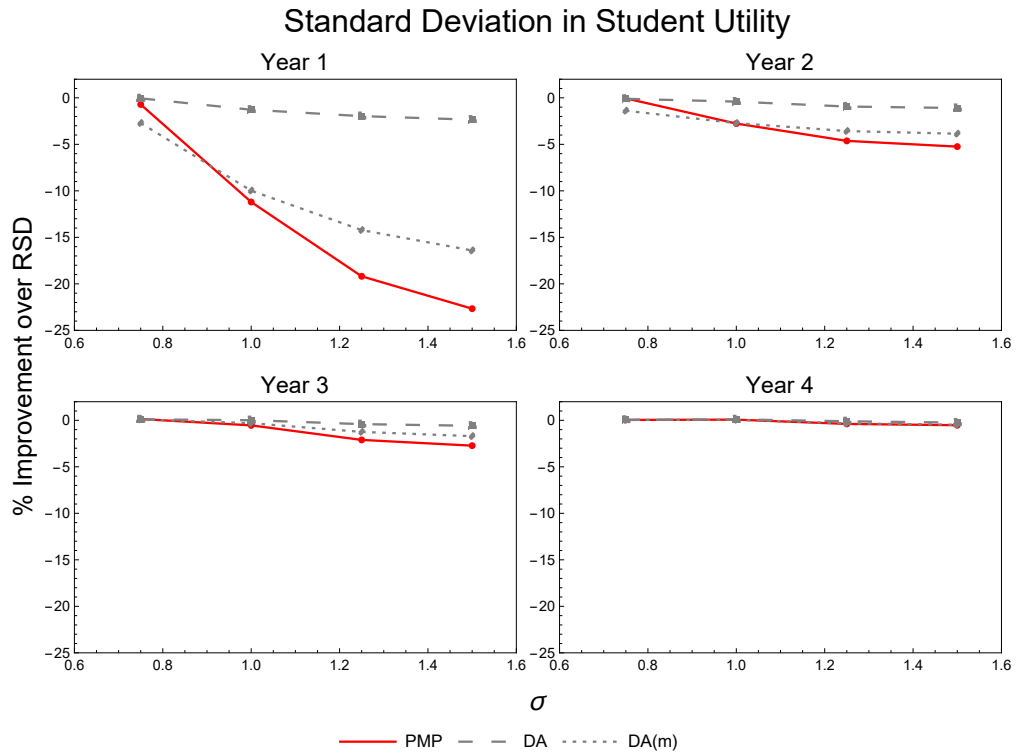
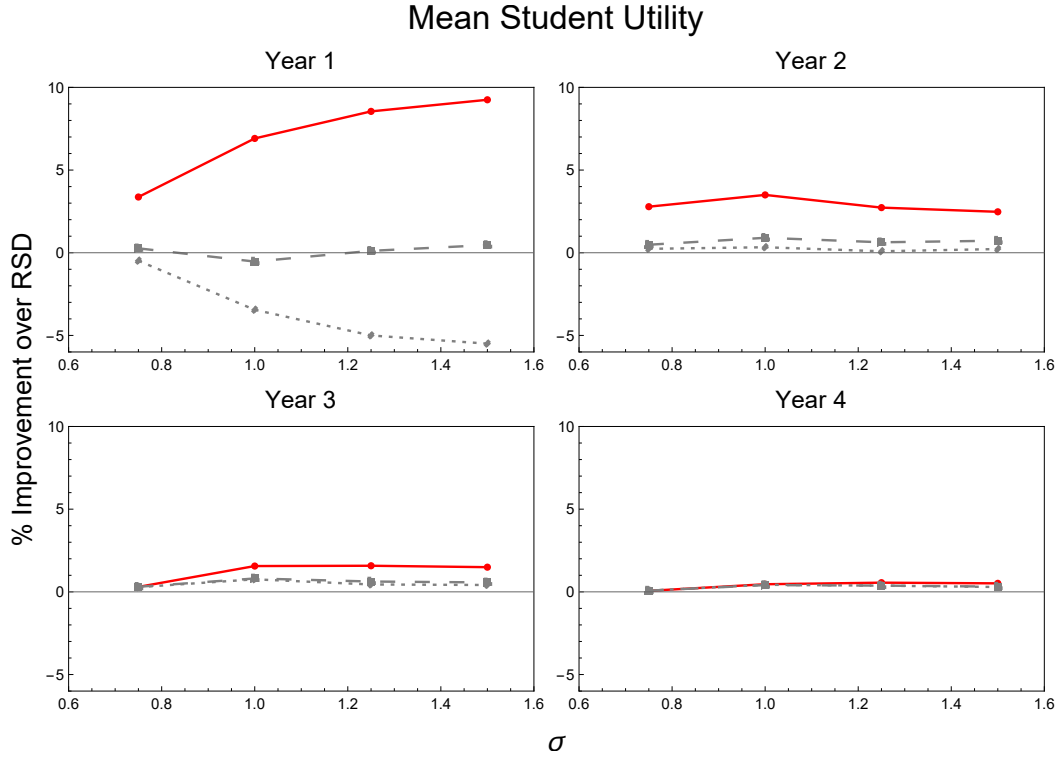


FIGURE D2: The performance of Pseudo-Market with Priorities (PMP), Deferred Acceptance with single tie-breaking (DA), and Deferred Acceptance with multiple tie-breaking DA(m) compared to the Random Serial Dictatorship (RSD) with optimal set-asides for $\sigma = 0.75, 1, 1.25, 1.5$. The results are based on 50 runs with different random utility draws for each parameter value.