

Undergraduate Course Allocation through Pseudo-Markets*

Daniel Kornbluth

Alexey Kushnir

April 20, 2022

Abstract

We consider a many-to-many matching problem with a priority structure such as the one in undergraduate course allocation. In order to incorporate course priorities, we develop a deterministic pseudo-market mechanism with priority-specific prices that is based on the approximate competitive equilibrium from equal incomes. This novel mechanism, the Pseudo-Market with Priorities mechanism, prevents justified course envy, prevents Pareto improvements among students respecting the priority structure, is strategy-proof in the large, bounds envy by a single course among students at the same level of priority, and maintains a small upper bound on the market-clearing error. In a simulated environment, we show that this mechanism increases student utility and outcome fairness when compared to the commonly used in practice Random Serial Dictatorship.

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

Keywords: market design, matching, many-to-many assignment, deterministic assignment, course allocation, competitive equilibrium

**Kornbluth:* Carnegie Mellon University, Pittsburgh, USA; dkornblu@andrew.cmu.edu; *Kushnir:* Carnegie Mellon University, Tepper School of Business, Pittsburgh, USA; akushnir@andrew.cmu.edu. We are grateful to Marek Pycia, Martin Bichler, and seminar participants at INFORMS Workshop on Market Design 2021 for their comments and suggestions.

1 Introduction

Every academic term, over 6500 post secondary institutions across America assign course schedules to a total of nearly twenty million students.¹ Based on timing and prerequisites, students have a limited number of schedules they can possibly be assigned in a term. Students also have quite heterogeneous preferences over courses. In the face of room size and teaching constraints, university registrars prioritize giving course seats to some students over others based on factors such as student seniority and major. As a result, university registrars face the challenging task of deciding how to allocate seats in overdemanded courses to undergraduate students.

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 *deterministic* allocation mechanism, the *Pseudo-Market with Priorities (PMP) mechanism*, that is based on the seminal Budish (2011)’s approximate competitive equilibrium from equal incomes (A-CEEI). This mechanism elicits student preferences over courses schedules, assigns almost equal budgets to students, and computes an approximate competitive equilibrium allocation and priority-specific prices. We require the equilibrium prices to respect the priority structure such that there is a "cutoff" priority level for each course, where students at the cutoff pay a non-negative price, students below the cutoff pay zero price, and students above the cutoff are not able to afford the course. We show that the PMP mechanism *prevents Pareto improvements respecting the priority structure* and *prevents justified course envy*, in that each student does not want to change her course assignment if she is given an opportunity to choose courses among courses in her assigned schedule and the assigned schedule of any student of lower priority. In addition, it *bounds envy by a single course among students at the same level of priority*, is *strategy-proof in the large*, and maintains the same upper bound for market-clearing error as in the case without priorities.

We highlight how university registrars would benefit by adopting the Pseudo-Market with Priorities mechanism for undergraduate course allocation. A typical university registration system uses a variant of the random serial dictatorship (RSD) mechanism, where students select courses in order of seniority and a priority structure based on student majors is enforced by setting aside seats in each course to be made available only to students in certain majors.² The PMP mechanism has two distinct advantages over the RSD mechanism. First, it endogenously computes course set-asides. In particular, the exact number of set-asides is typically hard to determine ex ante due to fluctuations in student demand for courses. On one hand, large set-asides lead to a large number of available seats demanded by students outside the major that the university registrar needs to assign man-

¹Based on National Center for Education Statistics, 2017-2018, <https://nces.ed.gov/>.

²Instead of a serial dictatorship, University of California system schools use an “enrollment pass” mechanism (see also Budish and Cantillon (2012)). The enrollment pass mechanism has the same problem in determining the exact number of set-asides.

ually with some ad hoc procedure after the main allocation process is over. On the other, small set-asides might lead to many students being rejected seats in courses that are required for their majors. The PMP mechanism relieves university registrars from having to correctly estimate the number of set-asides by computing them endogenously for each set of submitted student preferences and available seats in courses. This ensures that all students who require a course for their major are assigned these courses and limits remaining seats that need to be assigned manually. Second, the PMP mechanism allocates courses more fairly among students, as it bounds envy by a single course among students with at the same level of priority.

To provide additional support for our arguments, we compare the performance of the PMP and RSD mechanisms in a university-sized simulated environment. We conduct simulations with 5000 students and 1000 courses, where each course has 26 available seats and each student demands a schedule containing up to 10 courses, and investigate both major-specific course priority structures and common course priorities. We compare the PMP and RSD mechanism based on the mean and standard deviation of students' utility across 100 simulations. We also calculate the number of beneficial adjustments to the final assignments as explained below.

For major-specific priorities, students and courses are randomly and evenly distributed across fifty majors, and a binary priority structure in each course gives priority to students in a specific major. The PMP mechanism accounts for these priorities through priority-specific prices, while the RSD mechanism sets aside a number of seats in each course that can only be taken by students sharing the course's major. While the average student utility is only slightly larger for the PMP mechanism over the RSD mechanism, the standard deviation is significantly smaller. This suggests that the PMP mechanism can allocate courses more equally across students without sacrificing student utilities. We also account for two types of inefficiencies that arise in the RSD mechanism due to ex ante chosen set-asides. First, there may be unclaimed set-aside seats and students outside the course's major that want to receive these seats. Second, there may be students in a course's major who are not able to receive a seat in a course that a student outside the major is able to take. While the first situation is more prevalent for large set-asides, the second one is more prevalent for smaller set-asides. For an optimized choice of set-asides in the RSD mechanism, still nearly 18% of students on average would gain from an adjustment in the number of set-asides. In turn, the PMP mechanism endogenously finds set-asides with only a small number of possible beneficial adjustments. These beneficial adjustments may arise only because market clears approximately.

The paper is organized as follows. Section 2 introduces the environment of course allocation with priorities and presents some definitions. Section 3 develops the Pseudo-Market with Priorities mechanism and investigates its properties. Finally, Section 5 concludes. One should consult Kornbluth and Kushnir (2022) for *the review of the literature*. We avoid it here due to space limitation.

2 Environment

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

- $\mathcal{S} = \{1, \dots, S\}$ is a set of students; in reference to students, we use feminine pronouns.
- $\mathcal{C} = \{1, \dots, M\}$ is a set of courses.
- $Q = (q_1, \dots, q_M)$ is a vector of course capacities; each course 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 at most k courses in total. We assume that these restrictions are incorporated into student preferences, and, for simplicity of exposition, that $1 \leq k \leq M/2$. We also assume that student preferences over permissible course schedules are *strict*. Our results allow for the general case of substitutes and complements among courses.
- $\mathcal{R} = \{r_{s,c}\}_{s \in \mathcal{S}, c \in \mathcal{C}}$ is a course priority structure, where $r_{s,c} \in \{1, \dots, R\}$ is the priority of student s for course c , with a smaller number meaning a higher level of priority. The priority levels need not to be distinct and could be the same for several 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 $x_s \in 2^{\mathcal{C}}$ for each $s \in \mathcal{S}$. Allocation x is *feasible* if $\sum_{s \in \mathcal{S}} x_{s,c} \leq q_c$ for each course $c \in \mathcal{C}$. 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 $\{0, 1\}^M$.

We evaluate allocations based on fairness and efficiency. The seminal measure of fairness is envy-freeness, as introduced by [Foley \(1967\)](#): an allocation x *prevents envy* if there are no students $s, s' \in \mathcal{S}$ such that $x_{s'} \succ_s x_s$. The presence of priorities might allow some amount of envy among students that is justified.

Definition 1. An allocation x *prevents justified course envy* if there are no students $s, s' \in \mathcal{S}$ and course $c \in \mathcal{C}$ such that $r_{s,c} < r_{s',c}$, $c \notin x_s$, $c \in x'_{s'}$ and $c \in \arg \max_{\succsim_s} x_s \cup c$.

That is, if a student is willing to add a new course to her current schedule or replace one of her assigned courses with the new course, there should be no student of a lower level of priority for the new course who has a seat in the new course. The concept of preventing justified course envy directly corresponds to the absence of course-student blocking pairs (see [Roth and Sotomayor, 1990](#)).³ This concept also coincides with the concept of (pairwise) stability if no course seats remain unassigned.

³See [He, Miralles, Pycia, and Yan \(2018\)](#) and [Kesten and Ünver \(2015\)](#) for related envy concepts for random assignments.

In the case where courses have common priorities (i.e., each course has the same priorities over students), we also consider a weaker concept of an allocation *preventing justified schedule envy*, for which no student prefers the schedule of a lower priority student to her own schedule.

Additionally, we evaluate course allocations based on Pareto efficiency. 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 . As we explain below, we analyze approximate market equilibria, for which the market-clearing condition is satisfied with a small error. As a result, some unfilled course seats could be present. Hence, we distinguish between Pareto improvements among students and Pareto improvements between students and the university registrar that accounts for the possibility of unfilled seats. An allocation *prevents Pareto improvements among students* if there is no reallocation of the assigned seats in courses among students to obtain more preferred course schedules. An allocation *prevents Pareto improvements between students and the university registrar* if the university registrar cannot improve the course allocation by assigning students seats in courses with unfilled seats.

To account for the priority structure, we also introduce “course rankings” over allocations. Following [Schlegel and Mamageishvili \(2020\)](#), we define course rankings based on first-order stochastic dominance.⁴ We say that y dominates x for course c , and write $y_c \succeq_c x_c$, if for all $r = 1, \dots, R$ we have that $\sum_{s:r_{s,c} \leq r} y_{s,c} \geq \sum_{s:r_{s,c} \leq r} x_{s,c}$. Using this definition, we consider the following criterion of efficiency.

Definition 2. *An allocation x prevents Pareto improvements among students respecting the priority structure if for each assignment y that Pareto dominates x and has the same total number of assigned seats for each course, i.e., $\sum_{s \in S} y_{s,c} = \sum_{s \in S} x_{s,c}$ for all $c \in \mathcal{C}$, there is a course $c' \in \mathcal{C}$ for which $y_{c'} \not\succeq_{c'} x_{c'}$.*

The main difference between our concept of preventing Pareto improvements among students respecting the priority structure and priority-constrained efficiency, as introduced by [Schlegel and Mamageishvili \(2020\)](#), is that our concept requires any potential Pareto improvement to have the same total number of assigned seats in each course as currently assigned.⁵ This again accounts for the possibility of unfilled seats. Pareto improvements are possible between students and the university registrar when students can improve on their assigned schedule by adding courses with unfilled seats.⁶

Allocations are found through mechanisms, which systematically elicit student preferences over course schedules. Evidence from business schools demonstrates that mechanisms requiring strategic

⁴Note that course rankings coincide with course priorities in single-unit demand environments.

⁵We also consider only deterministic allocations.

⁶Note Definition 2 reduces to Budish’s approximate Pareto efficiency in environments without course priorities.

play on behalf of students can lead to large complications with efficiency (Budish and Cantillon 2012, Budish and Kessler 2021, and Sönmez and Ünver 2010).

Definition 3. A course allocation mechanism is **strategy-proof** if there is no student s , who, by reporting manipulated preferences $\succsim'_s$, receives an allocation she strictly prefers to the course schedule she would get by reporting $\succsim_s$.

We are mainly interested in mechanisms that are *strategy-proof in the large*; that is, strategy-proof in a limit market in which each student regards the “prices” in pseudo-market mechanisms, as introduced in the following section, as exogenous to her report (see Azevedo and Budish, 2019; Budish, 2011).⁷ We also assume that strategic play is only a concern for the students, as course priorities are typically set based on commonly observable factors such as student seniority or major.

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 (see Budish, 2011) to settings with course priorities. For this purpose, we allocate to each student s a budget of fake money b_s^* . To obtain good fairness properties, we assume that budgets are almost the same, with $1 \leq \min_s b_s^* \leq \max_s b_s^* \leq 1 + \beta$ for some small $\beta > 0$. The parameter β is interpreted as the maximum allowable budget inequality across students. We denote a positive number strictly greater than the largest allowed budget as $\bar{b} = 1 + \beta + \varepsilon$ for some $\varepsilon > 0$. We also allow for a slack in the market-clearing condition, which is bounded by $\alpha \geq 0$. To respect the priority structure, we require a special structure on priority-specific prices as defined below.

Definition 4. An allocation x^* , prices p^* , and budgets b^* constitute an (α, β) -Pseudo-Market Equilibrium with Priorities if

1. $x_s^* \in \max_{\succsim_s} \left\{ x' \in 2^{\mathcal{C}} : \sum_{c \in \mathcal{C}} p_{c,r_{s,c}}^* x'_c \leq b_s^* \right\}$ for each student $s \in \mathcal{S}$.
2. For each course $c \in \mathcal{C}$, there exists a cut-off 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

⁷See He, Miralles, Pycia, and Yan (2018) for asymptotic incentive compatibility in random assignment matching models.

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

4. $1 \leq \min_s b_s^* \leq \max_s b_s^* \leq 1 + \beta$.

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

Condition (1) ensures that the equilibrium allocation satisfies the priority structure, i.e., high priority students can obtain the course seats for free, the students at the cutoff level of priority face non-negative prices, and lower priority students cannot afford the course. A similar condition on prices first appeared in [He, Miralles, Pycia, and Yan \(2018\)](#) in the context of random allocation with priorities. In their setting, the pseudo-market equilibrium requires the market to clear exactly for courses with positive prices. In our setting, the market clears with some error. The market-clearing error for a course depends on its price for the lowest priority R , so that under-demand is not counted as an error if the price for the lowest priority group is zero. For all other cases, under-demand is counted towards the market-clearing error.⁸ Condition $\sum_{\{s \in \mathcal{S}: r_{s,c} < r_c^*\}} x_{s,c}^* < q_c$ also ensures that the cutoff is accurately defined, i.e., the cutoff cannot be lowered without causing under-demand for the course.

3.1 The Existence of Pseudo-Market Equilibrium with Priorities

Each course has up to R levels of priority for students. Hence, any price vector p will have MR components. This could potentially be problematic, as [Budish \(2011\)](#) establishes that the worst-case bound on the market-clearing error is proportional to the square root of the dimension of the price space. The main insight of the theorem presented below is that condition (1) allows us to reduce the effective dimension of the price space. In particular, all but one price for a given course are either zero or above the maximum allowable budget. Hence, there is only one relevant price component for each course. The theorem below shows there always exists a Pseudo-Market Equilibrium with Priorities with the same market-clearing error as in the case without priorities.⁹

⁸We want to mention that this requirement is a non-trivial extension of [Budish \(2011\)](#)'s condition to settings with course priorities. One consider an alternative definition where the market-clearing error for under-demanded courses with zero price for each level of priority is not counted as an error. In this case, one could use the cutoff level to artificially lower the market-clearing error. For example, consider the situation of a course with a price of zero at or below the cutoff level of priority and a price of $\bar{b}$ above 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 above the cutoff level was decreased.

⁹[Budish \(2011\)](#) shows the existence of an A-CEEI with market-clearing error at most $\sqrt{\sigma M}/2$, where $\sigma = \min\{2k, M\}$. Because we assume $k \leq M/2$, the worst-case bound on the market-clearing error is $\sqrt{kM}/2$.

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

Proof. See Kornbluth and Kushnir (2022). □

Overall, 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 priorities. The equilibrium has the same worst-case bound on the market-clearing error as the A-CEEI in the same environment without priorities.

3.2 Pseudo-Market with Priorities Mechanism

Using the Pseudo-Market Equilibrium with Priorities, we introduce the Pseudo-Market with Priorities (PMP) mechanism and investigate its properties in terms of fairness, efficiency, and the ability of students to manipulate its outcome.

Definition 5 (Pseudo-Market with Priorities Mechanism). *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 budget b_s^* that is a uniform random draw from $[1, 1 + \beta]$.*
3. *Compute prices $p^* = \{p_{c,r}^*\}_{c \in \mathcal{C}, r \in \mathcal{R}}$ and allocation $x^* = (x_1^*, \dots, x_S^*)$ such that*
 - *for each student $s \in \mathcal{S}$, allocation x_s^* maximizes her utility:*

$$x_s^* = \max_{\succsim_s} \left\{ x' \in 2^{\mathcal{C}} : \sum_{c \in \mathcal{C}} p_{c,r_{s,c}}^* x'_c \leq b_s^* \right\}$$

- *for each $c \in \mathcal{C}$, there exists a cutoff priority level r_c^* such that $\sum_{\{s \in \mathcal{S} : r_{s,c} < r_c^*\}} x_{s,c}^* < q_c$ and*

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

- *the market-clearing error is smaller than α , i.e., $\|z^*\|_2 \leq \alpha$, where $z^* = (z_c^*)_{c \in \mathcal{C}}$, and*

$$z_c^* = \begin{cases} \sum_{s \in \mathcal{S}} x_{s,c}^* - q_c & \text{if } p_{c,R}^* > 0 \\ \max[(\sum_{s \in \mathcal{S}} x_{s,c}^* - q_c), 0] & \text{if } p_{c,R}^* = 0 \end{cases}$$

Kornbluth and Kushnir (2022) establish that the PMP mechanism satisfied several important properties. For brevity of this submission we only state them without proofs. First, we show that the outcome of the PMP mechanism prevents justified course envy.

Theorem 2. *The outcome of the Pseudo-Market Equilibrium with Priorities mechanism prevents justified course envy.*

This property is closely related to (pairwise) stability in two-sided matching markets (see Echenique and Oviedo, 2006). The main difference is the possibility of blocking course-student pairs and individual courses due to the presence of under-subscribed and over-subscribed courses for the outcome of the PMP mechanism.

The PMP mechanism is not Pareto efficient. This is because it also needs to satisfy course priorities that might be in conflict with student preferences, similarly to how stable allocations in one-sided matching markets might not be Pareto efficient (see Roth and Sotomayor, 1990). However, the next result shows that the outcome of the PMP mechanism satisfies a form of constrained Pareto efficiency, i.e., students cannot trade courses among each other to improve their outcome without violating the priority structure.

Theorem 3. *The outcome of the Pseudo-Market Equilibrium with Priorities mechanism prevents Pareto improvements among students respecting the priority structure.*

Finally, there is almost no incentive for students to manipulate the PMP mechanism for large populations, which is typically the case for undergraduate course allocation.

Theorem 4. *The Pseudo-Market with Priorities mechanism is strategy-proof in the large.*

Next, we consider the performance of the PMP mechanism via simulations.

4 Simulations

We compare the performance of the Pseudo-Market with Priorities and Random Serial Dictatorship mechanisms through simulations of the course allocation process in a US college-sized environment. Our environment consists of 5000 students and 1000 courses, where each student demands a seat in five courses and each course has 26 seats.

We consider a setting with major-specific priorities.¹⁰ To create a priority structure, students and courses are randomly and evenly distributed across fifty majors. Each course gives priority to students that share its major, setting $r = 1$ for students with the same major (majors) and $r = 2$ for those who do not (non-majors).

¹⁰See Kornbluth and Kushnir (2022) for simulation in settings with major- and year-specific priorities.

Students are given a nonzero utility for ten out of the 1000 courses, including five courses within their major and five randomly selected others. This is a realistic assumption; even though a university may run hundreds of courses in a semester, any given student will only consider taking a small subset of the courses. Moreover, it allows our program to accommodate a large number of courses. For each student-course pair, this utility is made up of a non-random, random, and major-specific component. For course $c = 1, \dots, 1000$, if student s gives course c a nonzero utility, student s 's utility for taking course c is $c + \varepsilon_{s,c} + 250 \cdot [r_{s,c} = 1]$, where $\varepsilon_{s,c}$ is a random integer between -500 and 500 assigned to each student-course pair with positive utility, and $[r_{s,c} = 1]$ is an indicator variable which adds 250 utility if c is in s 's major. This utility design emphasizes that while some courses may be popular or unpopular among students, individual preferences vary. Preferences over course schedules are additively separable with respect to individual course utilities.

To find a Pseudo-Market Equilibrium with Priorities, we use a Tabu Search process similar to that of [Othman, Sandholm, and Budish \(2010\)](#) and [Budish, Cachon, Kessler, and Othman \(2017\)](#). Starting with a random vector of prices, each student's utility-maximizing schedule is found, and neighboring vectors are produced that adjust course prices based on over and under-demand. The neighboring vector with the largest reduction in market-clearing error is selected as a new starting point. The mechanism terminates after reaching the worst-case bound on the market-clearing error of $\alpha = \sqrt{kM/2} = 50$ and failing to improve significantly on the market-clearing error within five choices of a neighboring vector.

The RSD mechanism is conducted based on a randomly determined priority order among students. According to this order, students are assigned their most-preferred available course schedule. When a student selects a course within their major that has not filled all of its set-aside seats, the student will be assigned a set-aside seat. When no set-asides remain, students in the course's major take the remaining seats. Students outside the course major can only take seats outside set-asides. As a result, even if there are unfilled seats in a course, the course may not be available to a non-major student. In practice, this is used to aid the entry of students to courses within their majors, but can lead to inefficiency when too many or too few seats are reserved.

University registrars rely on information from previous years to estimate how many seats to reserve for majors in each course. To emulate this process, we randomly generate 25 environments, each consisting of student utilities, student and course majors, and a ranking across students. We utilize the deferred acceptance algorithm to determine an optimized number of set-asides for each course. This ensures that no non-major can receive a seat in a course when a major cannot, and no seat in a course is left empty when a student could benefit by taking it. Assigning each course set-asides equal to the number of seats it assigns to majors, it cannot be that a course has too many or too few set-asides; no non-major can receive a seat in a course when a major can, and no

Mechanism	Mean Student Utility	St. Dev. Student Utility	Distance to SOSM
RSD	3848	1202	876
DA with Random Priorities	4069	808	79
DA	4069	808	79
PMP	4069	808	79

Table 1: The mean and standard deviation of students’ utility, and the number of beneficial adjustments for the PMP and RSD mechanisms for major-specific priorities.

unclaimed set-aside seats prevent a non-major from taking the course. We take the average across the 25 groups (with each value rounded to the nearest integer) to be the set-asides we use in our simulations.

In turn, the PMP mechanism endogenously and optimally determines the number of set-asides. The PMP still allows for a possibility of a limited number of beneficial adjustments. When the market clears approximately, if there is under-demand when course prices are positive, there could exist unfilled seats that students at or above the cutoff level could benefit from taking. To provide a measure of beneficial adjustments in the PMP mechanism, we count the number of under-demanded seats when course prices are positive. Note that the under-demand will lead to beneficial adjustments only if there are students above the cutoff who would have demanded the course.

Table 1 presents results averaged across 100 simulations. As measures of student satisfaction and outcome fairness, we consider the mean and standard deviation of student total utility. We also count the number of possible beneficial adjustments in the RSD and PMP mechanisms as described above. The PMP mechanism is an improvement over the RSD in terms of students’ average utility and leads to a more fair outcome by reducing the standard deviation of students’ utilities. Additionally, the PMP mechanism leads to far less beneficial adjustments than the RSD mechanism. The 2734 beneficial adjustments in the RSD come from an average of 1450 different students, making up nearly 29% of the student body. In the PMP mechanism, under-demand is an upper bound on the true number of possible beneficial adjustments in the PMP mechanism; under-demand will only lead to beneficial adjustments if there are students above the cutoff would have demanded the course.

This suggests that reserving seats for students can lead to inefficiency even if the university registrar bases its set-asides on information from previous years. Because the PMP mechanism endogenously determines its cutoff groups, it can approximately determine the number of set-asides

necessary. In doing this, the PMP mechanism allows for a university to prioritize the entry of students into courses with less of a sacrifice to efficiency as the RSD.

5 Conclusion

In this paper, we explore a many-to-many matching problem that typically arises in undergraduate course allocation. We allow courses to prioritize students based on factors such as seniority and major, and design a deterministic allocation mechanism, the Pseudo-Market with Priorities mechanism, that respects this priority structure. The PMP mechanism is an extension of the A-CEEI mechanism to settings with course priorities and maintains its upper bound on the outcome's market-clearing error. We highlight how university registrars could benefit from adopting the PMP mechanism through computer simulations. First, the PMP mechanism respects course priorities and computes course set-asides endogenously, which relieves the burden of estimating them correctly from university registrars. In doing this, the PMP mechanism also ensures that students who require courses for their majors are assigned these courses and leaves almost no seats that need to be assigned manually. Second, the PMP mechanism leads to a more equal distribution of courses among students at the same level of priority than the serial dictatorship mechanism.

References

- ASHLAGI, I., AND P. SHI (2016): “Optimal allocation without money: An engineering approach,” *Management Science*, 62(4), 1078–1097.
- AZEVEDO, E., AND E. BUDISH (2019): “Strategy-proofness in the Large,” *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*.
- BUDISH, E. (2011): “The Combinatorial Assignment Problem: Approximate Competitive Equilibrium from Equal Incomes,” *Journal of Political Economy*, 119, 1061–1061.
- BUDISH, E., G. P. CACHON, J. B. 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., AND J. B. KESSLER (2021): “Can Market Participants Report Their Preferences Accurately (Enough)?,” *Management Science*.
- DIEBOLD, F., H. AZIZ, M. BICHLER, F. MATTHES, AND A. SCHNEIDER (2014): “Course allocation via stable matching,” *Business & Information Systems Engineering*, 6(2), 97–110.
- DIEBOLD, F., AND M. BICHLER (2017): “Matching with indifference: A comparison of algorithms in the context of course allocation,” *European Journal of Operational Research*, 260(1), 268–282.
- ECHENIQUE, F., A. MIRALLES, AND J. ZHANG (forthcoming): “Constrained pseudo-market equilibrium,” *American Economic Review*.
- ECHENIQUE, F., AND J. OVIEDO (2006): “A theory of stability in many-to-many matching markets,” *Theoretical Economics*, 1(2), 233–273.
- EHLERS, L., AND B. KLAUS (2003): “Coalitional strategy-proof and resource-monotonic solutions for multiple assignment problems,” *Social Choice and Welfare*, 21, 265–280.
- FOLEY, D. K. (1967): “Resource allocation and the public sector,” *Yale Economics Essays*.
- 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.
- KESTEN, O., AND M. U. ÜNVER (2015): “A theory of school-choice lotteries,” *Theoretical Economics*, 10(2), 543–595.
- KORNBLUTH, D., AND A. KUSHNIR (2022): “Undergraduate Course Allocation through Pseudo-Markets,” .
- MIRALLES, A., AND M. PYCIA (2021a): “Foundations of pseudomarkets: Walrasian equilibria for discrete resources,” Discussion Paper 385, University of Zurich, Department of Economics, Working Paper.
- (2021b): “Large vs. continuum assignment economies,” Discussion paper, University of Zurich.
- NGUYEN, T., A. PEIVANDI, AND R. VOHRA (2016): “Assignment problems with complementarities,” *Journal of Economic Theory*, 165, 209–241.
- OTHMAN, A., T. SANDHOLM, AND E. BUDISH (2010): “Finding approximate competitive equilibria: efficient and fair course allocation,” in *AAMAS*, vol. 10, pp. 873–880.
- PAPAI, S. (2001): “Strategyproof and Nonbossy Multiple Assignments,” *Journal of Public Economic Theory*, 3, 257–71.
- PYCIA, M. (2021): “Pseudomarkets,” in *Online and Matching-Based Market Design*, ed. by F. Echenique. Cambridge University Press.
- ROTH, A. E., AND M. A. O. SOTOMAYOR (1990): *Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis*, Econometric Society Monographs. Cambridge University Press.
- SCHLEGEL, J. C., AND A. MAMAGEISHVILI (2020): “Welfare theorems for random assignments with priorities,” *Games and Economic Behavior*, 124, 62–81.
- SÖNMEZ, T. (2014): “Bidding for Army Career Specialties: Improving the ROTC Branching Mechanism,” *Journal of Political Economy*, 121(1), 186–219.
- SÖNMEZ, T., AND T. B. SWITZER (2013): “Matching With (Branch-of-Choice) Contracts at the United States Military Academy,” *Econometrica*, 81(2), 451–488.
- SÖNMEZ, T., AND M. U. ÜNVER (2010): “Course bidding at business schools,” *International Economic Review*, 51(1), 99–123.