Course Allocation via Stable Matching

Authors

Franz Diebold*, Haris Aziz+, Martin Bichler*, Florian Matthes*, Alexander Schneider*

*Department of Informatics (I18), Technische Universität München, Germany

franz.diebold@tum.de, bichler@in.tum.de, matthes@in.tum.de, alexander.schneider@tum.de

+NICTA, 223 Anzac Parade, Sydney, 2033, Australia

Haris.aziz@nicta.com.au

The allocation of students to courses is a wide-spread and repeated task in higher education, often accomplished by a simple first-come first-served (FCFS) procedure. FCFS is neither stable nor strategy-proof, however. The Nobel Prize in Economic Sciences was awarded to Al Roth and Lloyd Shapley for their work on the theory of stable allocations. This theory was influential in many areas, but found surprisingly little application in course allocation as of yet. In this paper, we survey different approaches for course allocation with a focus on appropriate stable matching mechanisms. We will discuss two such mechanisms in more detail, the Gale-Shapley student optimal stable mechanism (SOSM) and the efficiency adjusted deferred acceptance mechanism (EADAM). EADAM can be seen as a fundamental recent contribution which recovers efficiency losses from SOSM at the expense of strategy-proofness. In addition to these two important mechanisms, we provide a survey of recent extensions with respect to the assignment of schedules of courses rather than individual courses. We complement the survey of the theoretical literature with results of a field experiment, which help understand the benefits of stable matching mechanisms in course allocation applications.

Keywords: Matching, stability, efficiency

1 Introduction

The Nobel Memorial Price in Economic Science in 2012 was awarded to Alvin E. Roth and Lloyd S. Shapley for "the theory of stable allocations and the practice of market design" as the Royal Swedish Academy of Sciences put it. Market design is an academic field at the intersection of computer science, economics, and the management sciences concerned with the design of market institutions (Roth 2002, pp. 1341-1378; Milgrom 2011, pp. 311-320), with many recent contributions from the Information Systems community (for example Weinhardt et al. (2003, pp. 635-640), Bapna et al. (2004, pp. 21-43), Bichler et al. (2009, pp. 111-117) and Bichler et al. (2011, pp. 688-699)).

Roth and Shapley were recognized for their work on stable matching between two sets of elements given preferences of each element over the other set. Those sets could be students and courses for example, and in contrast to auction markets no monetary transfers are allowed. A matching is stable, if there does not exist any alternative pairing in which both students and course organizers are better off. The problem of computing a stable matching is different from that of computing a maximum weight bipartite matching in the assignment problem, which is regularly taught in computer and management science. Gale and Shapley (1962, pp. 9-15) presented one of the most commonly used two-sided matching models, the marriage model, which matches a single man to a single woman (one-to-one). They suggested the *Gale-Shapley deferred acceptance algorithm* to find a stable matching.

Gale and Shapley (1962, pp. 9-15) show that, when preferences are *strict*, the deferred acceptance algorithm yields the unique stable matching in $O(n^2)$ time that is Pareto superior to any other stable matching from the viewpoint of the students. Therefore, the outcome of the student proposing deferred acceptance algorithm is also called the student optimal stable matching and the mechanism that associates the student optimal stable matching to any *one-to-many matching problem* as it can be found in school choice or college admission is known as the *student optimal stable mechanism* (SOSM). The underlying concept is the same as in the one-to-one Gale-Shapley deferred acceptance algorithm. Besides the fact that it gives the most efficient stable matching, another appealing feature of the SOSM is that it is *strategy-proof* (Roth 1982, pp. 617-628). Strategy-proofness means that no student has an incentive to misreport his true preferences and students have dominant strategies, clearly a very desirable property.

Although SOSM is strategy-proof and stable, the matching is not necessarily Pareto efficient, which is another important design desideratum. Only recently, Kesten (2010, pp. 1297-1348) proposed the *adjusted deferred acceptance mechanism* (EADAM) that allows for stability and Pareto efficiency at the expense of strategy-proofness. However, truth telling of students is a Bayesian Nash equilibrium in this mechanism. Even if there are no dominant strategies, possibilities to strategically misrepresent preferences are minimal in most applications. EADAM can be seen as a significant contribution to the literature. Apart from EADAM there have been a number of recent contributions with respect to the assignment of bundles of courses to a single student, and the theory of matching has drawn quite a bit of recent academic attention with many open problems which can be considered fundamental for Economics and the Management Sciences in general.

In this paper, we will focus on the one-to-many matching problem in the context of *course allocation*, as it is a wide-spread problem in higher education and beyond. We will discuss prime candidates for the solution of this problem and complement this survey with results from a field study, in which we compare two stable matching mechanisms to a *first-come-first-served* (FCFS) mechanism, a standard mechanism in many institutions. Given the wide-spread application of FCFS, it is interesting to understand the impact of stable matching mechanisms over FCFS in terms of stability and efficiency. Our field studies are the first to analyze EADAM in a real-world application to our knowledge, and we are not aware of empirical work comparing stable matching mechanisms such as SOSM or EADAM to FCFS. Apart from strategic properties of mechanisms such as strategy-proofness and stability, these field experiments shed light on secondary desiderata such as the average rank of students or the rank distribution that result from stable matching mechanisms and the incumbent FCFS. Note that we do not consider time tabling problems or capacity management with respect to rooms available in a school, because such decisions are typically made before students register for one or another course.

In Section 2 we will introduce stable matching and discuss various applications to provide an introduction to the field. In Section 3, we formally introduce the course allocation problem and provide relevant design desiderata as well as a succinct description of SOSM, EADAM, and FCFS. Section 4 introduces data and the results of two field experiments. This provides an understanding of how such mechanisms can be evaluated in practice and how the outcomes of FCFS differ from those of stable matching mechanisms. In Section 6 we discuss recent literature on the allocation of course schedules, before we conclude the paper in Section 7.

2 Stable matching and its applications

We will first give an overview of typical problems and successful applications before we focus on course allocation in more detail. In particular, the school choice problem has drawn a lot of attention and it shares many similarities with course allocation. In order to give parents the opportunity to choose the public school their child will attend, many U.S. states provide school choice systems. Each prospective student submits a list of preferences of schools to the central placement authority of the school district. On the other hand each school has a priority ordering of all students and a maximum capacity. This information is used to determine which student will be assigned to which school. Abdulkadiroğlu and Sönmez (2003, pp. 729-747) showed that matching mechanisms that have been in use in the US did not perform well in terms of efficiency, incentives, and stability. As a consequence, the Boston Public Schools replaced a priority mechanism (the Boston mechanism) with a deferred acceptance mechanism in 2005. Abdulkadiroglu, Pathak et al. (2006) present further arguments against the Boston mechanism. This mechanism does not exhibit a dominant strategy equilibrium and it is not stable. Some experimental work has focused on school choice comparing these two approaches. For example, in Chen and Sönmez (2006, pp. 202-231) experimental subjects play a one-shot game of incomplete information in which each participant is only informed about his own preferences, schools' capacities, and the matching mechanism. They find that from the perspective of students, the SOSM outperforms both the Boston mechanism and other alternatives. Featherstone and Niederle (2008) confirm this, but also discuss settings with only private information of subjects, where the Boston mechanism has advantages.

In Germany, the assignment of students to universities via the Zentralstelle für die Vergabe von Studienplätzen (ZVS) is a large-scale application of matching. One part of the capacity is reserved for excellent students and students with long waiting times via a Boston mechanism. The remaining places are allocated on the basis of universities‘ preferences via a university-proposing Gale-Shapley deferred acceptance mechanism. The two parts are administered sequentially in the aforementioned order (Westkamp 2012, pp. 561-589). Braun, Dwenger et al. (2007) present evidence from the field that some applicants behave strategically and not truthfully in this combined mechanism. Another widely cited application of one-to-many stable matching is the assignment of graduating medical students to their first hospital appointments by the National Resident Match Program in the US (Roth 1984, pp. 991-1016). A number of other applications can be found online (http://www.matching-in-practice.eu/).

The Gale-Shapley deferred acceptance mechanism has dominated the stable matching literature for many years and most comparisons are with the Boston mechanism, probably driven by the school choice problem. The outcomes of both mechanisms are not necessarily Pareto efficient, however. Unfortunately, Kesten (2010, pp. 1297-1348) shows that there cannot exist a mechanism always returns a Pareto optimal and stable matching.

There are Pareto efficient and strategy-proof (but not stable) mechanisms for school choice proposed by Abdulkadiroğlu and Sönmez (2003, pp. 729-747), which are based on the *top trading cycle* algorithm (Shapley and Scarf 1974, pp. 23-37). This mechanism enables a student to obtain a higher priority at a desirable school by trading his priority for a less desirable school with another student. In the context of school choice problems, the idea of trading priorities has been perceived as a problem. Alternatively, a simple *random serial dictatorship* can be used, where students are randomly ordered and then assigned to their first preference among the remaining schools. Also this mechanism is efficient and strategy-proof, but not stable. Similar to the top trading cycle, it has not been used for school choice or college admission as far as we know. Roth (2002, pp. 1341-1378) argues based on empirical observations that stability is a key feature of successful matching mechanisms in practice.

Gale and Shapley (1962, pp. 9-15) showed that the deferred acceptance algorithm leads to stable matchings if preferences are strict and complete. In practice, preference lists might be incomplete and the schools that are not on the preference list of a student might just not be acceptable to a student. Such problems can be accommodated by a version of SOSM (Manlove et al. 2002, pp. 261-279). It might also be that there are ties in the preferences or that unlisted schools have the lowest preference for students and they are tied but acceptable. A common practice in case of ties is to randomly break the ties. Erdil and Ergin (2008, pp. 669-689) show that there can be efficiency losses by SOSM when ties in priorities are broken in some random way. They introduce one mechanism that restores such artificial welfare losses. EADAM is an alternative way to recover welfare losses originating from random tie-breaking, which is another appealing feature of the mechanism and a reason, why we discuss it in this paper.

Note that there are also related versions of the stable matching problem that are computationally hard. For example, if there are ties, the preferences are incomplete, and the schools which are not listed are unacceptable, then finding a maximum cardinality stable matching is an *NP*-complete problem (Manlove et al. 2002, pp. 261-279). Related work has looked into approximation algorithms for hard stable matching problems (Halldórsson et al. 2003, pp. 431-447). We will not discuss questions of computational complexity further in this paper. In our field studies below we assume the preferences of students to be strict and courses not ranked by the students are considered unacceptable. We also refer the interested reader to Manlove (2013) for a comprehensive study of algorithmic aspects of matching problems.

Although the course allocation problem is similar to school choice it has not received as much attention in the literature, and different mechanisms are being used in practice. We are not aware of applications of the Boston mechanism, random serial dictatorship, or the top-trading cycle for course allocation, for example. Typically, monetary transfers are not allowed for course assignment in higher education so that auctions are not an option. However, some business schools in the USA use course bidding, where students are given a bid endowment in a virtual currency to allocate across the courses they consider taking. This virtual currency does not have outside value and there are various possibilities for manipulation (see Sönmez and Ünver (2010, pp. 99-123) and Krishna and Ünver (2008, pp. 262-282)). Krishna and Unver (2005) report on a field experiment with 535 students comparing course bidding with the Gale-Shapley stable matching mechanism and find that the latter could vastly improve efficiency.

3 Course allocation problem

Similar to school choice one can argue that stability is a desirable feature of course allocation mechanisms (Roth 2002, pp. 1341-1378). We want to discuss the relative merits of stable matching mechanisms for course allocation in this paper and compare them to FCFS. For this we will define the course allocation problem more formally in the next section and discuss prime candidates for stable course allocation mechanisms, as well as the FCFS procedure.

3.1 Matchings, matching mechanisms, and their properties

A *course allocation problem* consists of a finite set of students $S \equiv \{s_1, s_2, \dots, s_n\}$ and a finite set of courses $C \equiv \{c_1, c_2, \dots, c_m\}$ with the *maximum capacities* $\mathrm{q} = \left(\mathrm{q}_{\mathrm{c}_1}, \mathrm{q}_{\mathrm{c}_2}, \dots, \mathrm{q}_{\mathrm{c}_\mathrm{m}}\right)$. To ensure that a feasible matching exists we assume $q_c \geq 0$ for all $c \in C$ and $n \leq \sum_{c \in C} q_c$. Each student has a *preference relation* $\succsim_s$ over the courses C (called student preferences), each course (organizer) a preference relation $\succsim_\mathrm{c}$ over the students S (called course organizer preferences). These are essentially priority orderings of the course organizer over the students, which means that the course organizers will not be considered as strategic. We will assume strict preferences although SOSM and EADA extend easily for the case of indifferences.

The vectors for these relations are denoted $\succsim_S = (\succsim_s)_{s\in S}$ and $\succsim_C = (\succsim_c)_{c\in C}$. Let $\mathcal{P}$ denote the set of all possible preference relations over C and $\mathcal{P}^{|S|}$ the set of all preference vectors for all students.

Definition 1 (Matching). *A* matching *is a mapping* μ *of students* S *to courses* C *that satisfies:*

(i) $\mu(s) \in C$ *for all* $s \in S$
(ii) $\mu^{-1}(c) \subseteq S$ *for all* $c \in C$*, and*
(iii) for any $s \in S$ *and* $c \in C$*, we have* $\mu(s) = c$ *if and only if* $s \in \mu^{-1}(c)$

A matching is *feasible* if $|\mu^{-1}(c)| \leq q_c$ for all $c \in C$, which means that no course is overcrowded. The following simple example adapted from Kesten (2010, pp. 1297-1348) should illustrate these definitions.

Example 1. Consider the course allocation problem with four students $S \equiv \{s_1, s_2, s_3, s_4\}$ and four courses $C \equiv \{c_1, c_2, c_3, c_4\}$, each course having one seat. The course organizer preferences ($\succ_c$) and the student preferences ($\succ_s$) are given below:

Tab. 1 Course organizer preferences ($\succ_c$) and student preferences ($\succ_s$) for Example 1. Underlined preferences describe a matching.

$\succ_{c_1}$	$\succ_{c_2}$	$\succ_{c_3}$	$\succ_{c_4}$
s_4	s_2	s_3	$\underline{s_1}$
s_1	$\underline{s_3}$	$\underline{s_4}$	⋮
$\underline{s_2}$	⋮	⋮	⋮

$\succ_{s_1}$	$\succ_{s_2}$	$\succ_{s_3}$	$\succ_{s_4}$
c_1	$\underline{c_1}$	$\underline{c_2}$	$\underline{c_3}$
$\underline{c_4}$	c_2	c_3	c_1
⋮	⋮	⋮	⋮

One desirable property of matchings is *Pareto efficiency* such that no student can be made better off without making any other student worse off.

Definition 2 (Pareto efficiency of matchings). *A matching* μ *is* Pareto efficient *with respect to the students if there is no other feasible matching* μ' *such that* $\mu'(s) \succcurlyeq_s \mu(s)$ *for all students* $s \in S$ *and* $\mu'(s) \succ_s \mu(s)$ *for some* $s \in S$*.*

Stability means, that there should be no unmatched pair of a student and a course (s, c) where student s prefers course c to her current assignment and she has higher priority than some other student who is assigned to course c. Stability can be seen as capturing *no justified envy*.

Definition 3 (Stability). *A matching* μ *is* stable *if* $\mu(s') \succ_s \mu(s)$ *implies* $s' \succ_{\mu(s')} s$ *for all* $s, s' \in S$.

Next we will discuss mechanisms to compute matchings and their properties. A mechanism returns a matching for given preferences of students and courses. More formally, let $\mathcal{M}$ denote the set of all feasible matchings. A *matching mechanism* χ can be described as a function $\chi : \mathcal{P}^{|S|} \longrightarrow \mathcal{M}$ that returns a feasible matching of students to courses for every preference profile of the students. For a submitted preference profile $\succsim_S \in \mathcal{P}^{|S|}$ of the students, $\chi(\succsim_S)$ is the associated matching. For a student s the assigned course is $\chi_s(\succsim_S) \in C$. For a course c the set of the assigned students is $\chi_c(\succsim_S) \subseteq S$.

A mechanism is *Pareto efficient* if it always selects a Pareto-efficient matching. Also, a mechanism is *stable* if it always selects a stable matching. Another important property of a mechanism is *strategy-proofness*. This means, that there is no incentive for any student not to submit her truthful preferences, no matter what the other students submit.

Definition 4 (Strategy-proofness). *A mechanism* χ *is* strategy-proof *if for any* $\succsim_S \in \mathcal{P}^{|S|}$ *with* $s \in S$ *and* $\succsim_s' \in \mathcal{P}$ *we have* $\chi_s(\succsim_S) \succcurlyeq_s \chi_s(\succsim'_s, \succsim_{S\backslash\{s\}})$.

$\chi_s(\succsim'_s, \succsim_{S\backslash\{s\}})$ describes the preference profile, where the preferences of student s, $\succsim'_s$, differ from his true preferences $\succsim_s$. Unfortunately, there exists no strategy-proof matching mechanisms that is both efficient and stable (Kesten 2010, pp. 1297-1348).

For an illustration of the above definitions, we will introduce the *random serial dictatorship* mechanism, which is arguably one of the simplest matching mechanisms.

Algorithm 1.

```
Draw uniformly at random a permutation π of {1,2,…,n}.  of the
students in S . For x from 1 to n = |S| assign student c_π(x) to her
top choice among the remaining slots.
```

Consider the permutation

$$\pi = \begin{pmatrix} 1 & 2 & 3 & 4 \\ 2 & 1 & 3 & 4 \end{pmatrix}$$

which could be the result of a lottery. This would lead to the following ordering of the students $(s_2,\ s_1,\ s_3,\ s_4)$. Then the mechanism would work as described in Table 2.

Tab. 2 Example of the serial dictatorship mechanism

Step	Student to be assigned	c_1	c_2	c_3	c_4
1	s_2	$\boxed{s_2}$			
2	s_1	$\vdots$			$\boxed{s_1}$
3	s_3	$\vdots$	$\boxed{s_3}$		$\vdots$
4	s_4	$\boxed{s_2}$	$\boxed{s_3}$	$\boxed{s_4}$	$\boxed{s_1}$

The matching

$$\mu_1 = \begin{pmatrix} c_1 & c_2 & c_3 & c_4 \\ s_2 & s_3 & s_4 & s_1 \end{pmatrix}$$

would be the result. This result is underlined in Table 1.

The random serial dictatorship mechanism is *Pareto efficient* and *strategy-proof,* which is straightforward to see. Unfortunately, the mechanism is *not stable*, as the following shows. Student s_1 and course c_1 are unmatched, where student s_1 prefers course c_1 to her current assignment (c_4) and she has higher priority to c_1 than student s_2 who is assigned to course c_1. By setting $s = s_1$ and $s' = s_2$ we get

$$c_1 = \mu(s_2) = \mu(s') \succ_{s_1} \mu(s) = \mu(s_1) = c_4 \text{ and } s_2 = s' \prec_{c_1} s = s_1,$$

contradicting the definition of stability. Another disadvantage of random serial dictatorship is that *course-specific priorities* are not possible because there is only one single ordering of the students for all courses.

In the following subsections, we describe stable matching mechanisms to solve the course allocation problem. SOSM is well known, while the EADAM is a more recent approach, which gives up on strategy-proofness for a weaker game-theoretical solution concept, that of a Bayes-Nash equilibrium, to gain efficiency. Then we will briefly introduce FCFS as it used in many universities.

3.2 Gale-Shapley student-optimal stable mechanism (SOSM)

The *Gale-Shapley student-optimal stable mechanism (SOSM)* is a modified version of the *Gale-Shapley deferred acceptance algorithm* from (Gale and Shapley 1962, pp. 9-15), which allows for one-to-many assignments. This algorithm works as follows:

Algorithm 2.

Step 1: Each student proposes to her first choice course. For each course c with a capacity q_c, those q_c proposers who have the highest priority for c are tentatively assigned to c, the remaining proposers are rejected.

In general, at

Step k, $k \geq 2$: Each student who was rejected in the previous step $(k-1)$ proposes to her next choice course. For each course c, from the new proposers and those who were tentatively assigned at a previous step, the q_c with the highest priority are tentatively assigned to c, the rest is rejected.

The algorithm terminates when no student is rejected any more.

Example 2. Consider the problem given in Example 1:

Tab. 3 Course organizer preferences ($\succ_c$) and student preferences ($\succ_s$) for Example 2

$\succ_{c_1}$	$\succ_{c_2}$	$\succ_{c_3}$	$\succ_{c_4}$
$\cancel{s_4}$	$\cancel{s_2}$	$\cancel{s_3}$	$\cancel{s_1}$
s_1	s_3	s_4	$\vdots$
s_2	$\vdots$	$\vdots$	$\vdots$

$\succ_{s_1}$	$\succ_{s_2}$	$\succ_{s_3}$	$\succ_{s_4}$
c_1	c_1	c_2	c_3
$\cancel{c_4}$	$\cancel{c_2}$	$\cancel{c_3}$	$\cancel{c_1}$
$\vdots$	$\vdots$	$\vdots$	$\vdots$

The steps of the algorithm applied to this problem are shown in the following table. The students that are tentatively assigned to a course are shown in a box, the rejected students are not shown in a box.

Tab. 4 Example of the Gale-Shapley student-optimal stable mechanism

Step	c_1	c_2	c_3	c_4
1	$\boxed{s_1}, s_2$	$\boxed{s_3}$	$\boxed{s_4}$	
2	$\vdots$	$\boxed{s_2}, s_3$	$\vdots$	
3		$\vdots$	$\boxed{s_3}, s_4$	
4	$\boxed{s_4}, s_1$		$\vdots$	
5	$\boxed{s_4}$	$\boxed{s_2}$	$\boxed{s_3}$	$\boxed{s_1}$

The resulting matching is

$$\mu_2 = \begin{pmatrix} c_1 & c_2 & c_3 & c_4 \\ s_4 & s_2 & s_3 & s_1 \end{pmatrix}$$

and is underlined in Table 3.

As the name suggests, SOSM is *stable* (see Gale and Shapley (1962, pp. 9-15) for the proof) and *strategy-proof (Abdulkadiroğlu and Sönmez 2003, pp. 729-747)* (Proposition 2)). It is used in the public school systems of Boston and New York City, and both properties were used as an argument to switch to this mechanism.

One can easily verify, however, that the matching μ_2 is not *Pareto efficient*. If the students s_2, s_3 and s_4 would be assigned to their first choice course and s_1 would not be changed, then three students would be better off. This means, that SOSM may produce welfare losses on the students side and these losses can be significant. The *top trading cycle* algorithm is an alternative algorithm, which is strategy-proof and efficient, but not stable (Abdulkadiroğlu and Sönmez 2003), but in school choice applications stability has typically been preferred over efficiency.

The Boston mechanism is another algorithm that has been used for school choice. It works as follows. Let us assume that there are m schools. For round k=1 to m it does the following. In each round k, those students who have not been allocated a seat yet are considered to be allocated a seat at their k-th most preferred school. The seats are allocated according to the priorities of the schools as long as capacity is not used up. By the end of round m, each student has been allocated a school seat. The Boston mechanism is manipulable and truthful revelation is not a dominant strategy (Abdulkadiroğlu and Sönmez 2003, pp.729-747). Multiple sources discuss strategies that yield better outcomes than truth-telling such as not ranking unachievable courses (Abdulkadiroğlu et al. 2006). Also the EADAM mechanism introduced in the next subsection does not provide strategy-proofness, but truth-telling is a Bayes-Nash equilibrium strategy.

3.3 Efficiency adjusted deferred acceptance mechanism (EADAM)

Kesten (2010, pp. 1297-1348) introduced a matching mechanism, which reduces welfare losses on the student side as described above, but gives up on strategy-proofness. If we again look at Example 2, we can see that there is a *rejection chain* from step 1 to 4, which is initiated by student s_1 at course c_1 who rejects s_2 in step 1.

This rejection induces the rejections at the steps 2, 3, and 4, where s_1 is rejected from course c_1. Student s_1 does not benefit from being tentatively assigned to course c_1 from step 1 to step 4 and only hurts the other students (s_2, s_3 and s_4). Hence, Kesten calls any student like s_1 an *interrupter* and a pair like (s_1, c_1) an *interrupting pair*. If student s_1 would waive his priority for the critical course c_1 the other students would be assigned to their first choice course.

Definition 5 (Interrupter) (Kesten 2010). *Given a problem to which the DA algorithm is applied, let* s *be a student who is tentatively placed to a course* c *at some Step* t *and rejected from it at some later Step* t'*. If there is at least one other student who is rejected from course* c *after Step* $\mathrm{t}-1$ *and before Step* t'*, i.e., rejected at a step* $l \in \{t, t+1, \dots, t'-1\}$*, then we call student* s *an* interrupter *for course* c*, and the pair* (s, c) *an* interrupting pair *of Step* t'*.*

The following algorithm describes the above mechanism.

Algorithm 3 (Kesten 2010).

Step 0: Run the DA algorithm.

Step 1: Find the last step (of the DA algorithm run in Step 0) at which a consenting interrupter is rejected from the course for which he is an interrupter. Identify all interrupting pairs of that step each of which contains a consenting interrupter. If there are no interrupting pairs, then stop. For each identified interrupting pair (s, c), remove course c from the preferences of student s without changing the relative order of the remaining courses. Re-run the DA algorithm with the new preference profile.

In general,

Step $t, t \geq 2$: Find the last step (of the DA algorithm run in Step $t-1$) at which a consenting interrupter is rejected from the course for which he is an interrupter. Identify all interrupting pairs of that step each of which contains a consenting interrupter. If there are no interrupting pairs, then stop. For each identified interrupting pair (s, c) remove course c from the preferences of student s without changing the relative order of the remaining courses. Re-run the DA algorithm with the new preference profile.

Example 3. Again, we look at the problem given in Example 1, assuming for simplicity that all students consent.

Step 0*:* See Example 2.

Step 1*:* Since student s_1 is rejected from course c_1 at Step 4 and since student s_2 has been rejected from course c_1 while student s_1 was tentatively assigned to course c_1, we identify (s_1, c_1) as the last and the only interrupting pair. Suppose student s_1 consents. Then we remove course c_1 from student s_1's preferences

Tab. 5 Course organizer preferences ($\succ_c$) and updated student preferences ($\succ_s$) for Example 3

$\succ_{c_1}$	$\succ_{c_2}$	$\succ_{c_3}$	$\succ_{c_4}$
$\underline{s_4}$	$\underline{s_2}$	$\underline{s_3}$	$\underline{s_1}$
s_1	s_3	s_4	$\vdots$
s_2	$\vdots$	$\vdots$	$\vdots$

$\succ_{s_1}$	$\succ_{s_2}$	$\succ_{s_3}$	$\succ_{s_4}$
$\cancel{c_1}$	c_1	c_2	c_3
$\underline{c_4}$	$\underline{c_2}$	$\underline{c_3}$	$\underline{c_1}$
$\vdots$	$\vdots$	$\vdots$	$\vdots$

and re-run the DA algorithm with the new preference profile shown in Table 5.

Tab. 6 Example of the efficiency adjusted deferred acceptance mechanism

Step	c_1	c_2	c_3	c_4
1	$\boxed{s_2}$	$\boxed{s_3}$	$\boxed{s_4}$	$\boxed{s_1}$

The resulting matching is

$$\mu_3 = \begin{pmatrix} c_1 & c_2 & c_3 & c_4 \\ s_2 & s_3 & s_4 & s_1 \end{pmatrix}.$$

The outcomes of SOSM (underlined) and EADAM (in boxes) are shown in Table 5 and 6 resp. Kesten (2010, pp. 1297-1348) also describes a straightforward extension for the matching with indifferences. While EADAM is not strategy-proof, speculation is difficult and truth-telling is a Bayesian Nash equilibrium.

3.4 First-come-first-served course allocation system (FCFS)

Other than the two prominent stable matching mechanisms, we introduce a simple *first-come-first-served* course allocation system used at many universities. During the application period students can assign to exactly one course wherever free seats are available based on their time of arrival after the registration time starts. Changing the course assignment later is only possible *after* the existing course assignment is canceled by a student. Sometimes students just cannot be the first to register due to circumstances beyond their control and the outcome might be unstable and inefficient. Also, the registration process for groups of courses is often ordered sequentially, such that students have to decide in one period whether they register for a particular set of courses or wait for a later period with a course of higher priority. Now, students need to speculate on their probability of getting a seat in a more popular course with a later registration time, or they rather register for a less preferred course early.

4 Field experiments

Apart from the theoretical properties of the matching mechanisms described above it is interesting to understand, which differences in the outcome can be expected in the field. In what follows, we will discuss the data elicitation for field experiments and introduce the metrics, which allow for the comparison of outcomes.

4.1 Data collection

We collected preferences from students for two courses in the summer term 2012 and the winter term 2012/13.

The *first field experiment* was on a lecture with 136 students, who had to be assigned to 8 courses with a capacity of 17 seats. The students provided their preferences in the FCFS system and their ordinal preferences in a second system at about the same time. The comparison was only done with respect to the 136 students who provided their preferences in both systems. The students were told about SOSM and the fact that the mechanism is strategy-proof, and we would decide randomly on the outcome of one of the systems to finally determine the assignment.

We used the same data to compute efficient matchings in EADAM. While EADAM is not strategy-proof for students, truth-telling is an ordinal Bayesian Nash equilibrium and we suspect that students would not be able to game the mechanism strategically in such an environment. So, when we compute the EADAM outcomes, we assume that the students would also provide their preferences truthfully and consent to waive a certain priority since a consenting student causes himself no harm in EADAM.

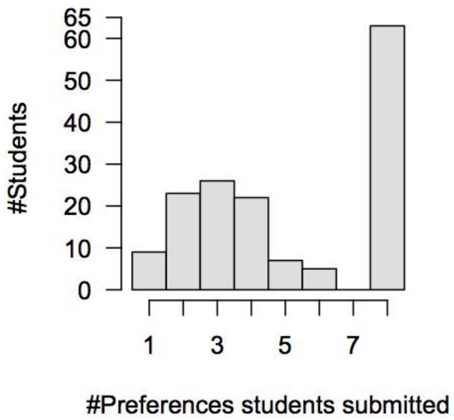

Figure 1 Histogram of number of preferences of the students in the first field experiment

Figure 1 shows the histogram of the number of preferences the students submitted. For instance 26 students submitted 3 preferences for the courses. We can also see that 63 students submitted preferences over all 8 courses.

The *second field experiment* in the summer term 2012 was on a large class with 794 students overall, and the students provided their preferences for 45 lab courses taking place at different times of the week. After the students registered in the first weeks of the semester in the FCFS system, we asked them to provide their preferences in an ordinal ranking within the subsequent 3 weeks. The sizes of the courses range from 14 to 22 students restricted by the room size. The 45 courses would have a capacity of 858 course seats.

The lecture was very large and for organizational reasons, we could not elicit all preferences as in the first field experiment. Overall, 418 of the 794 students provided their ordinal preferences, and only those students who provided their preferences and who were matched in the FCFS system were considered in a comparison. As a consequence, the group size of the courses was reduced proportionally by $\frac{418}{858} = 0{,}49 \approx \frac{1}{2}$ to allow for a comparison of the two approaches.

As indicated in the last section, in the FCFS system the students can often only register for groups of courses one after the other. This is done with large classes to avoid overload of the system. The registration process in this study began at 4 different points in time. The courses were classified by the weekday they were held. The time difference between these 4 points was 2 hours each. This can lead to speculation where students register for a group at an earlier time to prevent not getting a seat in any group.

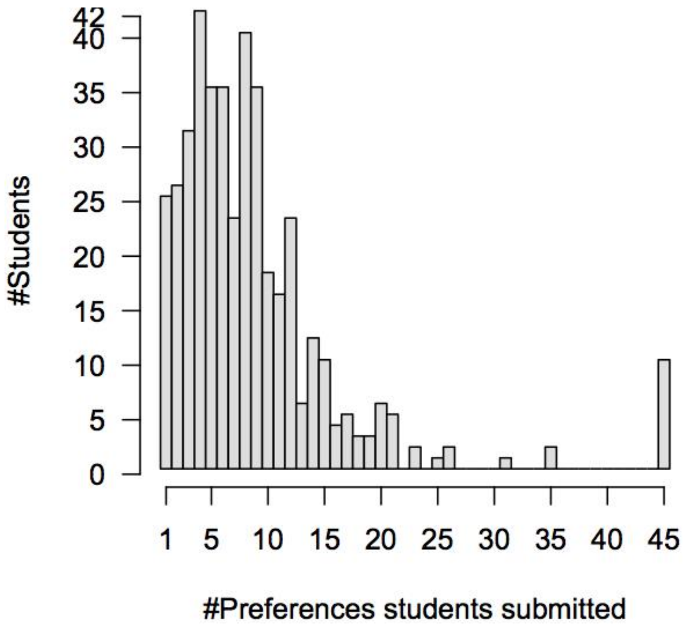

Figure 2 Histogram of number of preferences of the students in the second field experiment

The histogram of the number of preferences the students submitted in the second field experiment is shown in **Figure *2***. 10 students submitted preferences over all 45 courses. The 45 courses for this large class were distributed over 18 time slots with up to 5 courses. Although, we could not elicit all student preferences for this experiment for organizational reasons, we believe that the data provides another valuable observation and an estimate on how well stable matching performs as compared to an FCFS approach for large course allocation problems.

4.2 Generation of course organizers preferences

In both real applications course organizers did not use preferences on the assignment of students to different courses. We also wanted to get an understanding of the differences between SOSM and EADAM in case course organizers have preferences. In our department such preferences typically concern the type of study, the grades of a student, and taken pre-courses. In order to compare the matching mechanisms *with* preferences of course organizers we generated preferences of course organizers with indifferences across groups of students in four different ways, which are natural in our application domain.

First, we know the type of study (Computer Science (CS), Information Systems (IS), Mathematics, and Business) of all students participating in the experiments. In order to introduce realistic correlation among the student and the course organizer preferences we first looked at the top five student preferences for the different types of study. Depending on that we generated the type of study preferences for the courses. For example the course organizer of a course that is very popular among IS students would prefer IS students with a higher likelihood.

A second possibility for generating course organizers preferences are *pre-courses* being desired prerequisites for certain courses. We assumed five pre-courses that we determined for all of the courses according to the actual prerequisites. Based on real distributions on the taken pre-courses for the different types of study at our department we assigned the pre-courses to the students depending on their type of study. For instance a CS student has a higher probability to have taken a CS pre-course than a Business student. On the course organizers' side we randomly assigned the pre-courses according to the distribution of types of study in a course. As we had more CS students than Math students, more CS pre-course attendees were assigned to the courses. A course organizer prefers those students who have taken pre-courses for his course.

Third, we generated course organizers' preferences by looking at the *grades* of students. As we do not have the real grades of the students available due to privacy restrictions, we took the real distribution of average grades in our department and assigned grades to the students according to this distribution. Course organizers prefer higher grades across all courses, so their preferences are all the same for this preference.

Finally, we generated *combined preferences* by first assigning pre-course preferences to a course. Ties were broken by type of study and then by grade. We will report on the mean of 400 runs with different course organizer preferences, 100 for each method described above. This appears to be a likely type of preferences in our matching applications.

4.3 Metrics for matchings

In order to compare matching mechanisms, we will now introduce three metrics that allow for a comparison of the empirical results.

Average rank

Using the information on the ranks the students (courses) achieved is a standard way of gauging the welfare of the students (courses). In particular, the *average rank* has been used as a metric to gauge the difference in welfare of matching algorithms in Budish and Cantillon (2012, pp. 2237-2271) and Abdulkadiroğlu et. al. (2009, pp. 1954-1978), two of the few experimental papers on matching mechanisms. The following two tables show a comparison of SOSM (Example 2) and EADAM (Example 3) for students (Table 7) and courses (Table 8).

Tab. 7 Comparison of student average rank of SOSM and EADAM

	SOSM		EADAM	
Student	Assigned course	Rank of assigned course	Assigned course	Rank of assigned course
s_1	c_4	2	c_4	2
s_2	c_2	2	c_1	1
s_3	c_3	2	c_2	1
s_4	c_1	2	c_3	1
Average rank		**2**		**1.25**

The average rank for a single course with multiple seats is the average of the ranks of its assigned students. The average overall course rank is the mean of the average ranks per course.

Tab. 8 Comparison of course average rank of SOSM and EADAM

	SOSM		EADAM	
Course	Assigned student	Rank of assigned student	Assigned student	Rank of assigned student
c_1	s_4	1	s_2	3
c_2	s_2	1	s_3	2
c_3	s_3	1	s_4	2

c_4	s_1	1	s_1	1
Average rank		1		2

Popularity

As defined in Abraham et al. (2007, pp. 1030-1045), a matching μ' is *more popular than* another matching μ ($\mu' \succ \mu$), if the number of students (courses) that prefer μ' to μ exceeds the number of students (courses) that prefer μ to μ'. The preference of a course or student is the sign of the difference of the ranks. A comparison of the popularity of SOSM (Example 2) and EADAM (Example 3) for students is shown in Table 9. A similar popularity metric can be derived for the course organizers.

Tab. 9 Comparison of student popularity of SOSM and EADAM

	SOSM	EADAM		
Student	Rank of assigned course	Rank of assigned course	Rank difference	Preference
s_1	2	2	0	SOSM ~ EADAM
s_2	2	1	1	SOSM $\prec$ EADAM
s_3	2	1	1	SOSM $\prec$ EADAM
s_4	2	1	1	SOSM $\prec$ EADAM
			Popularity	**SOSM $\prec$ EADAM**

Rank distribution

The *rank distribution* compares how many students were assigned to their first choice, how many to their second choice, and so on. Similarly, for courses it compares how many students were assigned to the first choice of a course. Table 10 shows a comparison of SOSM (Example 2) and EADAM (Example 3) for students S (left) and courses C (right) with the rank distribution metric.

Tab. 10 Comparison of rank distribution of SOSM and EADAM for students and courses

Students

	SOSM	EADAM
Rank	#Students	#Students

Courses

	SOSM	EADAM
Rank	#Students	#Students

1	0	3
2	4	1

1	4	1
2	0	2
3	0	1

4.4 Results of the first field experiment

With the preferences of the 136 students from the first field experiment and the data about the courses we ran SOSM and EADAM in order to compute a matching of the students and courses. We first provide a comparison, where course organizers do not have preferences, and a second comparison, where we generated preferences of course organizers over students (see Section 3.2). Note that without preferences of group organizers there is no difference between SOSM and EADAM. Students will only be rejected when the group is full. This means that no student will be tentatively placed to a group when other students are rejected from that group and later on also be rejected from that group. If there are no interrupters, the preferences of the students will not be changed and EADAM will stop in step 1, thus returning the same result as SOSM. However, we decided to report both results (SOSM and EADAM) in both subsections for completeness.

4.4.1 No preferences of course organizers

The average rank metric (Table 11) shows that both SOSM and EADAM result in better matchings than the FCFS, even though the differences are small. The shorthand n/p refers to students being assigned to courses they have no preference for. If these courses are unacceptable, these students would not be matched.

Tab. 11 Student average ranks of FCFS, SOSM and EADAM without group preferences

	FCFS	SOSM	EADAM
Average rank	$1.31\ (5\ n/p)$	$1.25\ (6\ n/p)$	$1.25\ (6\ n/p)$

The popularity metric in Table 12 also shows the superiority of SOSM and EADAM.

Tab. 12 Student popularity of FCFS, SOSM and EADAM without group preferences

	FCFS vs. SOSM	FCFS vs. EADAM	SOSM vs. EADAM
#Students $\succ$	16	16	0
#Students $\prec$	20	20	0

Popularity	FCFS $\prec$ SOSM	FCFS $\prec$ EADAM	SOSM $\sim$ EADAM

The rank distribution metric (Table 13) demonstrates that 114 students would be matched to their first preference with SOSM or EADAM, whereas only 105 are matched to the first preference in the first-come-first-served course allocation system.

Tab. 13 Student rank distribution of FCFS, SOSM and EADAM without group preferences

Rank	FCFS		SOSM		EADAM	
1	105	(77.2%)	114	(83.8%)	114	(83.8%)
2	13	(9.6%)	5	(3.7%)	5	(3.7%)
3	12	(8.8%)	6	(4.4%)	6	(4.4%)
4	-		4	(2.9%)	4	(2.9%)
5	1	(0.7%)	1	(0.7%)	1	(0.7%)
n/p	5	(3.7%)	6	(4.4%)	6	(4.4%)

Note that the first-come-first-served course allocation system assigned 5 students to a group where the students have no preference (n/p) for, SOSM or EADAM would assign 6 students to a group where they have no preference for as they require a stable matching. Typically, it is students who submit only a small number of course preferences, who are assigned a course for which they have provided no preference. Nevertheless, this illustrates that FCFS can lead to mores students being matched to a course for which they have provided a preference, compared to SOSM and EADAM. Other than that, the metrics show that SOSM and EADAM result in better matchings than the first-come-first-served course allocation system when no group preferences exist.

4.4.2 With preferences of course organizers

In order to compare the matching mechanisms with preferences of course organizers we generated random group preferences with indifferences over the 136 students. The following results show the mean of 400 runs with the generated group preferences with indifferences.

As can be expected, FCFS yields unstable results. Unstable means that blocking pairs of unmatched students and schools exist, who would prefer to be matched with each other rather than their current assignment. We have identified 25 blocking pairs (18% of all students) on average.

Students

For the students the average rank metric (Table 14) shows that EADAM would provide the best results. On the same metric, SOSM would result in a better matching than FCFS.

Tab. 14 Student average ranks of FCFS, SOSM and EADAM with generated group preferences

	FCFS	SOSM	EADAM
Average rank	1.31 (5 n/p)	1.28 (6.88 n/p)	1.26 (6.88 n/p)

In contrast to the average rank metric, the popularity metric (see **Tab. *15***) shows a different result for the comparison of the FCFS course allocation system and SOSM. In this case FCFS is more popular than SOSM. This discrepancy is due to the fact that the 19.19 students (mean of 400 runs) who were better off with SOSM were matched to a much better ranked group than the 19.33 students being better off with FCFS. Hence the average rank for SOSM is lower than for FCFS. Still, EADAM shows the best results.

Tab. 15 Student popularity of FCFS, SOSM and EADAM with generated group preferences

	FCFS vs. SOSM	FCFS vs. EADAM	SOSM vs. EADAM
#Students $\succ$	19.33	18.23	0.00
#Students $\prec$	19.19	19.76	2.40
Popularity	**FCFS $\succ$ SOSM**	**FCFS $\prec$ EADAM**	**SOSM $\prec$ EADAM**

The same result as the average rank metric is found based on the rank distribution metric in Table 16 with EADAM matching the most students to their first preference.

Tab. 16 Student rank distribution of FCFS, SOSM and EADAM with generated group preferences

Rank	FCFS		SOSM		EADAM	
1	105	(77.2%)	106.4	(78.2%)	108.6	(79.8%)
2	13	(9.6%)	12.2	(9.0%)	10.5	(7.7%)
3	12	(8.8%)	8.3	(6.1%)	7.9	(5.8%)
4	-		1.9	(1.4%)	1.8	(1.3%)
5	1	(0.7%)	0.3	(0.2%)	0.3	(0.2%)
6	-		0.1	(0.1%)	0.1	(0.1%)
n/p	5	(3.7%)	6.9	(5.1%)	6.9	(5.1%)

With preferences of course organizers, EADAM results in the best matching for the students.

Courses

For the courses the average rank metric (Table 17) shows that SOSM produces the best matchings regarding the preferences of course organizers, but SOSM and EADAM would both be preferred to FCFS by course organizers.

Tab. 17 Group average ranks of FCFS, SOSM and EADAM with generated group preferences

	FCFS	SOSM	EADAM
Average rank	66.90 (0 n/p)	63.33 (0 n/p)	64.44 (0 n/p)

The same observation can also be made with regard to the popularity metric that is shown in Table 18.

Tab. 18 Group popularity of FCFS, SOSM and EADAM with generated group preferences

	FCFS vs. SOSM	FCFS vs. EADAM	SOSM vs. EADAM
#Courses $\succ$	3.62	3.69	2.02
#Courses $\prec$	4.38	4.31	0.84
Popularity	**FCFS $\prec$ SOSM**	**FCFS $\prec$ EADAM**	**SOSM $\succ$ EADAM**

The reason EADAM is less preferable to the course organizers is that it aims to increase the welfare of the students when it does improvements of the students. Also, the rank distribution metric shows that course organizers would prefer SOSM.

Tab. 19 Group rank distribution of FCFS, SOSM and EADAM with generated group preferences

Rank	FCFS		SOSM		EADAM	
1	1.1	(0.8%)	1.2	(0.9%)	1.1	(0.8%)
2	1.0	(0.7%)	1.1	(0.8%)	1.1	(0.8%)
3	0.9	(0.7%)	1.1	(0.8%)	1.1	(0.8%)
4	1.0	(0.7%)	1.0	(0.7%)	1.1	(0.8%)
5	1.1	(0.8%)	1.1	(0.8%)	1.2	(0.9%)
6	1.1	(0.8%)	1.1	(0.8%)	1.0	(0.7%)
7	1.1	(0.8%)	1.1	(0.8%)	1.1	(0.8%)
8	1.2	(0.9%)	1.3	(1.0%)	1.2	(0.9%)
9	1.0	(0.7%)	1.0	(0.7%)	1.0	(0.7%)
10	1.1	(0.8%)	1.1	(0.8%)	1.2	(0.9%)
11-20	10.8	(7.9%)	11.7	(8.6%)	11.3	(8.3%)
21-30	10.3	(7.6%)	11.5	(8.5%)	10.9	(8.0%)

31-50	19.9	(14.7%)	21.8	(16.1%)	21.2	(15.6%)
51-100	50.9	(37.4%)	50.9	(37.5%)	51.1	(37.6%)
101-136	33.5	(24.6%)	28.8	(21.2%)	30.2	(22.2%)

Overall, EADAM and SOSM would result in the better outcomes than FCFS for both students and course organizers, with or without preferences of course organizers in this field experiment.

4.5 Results of the second field experiment

In the second field experiment, we again analyze situations with or without preferences of group organizers. The second experiment is larger and we wanted to understand if the results from the first experiment carry over.

4.5.1 No preferences of course organizers

The average rank metric (Table 20) shows that with FCFS the students would be assigned to their 2.91th choice on average with 38 students being assigned to a course where they have no preference (n/p) for. However, with SOSM or EADAM the students would be assigned to their 1.78th choice and 53 students to a course they have no preference for.

Tab. 20 Student average ranks of FCFS, SOSM and EADAM without group preferences

	FCFS	SOSM	EADAM
Average rank	2.91 (38 n/p)	1.78 (53 n/p)	1.78 (53 n/p)

The popularity metric is shown in Table 21, which demonstrates that both SOSM and EADAM would be more popular than FCFS: Only 90 students would prefer FCFS to SOSM or EADAM, but 146 would prefer SOSM or EADAM to the first-come-first-served course allocation system.

Tab. 21 Student popularity of FCFS, SOSM and EADAM without group preferences

	FCFS vs. SOSM	FCFS vs. EADAM	SOSM vs. EADAM
#Students $\succ$	90	90	0
#Students $\prec$	146	146	0
Popularity	**FCFS $\prec$ SOSM**	**FCFS $\prec$ EADAM**	**SOSM $\sim$ EADAM**

If we look at the rank distribution metric (see Table 22) it shows that with SOSM or EADAM 280 students would be assigned to their first choice, whereas only 223 got their first choice with the first-come-first-served course allocation system. The first-come-first-served course allocation system assigned 38 students to a group where the students have no preference (n/p) for, SOSM or EADAM would assign 53 students to a group where they have no preference for. The reason for this could be the fact, that many students submitted very few preferences, as we saw in Figure 2.

Tab. 22 Student rank distribution of FCFS, SOSM and EADAM without group preferences

Rank	FCFS		SOSM		EADAM	
1	223	(53.3%)	280	(67.0%)	280	(67.0%)
2	30	(7.2%)	25	(6.0%)	25	(6.0%)
3	31	(7.4%)	18	(4.3%)	18	(4.3%)
4	30	(7.2%)	14	(3.3%)	14	(3.3%)
5	18	(4.3%)	7	(1.7%)	7	(1.7%)
6	8	(1.9%)	4	(1.0%)	4	(1.0%)
7	13	(3.1%)	7	(1.7%)	7	(1.7%)
8	4	(1.0%)	2	(0.5%)	2	(0.5%)
9	5	(1.2%)	1	(0.2%)	1	(0.2%)
10	6	(1.4%)	3	(0.7%)	3	(0.7%)
11	1	(0.2%)	1	(0.2%)	1	(0.2%)
12	1	(0.2%)	2	(0.5%)	2	(0.5%)
13	1	(0.2%)	1	(0.2%)	1	(0.2%)
16	2	(0.5%)	-		-	
17	1	(0.2%)	-		-	
19	1	(0.2%)	-		-	
20	2	(0.5%)	-		-	
23	1	(0.2%)	-		-	
35	1	(0.2%)	-		-	
41	1	(0.2%)	-		-	
n/p	38	(9.1%)	53	(12.7%)	53	(12.7%)

The three metrics average rank, popularity and rank distribution all show that without group preferences, SOSM or EADAM would result in a better matching for the students. Again, a disadvantage is the higher number of students matched to a group where they do not have preferences for, which is a consequence of the fact that students often only provided preferences for a few courses.

4.5.2 With preferences of course organizers

In this case FCFS induces 189 blocking pairs (45% of all students) on average, meaning that 189 pairs of student and school could improve their assignment by switching.

Students

The average ranks for students if course organizers have preferences are shown in Table 23. Using SOSM instead of FCFS the average rank improvement for each student would be 0.73. With EADAM rather than FCFS the average rank improvement would be 0.87. This does not mean that every student would improve by 0.87 ranks. Some might not improve or even worsen, and some would improve by more than one rank.

Tab. 23 Student average ranks of FCFS, SOSM and EADAM with generated group preferences

	FCFS	SOSM	EADAM
Average rank	2.91 (38 n/p)	2.18 (51.26 n/p)	2.04 (51.24 n/p)

The popularity metric (see Table 24) leads to the same result as the average rank metric with EADAM showing the best results for the students.

Tab. 24 Student popularity of FCFS, SOSM and EADAM with generated group preferences

	FCFS vs. SOSM	FCFS vs. EADAM	SOSM vs. EADAM
#Students $\succ$	125.81	115.99	0.00
#Students $\prec$	138.39	141.04	27.35
Popularity	**FCFS $\prec$ SOSM**	**FCFS $\prec$ EADAM**	**SOSM $\prec$ EADAM**

As shown in Table 25 the rank distribution metric provides the same results as the popularity metric. EADAM matches almost 75% of the students to their top three choice courses while SOSM shows worse results than FCFS. The comparison of FCFS and SOSM shows that with FCFS slightly more students are matched to their first choice course than with SOSM, but looking at the top three choices, SOSM shows the better results. Similar to the results in Section 5.1 with SOSM and EADAM more students would be matched to courses they do not have preferences for (n/p).

Tab. 25 Student rank distribution of FCFS, SOSM and EADAM with generated group preferences

Rank	FCFS		SOSM		EADAM	
1	223	(53.3%)	219.9	(52.6%)	236.5	(56.6%)
2	30	(7.2%)	52.5	(12.6%)	46.2	(11.1%)
3	31	(7.4%)	29.8	(7.1%)	27.0	(6.5%)
4	30	(7.2%)	23.4	(5.6%)	21.6	(5.2%)
5	18	(4.3%)	13.4	(3.2%)	11.4	(2.7%)
6	8	(1.9%)	8.8	(2.1%)	7.4	(1.8%)
7	13	(3.1%)	6.0	(1.4%)	5.4	(1.3%)
8	4	(1%)	4.6	(1.1%)	3.9	(0.9%)
9	5	(1.2%)	2.7	(0.6%)	2.3	(0.6%)
10	6	(1.4%)	2.3	(0.6%)	1.9	(0.5%)
11	1	(0.2%)	0.8	(0.2%)	0.6	(0.1%)
12	1	(0.2%)	1.1	(0.3%)	1.1	(0.3%)
13	1	(0.2%)	0.6	(0.1%)	0.5	(0.1%)
14	-		0.1	(0%)	0.1	(0%)
15	-		0.3	(0.1%)	0.2	(0%)
16	2	(0.5%)	0.1	(0%)	0.1	(0%)
17	1	(0.2%)	0.2	(0%)	0.2	(0%)
19	1	(0.2%)	0.1	(0%)	0.1	(0%)
20	2	(0.5%)	0.1	(0%)	0.1	(0%)
23	1	(0.2%)	-		-	
35	1	(0.2%)	-		-	
41	1	(0.2%)	-		-	
n/p	38	(9.1%)	51.3	(12.3%)	51.2	(12.3%)

The significant differences between the results of SOSM and EADAM emphasize the welfare losses for students that could be produced by SOSM.

The difference in the results for the students without group preferences (Section 5.1) is due to the fact that the group preferences affect the assignment process. For the case with group preferences a student could be rejected from a course after being tentatively assigned to it. This could not happen without group preferences.

Courses

The average rank metric (see Table 26) for courses shows that both SOSM and EADAM result in better matchings for the courses with SOSM as the better alternative. It is clear that SOSM and EADAM provide better results than the first-come-first-served course allocation system since they are considering the group preferences. The better result of SOSM versus EADAM can again be explained by the fact that EADAM is improving the welfare of the students, thus negatively affecting the course organizers.

Tab. 26 Group average ranks of FCFS, SOSM and EADAM with generated group preferences

	FCFS	SOSM	EADAM
Average rank	198.66 (0 n/p)	178.57 (0 n/p)	186.78 (0 n/p)

Table 27 provides the popularity metric for the courses. The outcomes are the same as we saw with the average rank metric. Accordingly, SOSM brings the best results for the courses and it is more popular than the other two mechanisms.

Tab. 27 Group popularity of FCFS, SOSM and EADAM with generated group preferences

	FCFS vs. SOSM	FCFS vs. EADAM	SOSM vs. EADAM
#Courses $\succ$	17.00	17.72	18.85
#Courses $\prec$	28.00	27.28	6.10
Popularity	**FCFS $\prec$ SOSM**	**FCFS $\prec$ EADAM**	**SOSM $\succ$ EADAM**

In Table 28 the rank distribution for courses is provided. For example using SOSM, 15.6 students (3.7%) would be matched to the top 10 ranks of a course. Also in that case, the results show that SOSM would be the best matching mechanism for the courses with EADAM providing better results than the first-come-first-served course allocation system.

Tab. 28 Group rank distribution of FCFS, SOSM and EADAM with generated group preferences

Rank	FCFS	SOSM	EADAM
1	1.3 (0.3%)	1.7 (0.4%)	1.4 (0.3%)
2	1.1 (0.3%)	1.6 (0.4%)	1.3 (0.3%)
3	1.2 (0.3%)	1.4 (0.3%)	1.3 (0.3%)
4	1.2 (0.3%)	1.6 (0.4%)	1.5 (0.4%)
5	1.3 (0.3%)	1.6 (0.4%)	1.5 (0.4%)
6	1.2 (0.3%)	1.4 (0.3%)	1.3 (0.3%)
7	1.2 (0.3%)	1.6 (0.4%)	1.2 (0.3%)

8	1.3	(0.3%)	1.6	(0.4%)	1.4	(0.3%)
9	1.2	(0.3%)	1.4	(0.3%)	1.3	(0.3%)
10	1.1	(0.3%)	1.7	(0.4%)	1.4	(0.3%)
11-20	11.7	(2.8%)	15.8	(3.8%)	13.4	(3.2%)
21-30	12.2	(2.9%)	16.8	(4%)	14.0	(3.3%)
31-50	23.1	(5.5%)	29.5	(7.1%)	26.3	(6.3%)
51-100	54.9	(13.1%)	61.0	(14.6%)	58.6	(14%)
101-150	52.7	(12.6%)	53.1	(12.7%)	53.3	(12.7%)
151-200	51.5	(12.3%)	50.8	(12.2%)	51.3	(12.3%)
201-300	99.9	(23.9%)	91.5	(21.9%)	96.6	(23.1%)
301-418	100.0	(23.9%)	84.0	(20.1%)	91.0	(21.8%)

Overall, EADAM resulted in better outcomes than FCFS for students and courses with or without preferences of course organizers. The experiments provide some useful information. First, the results suggest that EADAM is the preferred alternative, although the differences are not substantial. Second, there are trade-offs and the stable matching mechanism on average have more students unmatched, which might be an important criterion in some applications.

5 Towards multiunit and combinatorial assignments

Course assignment has become a popular research topic with many recent developments, as applications are wide-spread. One none-trivial extension that researchers have considered only recently are minimum quotas for courses. This is practically relevant if course organizers want to avoid situations where they have only a few students in their class. It is known that there is no strategy-proof mechanism that completely eliminates justified envy when minimum quotas are imposed (Hamada et al. 2011, pp. 180-191). Ueda et al. (2012, pp. 1327-1328) recently proposed strategy-proof mechanisms, which allows for minimum quotas and achieve efficiency with a weaker notion of stability.

Much of the new developments have been on the assignment of multiple course seats to one student, i.e., the *multiunit assignment problem* or *combinatorial assignment problem*, resp. Budish and Cantillon (2012, pp. 2237-2271) analyze *multiunit assignment problems*, where students want to get seats, not only a single out of many possible courses, and their preferences are ordinal and responsive. *Responsiveness* describes a form of separability of preferences about different courses: If a student prefers one course c over the other, c', he also prefers a bundle of courses $X \cup \{c\}$ to a bundle of courses $X \cup \{c'\}$. Theoretical analyses show that serial dictatorships are the only strategy-proof and efficient mechanisms for this problem (Pápai 2001, pp. 257-271; Ehlers and Klaus 2003, pp. 265-280). Unfortunately, they lead to highly unfair outcomes in which some students get all the courses they like most, whereas others only a few. Budish and Cantillon (2012, pp. 2237-2271) suggest a mechanism with proxy agents, which tries to get fairness of the distribution and efficiency at the expense of strategy-proofness.

The *combinatorial assignment problem* can be seen as an extension of the multiunit assignment problem, where a set of course seats is to be allocated amongst a set of agents with preferences over bundles of course seats in different courses. Each student wants to get 0 or 1 seat in a course, but his preferences for particular bundles of courses might be substitutes or complements as opposed to the responsive preferences assumed in the multiunit assignment problem. This is similar to a combinatorial auction in which no monetary transfers are allowed. In recent work by Budish (2011, pp. 1061-1103) suggests a mechanism, which adapts the idea of Competitive Equilibrium from Equal Incomes (Varian 1976, pp. 249-260) to environments with indivisible goods and proofs approximations of efficiency, strategy-proofness, and fairness properties. Othman et al. (2010, pp. 873-880) describe computational methods to implement the mechanism.

Another interesting extension of SOSM in the context of the hospitals/residents problem is the consideration of couples. The set of residents also includes couples who must be assigned together, either to the same hospital or to a specific pair of hospitals chosen by the couple. The addition of couples to the hospitals/residents problem renders the problem *NP*-complete (Gusfield and Irving 1989).

6 Conclusion

We have discussed established and more recent algorithmic developments for course allocation problems. The Gale-Shapley student optimal stable mechanism (SOSM) is a well-established matching mechanism which is applied for the assignment of students to public universities in the U.S. and in the school choice problems, for instance. While strategy-proof for students and stable, SOSM may not always be efficient. Kesten’s efficiency adjusted deferred acceptance mechanism (EADAM) (Kesten 2010, pp. 1297-1348) eliminates welfare losses on the students' side at the expense of strategy-proofness. It is a bit surprising that these algorithms have found little application in the large number of matching problems in university environments so far and first-come-first-served (FCFS) approaches are still in wide-spread use.

In this paper, we provided a brief survey of this field, which has seen a number of recent advances. In addition, we discuss the results of two field experiments, in which we compare an FCFS mechanism with SOSM and EADAM. The FCFS matching and the matchings of SOSM and EADAM were compared using average rank metrics, popularity, and rank distribution metrics. By and large, the outcomes of stable matching mechanisms were preferable to FCFS in both experiments, however, the differences were not substantial. Actually, the number of unmatched students was lower in FCFS throughout. For practitioners it is worthwhile to understand the different design desiderata and metrics before introducing a new matching mechanism. Still, one of the biggest advantages of SOSM and EADAM beyond the stability of the outcomes might be the fact that there are strong incentives for telling the truth and there are no benefits to being first during the registration period, which might lead to unfair outcomes as some students can just not be first due to circumstances beyond their control.

Bibliography

Abdulkadiroğlu A, Pathak P, Roth AE (2009) Strategy-proofness versus efficiency in matching with indifferences: Redesigning the NYC high school match. American Economic Review 99(5): 1954–1978

Abdulkadiroğlu A, Pathak P, Roth AE, Sonmez T (2006) Changing the Boston school choice mechanism. National Bureau of Economic Research

Abdulkadiroğlu A, Sönmez T (2003) School choice: A mechanism design approach. American Economic Review 93(3): 729-747

Abraham DJ, Irving RW, Kavitha T, Mehlhorn K (2007) Popular matchings. SIAM Journal on Computing 37(4): 1030-1045

Bapna R, Goes P, Gupta A, Jin Y (2004) User Heterogeneity and Its Impact on Electronic Auction Market Design: An Empirical Exploration. MIS Quarterly 28(1): 21-43

Bichler M, Pikovsky A, Setzer T (2009) An analysis of design problems in combinatorial procurement auctions. Business & Information Systems Engineering 51(1): 111-117

Bichler M, Gupta A, Ketter W (2011) Designing smart markets. INFORMS Information Systems Research 21: 688-699

Braun S, Dwenger N, Kübler D (2007). Telling the truth may not pay off: An empirical study of centralised university admissions in Germany. IZA

Budish E (2011) The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. Journal of Political Economy 119(6): 1061-1103

Budish E, Cantillon E (2012) The Multi-unit Assignment Problem: Theory and Evidence from Course Allocation at Harvard. American Economic Review 102(5): 2237-2271

Chen Y, Sönmez T (2006) School choice: an experimental study. Journal of Economic Theory 127(1): 202-231

Ehlers L, Klaus B (2003) Coalitional strategy-proof and resource-monotonic solutions for multiple assignment problems. Social Choice and Welfare 21(2): 265-280

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

Featherstone C, Niederle M (2008) Ex ante efficiency in school choice mechanisms: an experimental investigation. National Bureau of Economic Research

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

Gusfield D, Irving RW (1989) The Stable Marriage Problem: Structure and Algorithms. MIT Press, Cambridge
Halldórsson MM, Irving RW, Iwama K, Manlove DF, Miyazaki S, Morita Y, Scott S (2003) Approximability results for stable marriage problems with ties. Theoretical Computer Science 306(1): 431-447
Hamada K, Iwama K, Miyazaki S (2011) The hospitals/residents problem with quota lower bounds. In: Algorithms – ESA 2011 - 19th Annual European Symposium. Springer, Berlin Heidelberg, pp 180-191
Kesten O (2010) School choice with consent. The Quarterly Journal of Economics 125(3): 1297-1348
Krishna A, Ünver MU (2005) A Field Experiment on Course Bidding at Business Schools. EconWPA
Krishna A, Ünver MU (2008) Research Note—Improving the Efficiency of Course Bidding at Business Schools: Field and Laboratory Studies. Marketing Science 27(2): 262-282
Manlove D (2013). Algorithmics of matching under preferences. World Scientific Publishing, Singapore.
Manlove DF, Irving RW, Iwama K, Miyazaki S, Morita Y (2002) Hard variants of stable marriage. Theoretical Computer Science 276(1): 261-279
Milgrom P (2011) Critical issues in the practice of market design. Economic Inquiry 49(2): 311-320
Othman A, Sandholm T, Budish E (2010) Finding approximate competitive equilibria: Efficient and fair course allocation. In: Proc 9th International Conference on Autonomous Agents and Multiagent Systems: Volume 1, Toronto, pp 873-880
Pápai S (2001) Strategyproof and nonbossy multiple assignments. Journal of Public Economic Theory 3(3): 257-271
Roth AE (1982) The economics of matching: Stability and incentives. Mathematics of Operations Research 7(4): 617-628
Roth AE (1984) The evolution of the labor market for medical interns and residents: a case study in game theory. Journal of Political Economy 92: 991-1016
Roth AE (2002) The economist as engineer: Game theory, experimentation, and computation as tools for design economics. Econometrica 70(4): 1341-1378
Shapley L, Scarf H (1974) On cores and indivisibility. Journal of Mathematical Economics 1(1): 23-37
Sönmez T, Ünver MU (2010) Course Bidding at Business Schools. International Economic Review 51(1): 99-123
Ueda S, Fragiadakis D, Iwasaki A, Troyan P, Yokoo M (2012) Strategy-proof mechanisms for two-sided matching with minimum and maximum quotas. In: Proc 11th International Conference on Autonomous Agents and Multiagent Systems: Volume 3, Valencia, pp 1327-1328
Varian HR (1976) Two problems in the theory of fairness. Journal of Public Economics 5(3): 249-260
Weinhardt C, Holtmann C, Neumann D (2003) Market Engineering. Wirtschaftsinformatik 45(6): 635–640
Westkamp A (2013) An analysis of the German university admissions system. Economic Theory 53(3): 561-589