

A Field Experiment on Course Bidding at Business Schools*

Aradhna Krishna[†]

M. Utku Ünver[‡]

University of Michigan

Koç University

March 2005

Abstract

Allocation of course seats to students is a challenging task for registrars' offices in universities. Since demand exceeds supply for many courses, course allocation needs to be done equitably and efficiently. Many schools use bidding systems where student bids are used both to infer preferences over courses and to determine student priorities for courses. However, this dual role of bids can result in course allocations not being market outcomes and unnecessary efficiency loss, which can potentially be avoided with the use of an appropriate market mechanism. We report a field experiment done at the University of Michigan Business School in Spring 2004 comparing its typical course bidding mechanism with the alternate Gale-Shapley Pareto-dominant market

*We are indebted to Tayfun Sönmez for extensive comments and suggestions. Preliminary versions of this study were presented at Society of Economic Design Conference in Mallorca, Spain, World Congress of Game Theory in Marseille, France, Social Choice and Welfare Conference in Osaka, Japan, and Murat Sertel Theory Conference in İstanbul, Turkey. We would like to thank the participants for comments. Ünver gratefully acknowledges support from the NSF. All errors are our own responsibility. The field experiment data can be downloaded from the world-wide web at <http://home.ku.edu.tr/~uunver/research/UMBSexperimentdata.xls>.

[†]University of Michigan Business School, 701 Tappan St. Ann Arbor, MI 48109-1234.

[‡]Koç University, Department of Economics, Rumelifeneri Yolu, Sarıyer, 34450, İstanbul, Turkey.

mechanism. Our results suggest that using the latter could vastly improve efficiency of course allocation systems while facilitating market outcomes.

1 Introduction

Registrars' offices at most universities face the daunting task of allocating course seats to students. Since the learning experience of students is a direct function of the courses they take, and demand for many courses often exceeds supply, it is important that courses are allocated equitably and efficiently across students. While many schools use preference revelation (preference ranking) mechanisms to allocate courses (e.g., Stanford Business School, Harvard Business School), others use bidding mechanisms for the same purpose (e.g., Columbia Business School, Yale School of Management). Several schools have recently moved from preference revelation mechanisms to bidding mechanisms (e.g., the University of Michigan, Ross School of Business, henceforth referred to as UMBS), considering the latter superior in terms of efficiency. However, under the bidding mechanisms, bids are not only used to infer student preferences but also to determine student priorities for courses. This dual role of bids results in the schools' course allocations not being market outcomes, that is, the announced "prices" for courses and the announced "course allocation" do not actually clear the market. That is, students can be better off at other schedules versus their allotted course allocation and they can afford to "buy" these schedules at the announced course prices using their bid distribution. Typically, add-drop periods at the beginning of semesters are supposed to correct such failures, if there are any, however these periods can be congested and furthermore, in theory, it may not be possible to correct such a failure after it occurs. Thus, the current bidding mechanism results in unnecessary efficiency loss (Sönmez and Ünver 2003). Although theory predicts this efficiency loss, its existence and magnitude has not yet been tested.

As such, we do not know if policy makers should spend the effort to move away from the current bidding mechanisms.

We report a field experiment carried out at UMBS immediately after the bidding period for the Spring 2004 semester. The experiment is designed to test whether typical bidding systems (in particular the UMBS mechanism) result in efficiency loss in real-life applications, and if so how much efficiency improvement can be obtained through a transition to a market mechanism such as the Gale-Shapley Pareto-dominant market mechanism. We show that the current systems in place can be vastly improved in terms of (Pareto) efficiency, making a large proportion of students (approximately 20% in our study) better off. The key is “separating” the two roles of the bids by simply asking students to submit their preference ranking of the courses in addition to bidding for the courses. Our results have the potential to affect the learning experience of very large numbers of students enrolling in business schools and other institutions which use similar bidding mechanisms for course allocation.

The rest of the paper is organized as follows. We first elaborate on literature background related to this paper. Then we describe currently used course bidding mechanisms. This is followed by a detailed description of course bidding at UMBS where we did our field study. We next describe the alternative GS mechanism that we test, and then the results from the field study. We end with implications and limitations of the research.

2 Literature Background and Related Research on Matching Market Design

This research falls into the area of market design, more specifically mechanism design for real-life “matching problems.” For example, the new hospital-intern matching mechanism proposed by Roth and Peranson (1999) was adopted in 1997 by the National Resident Matching Program, a centralized clearinghouse for the entry-level labor market for new physicians in the United States. Roth (2002) gives an extensive survey of “engineering” approach in mechanism design for real-life markets. For example, see Abdulkadiroğlu, Pathak and Roth (2005), Abdulkadiroğlu, Pathak, Roth and Sönmez (2005), Niederle and Roth (2005), and Roth, Sönmez and Ünver (2005) for current issues on adoption of market designs in several matching markets: New York City high school match, Boston public school admissions, resident matching to gastroenterology specialty fellowships, and kidney exchange for paired donations for organ transplants, respectively. Many of the designs have been supported by experimental studies which test the predictions of the theoretical design by emulating markets before and after the design (Kagel and Roth 2000, Ünver 2001, Haruvy, Roth and Ünver 2001, Chen and Sönmez 2002 and 2003, McKinney, Niederle and Roth 2004, and Niederle and Roth 2004). We hope that our study which tests the Sönmez and Ünver (2003) theory in the field will also influence policy makers to consider a new market design. Moreover, our experiment is the first (controlled) field experiment in the matching literature.

Next, we give a brief background on theoretical matching models and some of the mechanisms that are most related to our research. One of the most commonly used matching models is due to Gale and Shapley (1962), known as the marriage model or two-sided matching model (see Roth and Sotomayor 1990 for an extended survey of two-sided matching models). The two-sided matching

problem consists of two sets of players – firms and workers – that need to be matched with each other using preferences of firms over workers and of workers over firms. The central solution concept in this domain is stability, i.e. finding matchings of firms with workers, and the same workers with the same firms, so that no firm-worker pair would rather be matched with each other versus their allotted partners (and no firm or worker would rather stay unmatched than be with their allotted partner). Gale and Shapley also proposed two stable matching mechanisms in their seminal paper. Many real-life markets have adopted mechanisms based on Gale-Shapley mechanisms: Roth and Peranson’s (1999) redesign for the American hospital-intern matching mechanism is a variant of the Gale-Shapley (1962) intern-optimal stable mechanism. The previously used mechanism in the American market and the mechanisms in some regional British hospital-intern markets are variants of the Gale-Shapley mechanisms as well (Roth 1984 and 1991).

Our problem, allocation of course seats to students, substantially differs from a two-sided matching problem. This problem is a variant of the house allocation problem, in which indivisible objects (houses) need to be assigned to agents each of whom has preferences over these objects – agents are the only players in this model, objects are not players. Random serial dictatorship is one of the most used mechanisms for this problem, in which agents are randomly ordered in a linear order and agents choose their favorite object among the available ones, one at a time, according to this order. In many North American college campuses, systems variants of this mechanism are used to allocate dormitory rooms to students (Abdulkadiroğlu and Sönmez 1999). Course allocations in Stanford Graduate School of Business School and Harvard Business School are done using systems that are variants of this approach. Starting with Balinski and Sönmez (1999), many studies showed similarities between plausible mechanisms for the house allocation problem and plausible mechanisms for the two-sided matching problem.

If there is additional structure to the house allocation problem, such as priorities of agents over the houses, then mechanisms that are invented for two-sided matching can be used for these problems as well, whenever a normative criterion similar to “stability” needs to be respected. The underlying idea is converting the object allocation problem to an induced two-sided matching problem by treating the priorities of the agents as “the preferences of the objects over agents.” Balinski and Sönmez showed that the Turkish college admissions mechanism used to place high school students to colleges in Turkey (based on exam scores determined by a centralized college admissions test) is equivalent to the Gale-Shapley college-optimal stable mechanism for the induced two-sided matching problem. Similarly, Abdulkadiroğlu and Sönmez (2003) proposed the Gale-Shapley student-optimal stable mechanism for student admissions to K-12 schools in U.S. public schools, where students have priorities for schools in their neighborhoods constituted by the U.S. constitution. These priorities can be used to induce preferences of schools over students in order to create an induced two-sided matching market; then the Gale-Shapley student-optimal stable mechanism can be used in this induced two-sided matching market to find a stable matching. A version of this proposal was adopted by New York City high schools in 2004.¹

When bidding is used as a tool in the course allocation problem with the intent of reaching market outcomes (as done in many business schools), we can induce a two-sided matching market using student bids for each course as induced preferences of the courses, i.e., the courses are assumed to prefer students who bid a higher amount for them (Sönmez and Ünver 2003). Sönmez and Ünver proposed the Gale-Shapley student-optimal stable mechanism as an alternative to the current bidding mechanisms used in many business schools. They showed that this mechanism is not only a market mechanism but also the Pareto-dominant one among all market mechanisms, while UMBS

¹Also see Ergin (2002) and Kesten (2002) on priority allocations.

mechanism is not a market mechanism.² In this paper, we test this extension of the Gale-Shapley student-optimal stable mechanism (that will be referred to as Gale-Shapley Pareto-dominant market mechanism or simply GS mechanism and will be explained in Section 5) in a controlled field experiment.

Within the management sciences, studies using systematic experimental tests include Amaldoss, Meyer, Raju and Rapoport (2000), Amaldoss and Jain (2002), Katok and Roth (2004), Ho and Weigelt (1996 and 2004), Murnighan and Roth (1977), Rapoport, Erev and Zwick (1995), Srivastava, Chakravarti and Rapoport (2000), and Zwick and Chen (1999).

Next, we discuss currently used bidding mechanisms.

3 Currently Used Course Bidding Mechanisms

Most business schools have historically used either bid-based or preference-based mechanisms for allocating courses. For example, Ross School of Business at the University of Michigan (UMBS), Kellogg Graduate School of Management at Northwestern, Johnson Graduate School of Management at Cornell, Columbia Business School, Haas School of Business at UC Berkeley, Yale School of Management, and INSEAD rely on versions of a bidding mechanism that we refer to as the **UMBS mechanism**. Some law schools too rely on bidding systems for course allocation, e.g., the School of Law at the University of Colorado at Boulder. Harvard Business School and Stanford Graduate School of Business rely on preference-based course allocations. Recently, there appears to be a shift from preference-based to bid-based allocations, UMBS being an example. In this paper, we focus on bid-based mechanisms. However, some details of the preference-based allocation systems

²It turns out that a course allocation is a market outcome (i.e., the allocation determined in a market equilibrium) in the course bidding problem if and only if it is (pairwise) stable in the induced two-sided matching problem.

are provided in Appendix A. Appendix A also describes some variants of the bid-based course allocation mechanisms.

Under the UMBS bid-based mechanism, each student is given a bid endowment $B > 0$ at the beginning of each semester. In order to keep the notation at a minimum, we assume that the bid endowment is the same for each student. This is the case at UMBS where we conducted our field experiment. Each student is asked to allocate her bid endowment among the courses, and once all bids are submitted, course seats are assigned to students as follows:

1. All bids for all courses and all students are ordered in a *single* list from highest to smallest. A tie-breaking lottery is used to determine the relative ordering of two bids of the same size. Thus, if each student i , places k_i bids of exactly x points, the tie-breaking lottery determines the order of these $\sum_i k_i$ bids.
2. Each bid is considered one at a time following the order in the list. When it is the turn of bid b_{ic} of student i for course c , the bid is **successful** if (a) course c still has unfilled seats, (b) student i still has unfilled slots in her schedule, and (c) course c does not conflict with any of the courses that are assigned to student i so far. If the bid is successful, then student i is assigned a seat at course c (i.e. the bid is honored) and the process proceeds with the next bid in the list. Otherwise student i is declined a seat at course c and the process continues with the next bid in the list.
3. When all bids are handled, a *schedule* is obtained for each student and a *course allocation* is hence obtained.

Market clearing bid or price for each course is the lowest successful bid if all course seats are allocated, and zero otherwise.

Bids have two roles under the UMBS mechanism:

1. Bids are used to infer student preferences over the courses. Consider the following statement from the guidelines for Allocation of Places in Oversubscribed Courses and Sections at the School of Law, University of Colorado at Boulder:

Each student has 100 bidding points for each semester. You can put all your points in one course, section or seminar, or you can allocate points among several. By this means, you express the strength of your preferences.

2. Bids are also used to determine who has a bigger claim on each course seat and therefore choice of a bid-vector is an important strategic tool.

The following statement from the Bidding Instructions of both Columbia Business School and Haas School of Business at UC Berkeley shows that these two roles may easily conflict.

If you do not think a course will fill up, you may bid a token point or two, saving the rest for courses you think will be harder to get into.

What happens when bids are used for both purposes (infer preferences and determine seat claims) is that students may bid a high number of points on more popular courses, and bid few points on less popular courses, even if they prefer the latter courses. However, it is easy to see that this conflict may result in efficiency loss because a student may be declined a seat at one or more of her preferred courses, despite “clearing the market” (i.e. although her bid is high enough), simply because she clears the market in too many other less preferred courses for which she has submitted higher bids.

Sönmez and Ünver (2003) shows that efficiency may be lost even if the students are expected utility maximizers, and therefore there is no reason to expect that such efficiency loss is a rare

event, or a mistake. Their Example 2 about this point is reproduced in Appendix B. Sönmez and Ünver (2003) also describe how the above mentioned efficiency loss can be avoided. The key is “separating” the two roles of the bids by simply asking students to submit their preference ranking of the courses in addition to bidding for the courses. In this way the registrar’s office no longer needs to “guess” what student preferences are. Once the bids and preference ranks are obtained, the *Gale-Shapley Pareto-dominant market mechanism*, may be used for improved efficiency, which we describe in Section 5.

We next describe in detail the bidding environment at the Ross School of Business at the University of Michigan where we conducted our field study.

4 Course Bidding at UMBS: Field Study

The question now is whether the potential efficiency loss under the UMBS mechanism occurs in practice and how prevalent it is. At UMBS, prior to, and during the bidding period, students review course descriptions and time schedules on the UMBS intranet. Professors may also make available syllabi or additional information on the course. Once the bidding period begins, students can visit a webpage within the UMBS intranet which lists all courses available to them. The webpage also contains information on the bidding system, timetables, “Tips & Tricks,” rules and regulations, etc. In addition, it has information on previous market clearing prices for all the courses.

Each student is allocated 1000 bid points for the semester. On the webpage, each course has a place to enter a bid value. As a student bids on a course from her allocated 1000 bid points, the bid amount is deducted from the 1000 allocated points and the balance is shown. Once 1000 points have been allocated, the student is prevented from entering any more bids. Students can

adjust these bids as they wish (by deducting from one, then adding to another), until the bidding period closes.

We got permission from UMBS to collect rank data from students in addition to the bid data. So as not to contaminate the bidding process in place, we collected rank data *after* bidding was over but *before* course allocations were made. This was a (very short) one week window of time.

4.1 Students and Courses

We have a sample of $n_I = 535$ students who bid for $n_C = 135$ classes scheduled for Spring 2004 semester. Let $I = \{i_1, i_2, \dots, i_{535}\}$ be the set of students and $C = \{c_1, c_2, \dots, c_{135}\}$ be the set of classes. Each class is either the sole section of a course or one of the multiple sections of a course. Therefore, we will refer to each class as a **section**. At UMBS, there are two mini-semesters in each semester. Each mini-semester lasts about 7 weeks, and the semester lasts about 14 weeks. Sections can be scheduled for the whole semester, for the first mini-semester, or for the second mini-semester. Our section sample consists of 57 first mini-semester sections, 47 second mini-semester sections, and 31 full semester sections. Each mini-semester-long section is worth 1.5 credits and each full semester-long section is worth 3 credits.

4.2 Feasibility Conditions

In the context of course-bidding, there are feasibility conditions on individual schedules as well as feasibility conditions on the course allocation. While a student can bid for as many sections as she wishes, she can be registered in

1. no more than 9 credits (or an equivalent of 6 mini-semester-long sections) for the first mini-semester,

2. no more than 9 credits (or an equivalent of 6 mini-semester-long sections) for the second mini-semester, and
3. no more than 16.5 credits (or an equivalent of 11 mini-semester-long sections) for the whole semester.³

An additional feasibility constraint on individual schedules rules out any conflict within a schedule. Two sections **conflict** if either they are both sections of the same course or their weekly meeting times overlap. While a student can bid for two conflicting sections, she cannot be registered in both of them. In our sample, 497 pairs of sections conflict out of 9045 section pairs. We refer to any set of sections that satisfy these feasibility constraints as a **schedule**.

The last feasibility condition pertains not to individual schedules, but concerns the **course allocation**. Each section has a **capacity** which is a cap on the number of students who can be registered in the section. Let $q = (q_c)_{c \in C}$ denote the capacity vector of the sections. In our sample, the smallest capacity is 5 and the largest capacity is 430. The most common capacities are 65 and 30, applying to 58 sections and 20 sections respectively. In our sample, 35 courses received more bids than their capacities.

4.3 Bids

As we already indicated, each student is endowed with 1000 bid points that she can use to bid across desired sections. Bid points cannot be transferred between semesters and bids should be integer values. A student should bid at least 1 point in order to be registered in a section. Students can bid for as many sections as they wish including conflicting sections. Let $B = [b_{ic}]_{i \in I, c \in C}$ denote the

³When a student registers for a full semester section, she “consumes” 1.5 credits from the first mini-semester and 1.5 credits from the second mini-semester.

bid matrix. Here, b_{ic} is the submitted bid of student i for section c , and $b_{ic} = 0$ if student i did not bid for section c . There are 5665 positive bids in our sample. Many students submitted the same magnitude of bid for multiple sections. Similarly, many sections received the same magnitude of bid from multiple students. We find the most repeated bid to be ‘1’, followed by ‘100’, ‘2’, ‘50’, ‘150’, ‘200’, ‘5’, ‘10’, ‘13’, and ‘20’. The top-10 bids are used a total of 2135 times. A strict bid ordering is needed to implement the UMBS mechanism and the administrators at UMBS rely on a tie-breaking lottery for this purpose. A random real number $\phi_{ic} \in (0, 1)$ is drawn from the uniform distribution for each student-section pair (i, c) , and each positive bid $b_{ic} > 0$ is modified as $b'_{ic} = b_{ic} + \phi_{ic}$ in order to break ties. Let $b'_{ic} = b_{ic}$ whenever $b_{ic} = 0$ and let $B' = [b'_{ic}]_{i \in I, c \in C}$ be the modified bid matrix. The administrators at UMBS provided us with their tie-breaking lottery draw for Spring 2004.

4.4 Preferences over Sections

We surveyed students to learn their preferences over sections. Within a few hours of the official closure of bidding, we sent each of the 535 students who had submitted bids a customized e-mail asking each student to rank the sections she bid upon. Each e-mail contained an explanation of our study and a list of all the sections that the student had bid upon. The sections were listed by descending order of bid points, but the actual bid points were left off.⁴ A permission from the Associate Dean was obtained to use his name as the sender of this e-mail in order to lend credibility and urgency to the survey. Two reminder e-mails were sent to students within the same week (see Appendix C for the original and subsequent e-mail messages). We received 489 responses out of a total of 535 students. In addition to 46 missing responses, 32 students submitted preferences with

⁴This practice is in favor of UMBS mechanism, since non-motivated students may just give ranks 1,2,3. Bids lining up with ranks, i.e., bid-monotonic preferences, favor the UMBS mechanism.

indifferences (although they were specifically asked not to). In order to measure the efficiency loss under the UMBS mechanism, we need strict preferences for all students. Using the 489 responses, we constructed a strict preference ranking $P = (P_i)_{i \in I}$ for all 535 students to analyze the *best case scenario* for the UMBS mechanism:

1. For each of the 32 students who indicated indifferences, we broke indifferences in favor of sections for which the student had a higher bid based on the modified bid matrix B' .
2. For each of the 46 students with missing preferences, we assumed that sections with higher bids were preferred to sections with lower bids.

Formally, for any student i with missing preferences and for any two sections c, d where student i submitted positive bids we assumed that

$$cP_id \quad \text{if and only if} \quad b'_{ic} > b'_{id}.$$

This preference construction results in the *lower bound* of the efficiency loss under the UMBS mechanism. That is the case because any efficiency loss here is an implication of students possibly preferring sections for which they have lower bids to sections for which they have higher bids. We say that a student i has **bid-monotonic** preferences if for any two sections $c, d \in C$,

$$cP_id \quad \text{if and only if} \quad b'_{ic} > b'_{id}.$$

As we have already emphasized, we assumed that each of the 46 students with missing preferences has bid-monotonic preferences. In addition among the 489 students who responded to the survey, 82 submitted bid-monotonic preferences. A vast majority of the students, 375 of them, submitted preferences that are *not* bid-monotonic, indicating the role of the strategic aspect of bidding in the UMBS process. This suggests that the efficiency loss can be quite significant under the UMBS mechanism.

4.5 Preferences over Schedules

We need a way to compare alternative course schedules for a student and determine which will be preferable to her. We take a conservative approach here, making the safest assumption.

Given her preferences over courses P_i , a student i **unambiguously prefers** a schedule S to another schedule S' if and only if

1. schedule S has at least as many credits as schedule S' does, and
2. each section in $S \setminus S'$ is strictly preferred to each section in $S' \setminus S$ based on the preference ranking P_i .

Therefore, we will only conclude that a student has an improvement in her schedule if she receives at least as many credits and also if any replacement in her schedule is a favorable one. Clearly, we will not be able to compare many pairs of schedules and in such cases we call the welfare comparison **ambiguous**.

Next we describe Gale-Shapley Pareto-dominant Market Mechanism, the best market mechanism for the course bidding problem.

5 The Gale-Shapley Pareto-dominant Market Mechanism

The efficiency of course-bidding can be improved by adopting the **Gale-Shapley Pareto-dominant market mechanism**. The following notation is useful to ease the description of this mechanism:

Given a preference ranking P_i and a subset of sections $D \subseteq C$, construct the **best schedule** $\mathcal{B}(D)$ as follows: Start by including the best section among sections in D based on the preference

ranking P_i . Next, add the second best section under P_i provided that neither the credit requirements are violated, nor is there any conflict with the section that is already included. Next, add the third best section under P_i provided that neither the credit requirements are violated nor are there any conflicts with one or more of the sections that are already included in $\mathcal{B}(D)$. Proceed in a similar way until either the credit requirements do not allow for any addition or there are no sections left with a positive bid. Define $\mathcal{B}(\emptyset) = \emptyset$. The best schedule is preferred to any other schedule. Under the Pareto-dominant market mechanism, students are asked not only to submit their bids but also their preference ranking over the courses they bid upon. Based on the student preferences $P = (P_i)_{i \in I}$, the modified bid matrix B' , and the capacity vector $q = (q_c)_{c \in C}$, the outcome of the Pareto-dominant market mechanism (which will be referred to as **GS mechanism** in short) can be obtained via the following version of the **deferred acceptance algorithm** (Gale and Shapley 1962, Kelso and Crawford 1982):

Step 1. Each student proposes to all sections in her best schedule from the set of all sections C . Each course c rejects all but the highest bidding q_c students among those who have proposed. Those who are not rejected are **kept on hold**.

In general, at

Step t . Each student proposes to all sections in her best schedule from the set of sections which have not rejected her in earlier steps. Each course c rejects all but the highest bidding q_c students among those who have proposed. Those who are not rejected are **kept on hold**.

The procedure terminates when no proposal is rejected and at this stage the course assignments are finalized by assigning each student the courses which keeps her on hold. The market clearing bid or price of each course is the lowest successful bid, if all course seats are filled, and zero, otherwise.

The course bidding market we consider consists of preferences of students over schedules and a

bid matrix of students. A course allocation and a price vector pair is a *market equilibrium* of this *course bidding market* if for each student there is no other schedule that the student prefers to her assigned schedule where she can “afford” each course in this new schedule at this price vector. The course allocation at a market equilibrium is called a *market outcome*. The price vector associated with a market outcome is a *competitive price vector*. As Sönmez and Ünver (2003) prove, (i) GS mechanism eliminates inefficiencies that result from registrar’s offices using bids as a proxy of the strength of the preferences; and (ii) GS mechanism is a market mechanism whose outcome Pareto-dominates any other market outcome.⁵ As the example in Appendix B shows, the UMBS mechanism is not a market mechanism although it is promoted as one by many business schools. That it is promoted as a market mechanism can be inferred from the following question and its answer borrowed from UMBS, Course Bidding Tips and Tricks:

Q. How do I get into a course?

A. If you bid enough points to make market clear, a seat will be reserved for you in that section of the course, up to class capacity.

We now discuss the results from the field study and see how large the loss in efficiency can be from using the UMBS versus the GS mechanism.

⁵Sönmez and Ünver (2003) also prove that for existence of market equilibria, it is sufficient for the preferences over schedules to satisfy *substitutability* (Kelso and Crawford 1982) which simply means that if two courses are both in the best schedule from a set of available courses and if one of the courses becomes unavailable, then the other one is still in the best schedule from the smaller set of available courses.

6 Results from the Field Study

We provide two sets of results. In the first set of results, we compare the efficiency of the two mechanisms for the Spring 2004 UMBS tie-breaker lottery draw and the modified bid matrix B' . In the second set of results we provide a robustness check with Monte Carlo simulation.

6.1 Analysis Using the UMBS Tie-Breaker Draw

Our analysis reveals that a potential transition to the GS mechanism is likely to result in significant efficiency improvement.

We observe that 10 first-quarter sections, 12 second-quarter sections, and 5 full semester sections filled their capacities under the GS mechanism whereas 9 first-quarter sections, 11 second-quarter sections, and 5 full semester sections filled their capacities under the UMBS mechanism.

<i>Full Sections</i>	GS	UMBS
Quarter 1	10	9
Quarter 2	12	11
Regular Semester	5	5

Table 1: Number of sections that are full under the UMBS and GS allocations using the UMBS tie-breaker lottery draw.

Out of the 489 students who responded to the survey, 456 receive the same credit load under both mechanisms. Among them

- each of 366 students is assigned the same schedule under both mechanisms,
- each of 83 students unambiguously prefers her schedule under the GS mechanism to her schedule under the UMBS mechanism,

- no student unambiguously prefers her schedule under the UMBS mechanism to her schedule under the GS mechanism, and
- the welfare comparison is ambiguous for 7 of these students.

Out of the 489 students who responded, 21 students receive more credits under the GS mechanism and

- each of 18 of them unambiguously prefers her schedule under the GS mechanism to her schedule under the UMBS mechanism, whereas
- the welfare comparison is ambiguous for the remaining 3.

Out of the 489 students who responded, 12 students receive more credits under the UMBS mechanism and

- each of 2 of them unambiguously prefers her schedule under the UMBS mechanism to her schedule under the GS mechanism, whereas
- the welfare comparison is ambiguous for the remaining 10.

So altogether, 366 of the 489 students who responded are indifferent between the two mechanisms, 101 of them unambiguously prefer the GS mechanism, 2 of them unambiguously prefer the UMBS mechanism, and a conclusion cannot be drawn for 20 students.

One last point deserves clarification: It is clear why many students prefer the GS mechanism to the UMBS mechanism. This is because they get the courses they really want and do not end up with a situation where they get more popular courses that they strategically bid more on, but like less (and get closed out of courses they like more but bid less on). What may be less clear is

<i>Valid Responses</i>	Indifferent	Prefers GS	Prefers UMBS	Ambiguous	Total
Same Credit Load	366	83	0	7	456
More Credits under GS	-	18	-	3	21
More Credits under UMBS	-	-	2	10	12
Total	366	101	2	20	489

Table 2: Among students who have responded, comparison of student preferences over the UMBS and GS allocations using the UMBS tie-breaker lottery draw.

why two students prefer the UMBS mechanism to the GS mechanism. The reason is quite simple. These two students got “lucky” under the UMBS mechanism and were assigned one or more courses despite their relatively low bids because some other students with higher bids were denied seats and were instead assigned seats at their less preferred courses (where they bid even higher). Since such “mistakes” are corrected under the GS mechanism, and the courses filled up, the two lucky students suffer a welfare loss under the GS mechanism.

6.2 Robustness Check: Simulation for the Tie-Breaker Lottery

As we have reported earlier, the indifferences are broken with a tie-breaking lottery at UMBS and we have already reported the results for the Spring 2004 lottery draw. In the following table we report the results of a robustness test where we draw random lotteries to break ties, and repeat this 1000 times. The results are virtually the same as the UMBS lottery draw.

6.3 Improvement in Ranks

We also computed the improvement in mean ranks among all courses allocated by the UMBS versus GS mechanisms for the modified bid matrix B' . In order to do this across both full semester and mini-semester courses, all full semester courses were treated as a package of two mini-semester

<i>Valid Responses</i>	Indifferent	Prefers GS	Prefers UMBS	Ambiguous	Total
Same Credit Load	367.853 (3.221)	82.590 (3.265)	0.247 (0.431)	7.642 (1.247)	458.332 (2.225)
More Credits under GS	-	16.953 (1.674)	-	1.366 (0.527)	18.319 (1.632)
More Credits under UMBS	-	-	2.217 (1.018)	10.132 (1.217)	12.349 (1.631)
Total	367.853 (3.221)	99.543 (3.189)	2.464 (1.082)	19.14 (1.804)	489

Table 3: Among students who have responded, comparison of student preferences over the UMBS and GS allocations using the Monte-Carlo simulation. Averages and standard errors (in parentheses below the averages) are reported for the sample.

courses with the same rank applying to both. We explain this with a simple example:

Example: Suppose a student bids for 13 sections and only 4 of these are full-semester sections, which are the 2nd, 3rd, 10th and 11th choices in her preference ranking. After we convert her preferences to their mini-semester equivalent, her preferences include 17 mini-sections. Each of the two mini-semester equivalents for the full-semester section is given the same preference ranking. Thus, the two mini-sections of the full-semester sections are ranked as the 2nd, 4th, 12th and 14th choices, amongst the 17 choices. Now, suppose under one of the mechanisms, she gets enrolled in 9 mini-sections although she could have registered in up to 11 mini-sections due to her credit limit – these 9 are her 1st, 2nd, 2nd, 7th, 9th, 12th, 12th, 14th, and 14th ranked mini-sections. The two unregistered slots in her maximal possible schedule are counted as if she were registered in her 18th choice, remaining unmatched. Note that not accounting for unregistered slots, i.e., slots below

credit limit, would needlessly favor a system where slots were left vacant. Thus, the mean rank of this schedule is calculated as

$$(1 + 2 + 2 + 7 + 9 + 12 + 12 + 14 + 14 + 18 + 18)/11 = 9.91.$$

■

Across all 489 students submitting a preference ranking, the average rank improvement by using GS versus UMBS is 1.1053. Across the 123 students who submitted a submitted preference ranking and were assigned different schedules under GS vs. UMBS, the average rank improvement is 4.3943.

An alternative method is comparing only the students who have the same credit load in the two schedules (GS and UMBS) and not taking into account unmatched slots in students' schedules. Across the 456 students who have the same credit load in the two schedules (and who submitted a preference ranking), the average rank improvement by using GS versus UMBS is 1.0647. Across the 90 students who submitted a preference ranking and were assigned different schedules under GS vs. UMBS with the same credit load, the average rank improvement is 5.3944.

7 Conclusion

In this paper, we draw attention to bid-based course allocation systems used in universities. Sönmez and Ünver (2003) theoretically show that bid-based allocation systems currently in use can result in allocations that are not market outcomes and result in unnecessary loss of efficiency. They propose an alternate course allocation mechanism which they show results in the Pareto-dominant market outcome. We use a controlled field experiment to test the extent of the efficiency loss created by the UMBS mechanism with respect to the best market mechanism, the one proposed by

Sönmez and Ünver. We believe that the field test carried out at UMBS shows quite clearly that the UMBS mechanism results in significant efficiency loss due to the two possibly conflicting roles of the student bids. The Gale-Shapley Pareto-dominant market mechanism has the potential to make a substantially larger proportion of students better off (approximately 20% in our study) than the UMBS mechanism which is currently in place at many schools. Thus, schools should consider adoption of the Gale-Shapley mechanism, given that the change required is relatively minor and the potential benefits are quite large.

We should mention that during the add-drop period in the UMBS system, some of the students would be able to transfer from a less preferred course that they bid high points on and got into, to a more preferred course that they bid less on and did not get assigned. However, sometimes this will not be possible, since this more preferred course will be already filled during the course allocation. Moreover, when it is possible, unnecessary add-drop creates a lot of upheaval in the registrar's office and even worse is the fact that the class takes longer to get settled, reducing the quality of education in the early weeks of the semester. Perhaps, it is due to the inefficient allocation process that UMBS allows two full weeks for add-drop in a mini-semester that is only 7 weeks long. Clearly, this needs to be changed.

References

- [1] Abdulkadiroğlu, A., P. A. Pathak and A. E. Roth, "The New York City High School Match." *American Economic Review, Papers and Proceedings*, May, 2005, forthcoming.
- [2] Abdulkadiroğlu, A., P. A. Pathak, A. E. Roth and T. Sönmez, "The Boston Public School Match." *American Economic Review, Papers and Proceedings*, May, 2005, forthcoming.

- [3] Abdulkadirođlu, A. and T. Sönmez, “House Allocation with Existing Tenants.” *Journal of Economic Theory*, 1999, 88: 233-260.
- [4] Abdulkadirođlu, A. and T. Sönmez, “School Choice: A Mechanism Design Approach.” *American Economic Review*, 2003, 93: 729-747.
- [5] Amaldoss, W., R. J. Meyer, J. S. Raju and A. Rapoport, “Collaborating to Compete.” *Marketing Science*, 2000, 19: 105-126.
- [6] Amaldoss, W. and S. Jain, “David vs. Goliath: an Analysis of Asymmetric Mixed-strategy Games and Experimental Evidence.” *Management Science*, 2002, 48: 972-991.
- [7] Balinski, M. and T. Sönmez, “A Tale of Two Mechanisms: Student Placement.” *Journal of Economic Theory*, 1999, 84: 73-94.
- [8] Chen, Y. and T. Sönmez, “Improving Efficiency of On-Campus Housing: An Experimental Study.” *American Economic Review*, 2002, 92: 1669-1686.
- [9] Chen, Y. and T. Sönmez, “School Choice: An Experimental Study.” Working paper, University of Michigan and Koç University, 2003. *Journal of Economic Theory*, forthcoming.
- [10] Ergin H., “Efficient Resource Allocation on the Basis of Priorities.” *Econometrica*, 2002, 70: 2489-2497.
- [11] Gale, D. and L. Shapley, “College Admissions and the Stability of Marriage.” *American Mathematical Monthly*, 1962, 69: 9-15.
- [12] Haruvy, E., A. E. Roth and M. U. Ünver, “Dynamics of Law Clerk Matching: An Experimental and Computational Investigations of Proposals for Reform of the Market.” Working

- paper, University of Texas at Dallas, Harvard University and Koç University, 2001. *Journal of Economic Dynamics and Control*, forthcoming.
- [13] Ho, Teck-Hua and Keith Weigelt. “Task Complexity, Equilibrium Selection, and Learning: an Experimental Study”. *Management Science*, May 1996. 42(5), 659 – 679.
- [14] Ho, T-H. and K. Weigelt, "Trust Building Among Strangers," Working Paper 2004, *Management Science*, forthcoming.
- [15] Kagel, J. H. and A. E. Roth, “The Dynamics of Reorganization in Matching Markets: A Laboratory Experiment Motivated by a Natural Experiment.” *Quarterly Journal of Economics*, 2000, 115: 201–235.
- [16] Katok, E. and A. E. Roth, “Auctions of Homogeneous Goods with Increasing Returns: Experimental Comparison of Alternative ‘Dutch’ Auctions.” *Management Science*, 2004, 50: 1044–1063.
- [17] Kelso, A. S. and V. P. Crawford, “Job Matching, Coalition Formation, and Gross Substitutes.” *Econometrica*, 1982, 50: 1483-1504.
- [18] Kesten, O., “On Two Competing Mechanisms for Priority Based Allocation Problems.” Working paper, University of Rochester, 2002. *Journal of Economic Theory*, forthcoming.
- [19] McKinney, C. N., M. Niederle and A. E. Roth, “The Collapse of a Medical Labor Clearinghouse (and Why Such Failures are Rare).” 2004. *American Economic Review*, forthcoming.
- [20] Murnighan, J. K. and A. E. Roth, “The Effects of Communication and Information Availability in an Experimental Study of a Three-Person Game.” *Management Science*, 1977, 23: 1336–1348.

- [21] Niederle, M. and A. E. Roth, "Market Culture: How Norms Governing Exploding Offers Affect Market Performance." Working paper, Stanford University and Harvard University, 2004.
- [22] Niederle, M. and A. E. Roth, "The Gastroenterology Fellowship Match: Should There Be a Match?" *American Economic Review, Papers and Proceedings*, May, 2005, forthcoming.
- [23] Rapoport, A., I. Erev and R. Zwick, "An experimental study of buyer-seller negotiation with one-sided incomplete information and time discounting." *Management Science*, 1995, 41: 377-394.
- [24] Roth, A. E., "The Evolution of Labor Market for Medical Interns and Residents: A Case Study in Game Theory." *Journal of Political Economy* 1984, 92: 991-1016.
- [25] Roth, A. E., "A Natural Experiment in the Organization of Entry-Level Labor Markets: Regional Markets for New Physicians and Surgeons in the United Kingdom." *American Economic Review*, 1991, 81: 414-440.
- [26] Roth, A. E., "The Economist as Engineer: Game Theory, Experimentation, and Computation as Tools for Design Economics." *Econometrica* 2002, 70: 1341-1378.
- [27] Roth, A. E. and E. Peranson, "The Redesign of the Matching Market for American Physicians: Some Engineering Aspects of Economic Design." *American Economic Review*, 1999, 89: 748-780.
- [28] Roth, A. E., T. Sönmez and M. U. Ünver "A Kidney Exchange Clearinghouse in New England." *American Economic Review, Papers and Proceedings*, May, 2005, forthcoming.
- [29] Roth, A. E. and M. Sotomayor, *Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis*, Cambridge: Cambridge University Press, 1990.

- [30] Sönmez, T. and M. U. Ünver, “Course Bidding at Business Schools.” Working paper, Koç University, 2003.
- [31] Ünver, M. U. “On the Survival of Some Unstable Two-Sided Matching Mechanisms.” Working paper, Koç University, 2001. *International Journal of Game Theory*, forthcoming.
- [32] Zwick, R. and X. Chen, “What Price Fairness: A Bargaining Study.” *Management Science*, 1999, 45: 804-823.

8 Appendices

A Variants of Course Allocation Mechanisms

A.1 Some Variants of UMBS Course-Bidding Mechanism

Yale School of Management:

Uses the same mechanism as the University of Michigan Business School except that students can only bid for only five courses (and the normal course load is four courses).

Columbia Business School:

- The real-life version of UMBS course-bidding mechanism is used for two rounds.
- The first round is the “main” round whereas in Round 2 students are expected to fill the gaps in their first round schedule.
- Unsuccessful bids from Round 1 are returned to students to be used in Round 2.
- Students can only bid for undersubscribed courses in Round 2.

Haas School of Business, UC Berkeley:

Uses the same two-round version as the Columbia Business School except that students cannot bid for more than a fixed number of units.

Kellogg School of Management, Northwestern University:

- The bid endowment should be used over two quarters by first year MBA students and over three quarters by second year MBA students. Points not used in first year do not carry over to the second year.
- Each quarter there are two rounds of bidding similar to the bidding at Columbia Business School, except that
 - students can bid for at most five courses (where the normal course load is four courses),
 - students are charged for the market clearing bids, not their own bids, and
 - bids from the second rounds carry over to the next quarter unless bidding is for the last quarter of the year.
- Hence bidding for the second quarter of the first year and the third quarter of the second year is analogous to course bidding at Columbia and Haas.

Princeton University:

- Undergraduate students cluster alternate courses together and strictly rank the courses within each cluster. Students will be assigned no more than one course from each cluster.
- Students allocate their bid endowment over clusters (as opposed to individual courses). The bid for each course in a cluster is equated to the bid for the cluster. Based on these bids, course allocation is implemented via a variant of UMBS course-bidding mechanism where

- the bids of a student for courses in a cluster are ordered subsequently based on the ranking within the cluster, and
- once a bid of a student is successful for a course in a cluster, her bids for all lower ranked courses in the same cluster are dropped.

A.2 Examples of Preference-based Course Allocation Mechanisms

Harvard Business School Course Allocation Mechanism (also adopted in Stanford Graduate School of Business):

- Students are strictly ordered in a single priority list with a random lottery.
- Each student submits a preference ranking of the courses.
- The assignment of the first course seat for each student is obtained with the serial dictatorship that is induced by the priority ordering of students: The first student is assigned a seat at her top choice, the next student is assigned a seat at her top choice among classes with still available seats, and so on.
- Once the assignment of the first seats are finalized (or equivalently the first cycle is completed), the assignment of second course seats are determined in a similar way using the reverse priority order, next the third course seats are determined in a similar way using the initial priority order, and so on.

B An Example When Students are Expected Utility Maximizers and Price-Takers

Example: (Sönmez and Ünver 2003) Consider a student i who shall register up to $q_I = 5$ courses and suppose there are six courses. Her utility for each individual course is given in the following table

Course	c_1	c_2	c_3	c_4	c_5	c_6
Utility	150	100	100	100	100	100

and her utility for a schedule s , consisting of no more than 5 courses, is additively-separable

$$U_i(s) = \sum_{c \in s} U_i(c).$$

Student i has a total of $B = 1001$ points to bid over courses $c_1 - c_6$ and the minimum acceptable bid is 1. Based on previous years' bid-data, student i has the following belief on the market clearing bids:

- Market clearing bid for course c_1 will be 0 with probability 1.
- Market clearing bids for the courses in $c_2 - c_6$ have independent identical cumulative distribution functions and for any of these courses c , the cdf F_c^i is strictly concave with $F_c^i(200) = 0.7$, $F_c^i(250) = 0.8$, and $F_c^i(1001) = 1$. That is, for each of the courses $c_2 - c_6$, student i believes that the market-clearing bid will be no more than 200 with 70% probability and no more than 250 with 80% probability.

Assuming that she is an expected utility maximizer, we next find the optimal bid-vector for student i : By first order necessary conditions and symmetry, student i

- shall bid 1 for course c_1 , and

- the same value for each course $c \in \{c_2, c_3, c_4, c_5, c_6\}$ for which she devotes a positive bid.

Therefore the optimal bid-vector is in the form: $b_{ic_1} = 1$, $b_{ic} = 1000/k$ for any k of courses $c_2 - c_6$. We next derive the expected utility of each such possibility.

Case 1: $b_{ic_1}^1 = 1$, $b_{ic_2}^1 = b_{ic_3}^1 = b_{ic_4}^1 = b_{ic_5}^1 = b_{ic_6}^1 = 200$.

$$\begin{aligned}
u^1 &= \Pr\{p_{c_2} \leq 200, p_{c_3} \leq 200, p_{c_4} \leq 200, p_{c_5} \leq 200, p_{c_6} \leq 200\} \times U_i(\{c_2, c_3, c_4, c_5, c_6\}) \\
&\quad + 5 \Pr\{p_{c_2} > 200, p_{c_3} \leq 200, p_{c_4} \leq 200, p_{c_5} \leq 200, p_{c_6} \leq 200\} \times U_i(\{c_1, c_3, c_4, c_5, c_6\}) \\
&\quad + 10 \Pr\{p_{c_2} > 200, p_{c_3} > 200, p_{c_4} \leq 200, p_{c_5} \leq 200, p_{c_6} \leq 200\} \times U_i(\{c_1, c_4, c_5, c_6\}) \\
&\quad + 10 \Pr\{p_{c_2} > 200, p_{c_3} > 200, p_{c_4} > 200, p_{c_5} \leq 200, p_{c_6} \leq 200\} \times U_i(\{c_1, c_5, c_6\}) \\
&\quad + 5 \Pr\{p_{c_2} > 200, p_{c_3} > 200, p_{c_4} > 200, p_{c_5} > 200, p_{c_6} \leq 200\} \times U_i(\{c_1, c_6\}) \\
&\quad + \Pr\{p_{c_2} > 200, p_{c_3} > 200, p_{c_4} > 200, p_{c_5} > 200, p_{c_6} > 200\} \times U_i(\{c_1\}) \\
&= 0.7^5 \times 500 + 5 \times 0.7^4(1 - 0.7) \times 550 + 10 \times 0.7^3(1 - 0.7)^2 \times 450 \\
&\quad + 10 \times 0.7^2(1 - 0.7)^3 \times 350 + 5 \times 0.7(1 - 0.7)^4 \times 250 + (1 - 0.7)^5 \times 150 = 474.79
\end{aligned}$$

Case 2: $b_{ic_1}^2 = 1$, $b_{ic_2}^2 = b_{ic_3}^2 = b_{ic_4}^2 = b_{ic_5}^2 = 250$, $b_{ic_6}^2 = 0$.

$$\begin{aligned}
u^2 &= \Pr\{p_{c_2} \leq 250, p_{c_3} \leq 250, p_{c_4} \leq 250, p_{c_5} \leq 250\} \times U_i(\{c_1, c_2, c_3, c_4, c_5\}) \\
&\quad + 4 \Pr\{p_{c_2} > 250, p_{c_3} \leq 250, p_{c_4} \leq 250, p_{c_5} \leq 250\} \times U_i(\{c_1, c_3, c_4, c_5\}) \\
&\quad + 6 \Pr\{p_{c_2} > 250, p_{c_3} > 250, p_{c_4} \leq 250, p_{c_5} \leq 250\} \times U_i(\{c_1, c_4, c_5\}) \\
&\quad + 4 \Pr\{p_{c_2} > 250, p_{c_3} > 250, p_{c_4} > 250, p_{c_5} \leq 250\} \times U_i(\{c_1, c_5\}) \\
&\quad + \Pr\{p_{c_2} > 250, p_{c_3} > 250, p_{c_4} > 250, p_{c_5} > 250\} \times U_i(\{c_1\}) \\
&= 0.8^4 \times 550 + 4 \times 0.8^3 \times (1 - 0.8) \times 450 + 6 \times 0.8^2 \times (1 - 0.8)^2 \times 350 \\
&\quad + 4 \times 0.8 \times (1 - 0.8)^3 \times 250 + (1 - 0.8)^4 \times 150 = 470.0
\end{aligned}$$

Since expected utility of bidding for three or less of courses $c_2 - c_6$ can be no more than $150 + 3 \times 100 = 450$, the optimal bid vector for student i is b_i^1 with an expected utility of 474.79. There are two important observations we shall make. The first one is an obvious one: The optimal bid for the most deserved course c_1 is the smallest bid violating bid-monotonicity. The second point is less obvious but more important: Under the optimal bid b_i^1 , student i is assigned the schedule $s = \{c_2, c_3, c_4, c_5, c_6\}$ with probability $0.7^5 = 0.168$. So although the bid $b_{ic_1}^1 = 1$ is high enough to claim a seat at course c_1 , since it is the lowest bid, student i is not assigned a seat in an available course under UMBS course-bidding mechanism. Therefore the outcome of UMBS course-bidding mechanism cannot be supported as a market outcome and this weakness is a direct source of efficiency loss. To summarize:

1. how much a student bids for a course under UMBS course-bidding mechanism is not necessarily a good indication of how much a student wants that course,
2. as an implication the outcome of UMBS course-bidding mechanism cannot always be supported as a market outcome, and
3. UMBS course-bidding mechanism may result in unnecessary efficiency loss due to not seeking direct information on student preferences. ■