

Course Bidding at Business Schools

Authors: Tayfun Sönmez, M. Ünver

This work is posted on [eScholarship@BC](#),
Boston College University Libraries.

Boston College Working Papers in Economics, 2005

Originally posted on: <http://ideas.repec.org/p/boc/bocoec/618.html>

Course Bidding at Business Schools*

Tayfun Sönmez

Koç University

and

M. Utku Ünver

Koç University

Abstract

Mechanisms that rely on course bidding are widely used at Business Schools in order to allocate seats at oversubscribed courses. Bids play two key roles under these mechanisms: Bids are used to infer student preferences and bids are used to determine who have bigger claims on course seats. We show that these two roles may easily conflict and preferences induced from bids may significantly differ from the true preferences. Therefore while these mechanisms are promoted as market mechanisms, they do not necessarily yield market outcomes. The two conflicting roles of bids is a potential source of efficiency loss part of which can be avoided simply by asking students to state their preferences in addition to bidding and thus “separating” the two roles of the bids. While there may be multiple market outcomes under this proposal, there is a market outcome which Pareto dominates any other market outcome.

*We would like to thank seminar participants at London Business School and SITE 2003 Workshop on Matching Markets for their insightful comments. Sönmez gratefully acknowledges the research support of KoçBank via the KoçBank scholar program and Turkish Academy of Sciences in the framework of the Young Scientist Award Program via grant TS/TÜBA-GEBİP/2002-1-19. Any errors are our own responsibility.

1 Introduction

Allocation of course seats to students is one of the major tasks of registrar's offices at most universities. Since demand exceeds supply for many courses, it is important to design mechanisms to allocate course seats equitably and efficiently. Many business and law schools rely on mechanisms based on *course bidding* to serve this purpose. The following statement is from Kellogg Course Bidding System Rules:¹

The bidding is designed to achieve an equitable and efficient allocation of seats in classes when demand exceeds supply.

While not all schools use the same version, the following simplest version captures the main features of a vast majority of these mechanisms:

1. Each student is given a positive *bid endowment* to allocate across the courses he considers taking.
2. All bids for all courses and all students are ordered in a single list and processed one at a time starting with the highest bid. When it is the turn of a bid, it is *honored* if and only if the student has not filled his schedule and the course has not filled all its seats.

When the process terminates, a schedule is obtained for each student. Similarly a market-clearing "price" is obtained for each course which is simply the lowest honored bid unless the course has empty seats and in that case the price is zero. The version we describe is closest to the version used by the University of Michigan Business School and thus we refer it as *UMBS course-bidding mechanism*. Schools that rely on this mechanism and its variants include Columbia Business School, Haas School of Business at UC Berkeley, Kellogg Graduate School of Management at Northwestern, Princeton University, and Yale School of Management.

UMBS course-bidding mechanism is inspired by the market mechanism and schools that rely on this mechanism promote it as a market mechanism. Consider the following question and its answer borrowed from University of Michigan Business School, Course Bidding Tips and Tricks:²

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.

In this paper we show that, UMBS course-bidding mechanism does not necessarily yield a market outcome and this is a potential source of efficiency loss part of which can be avoided by an

¹See http://www.kellogg.nwu.edu/script_html/CBSDEMO/cbs_demo.htm.

²See <http://webuser.bus.umich.edu/Departments/Admissions/AcademicServices/CurrentUpdates/BiddingTipsTricks.htm>.

appropriate choice of a market mechanism. While UMBS course-bidding mechanism “resembles” the market mechanism, there is one major aspect that they differ: Under UMBS course-bidding mechanism, students do not provide direct information on their preferences and consequently their schedules are determined under the implicit assumption that courses with higher bids are necessarily preferred to courses with lower bids. For example 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:³

The second rule is that places are allocated by the bidding system. 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.

The entire strategic aspect of course bidding is overseen under this interpretation of the role of the bids. While the choice of bids is clearly affected by the preferences, it is not adequate to use them as a proxy for the strength of the preferences. For example, if a student believes that the “market clearing” price of a course will be low, it is suboptimal for him to bid highly for that course regardless of how much he desires to be assigned a seat at that course. Indeed this point is often made by the registrar’s offices. The following statement appears in the Bidding Instructions of both Columbia Business School and Haas School of Business at UC Berkeley:⁴

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.

Here is the crucial mistake made under the UMBS course-bidding mechanism: Bids play two important roles under this mechanism.

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

These two roles may easily conflict: For example a student may be declined a seat at one of his favorite courses, despite clearing the market, simply because he clears the market in “too many” other less favorite courses. Indeed such bidding behavior is consistent with expected utility maximization and thus it cannot be considered to be a mistake.

Once we understand what is wrong with UMBS course-bidding mechanism, it is relatively easy “fixing” it: The key is “separating” the two roles of the bids and asking students to

³See http://www.colorado.edu/law/wait_list.html.

⁴See www.gsb.columbia.edu/students/biddinginstructionsummer.html and <http://web.haas.berkeley.edu/Registrar>.

1. submit their preferences, in addition to
2. allocating their bid endowment across the courses.

In this way the registrar’s office no longer needs to “guess” what student preferences are. While there may be several market outcomes in the context of course bidding, choosing the “right” one is easy because there is a market outcome which Pareto-dominates any other market outcome. We show this by relating course bidding to *two-sided matching markets* (Gale and Shapley [1962]). The Pareto-dominant market outcome can be obtained via an extension of the celebrated *Gale-Shapley student-proposing deferred acceptance algorithm*.

The mechanism design approach has recently been very fruitful in similar real-life resource allocation problems. Two important examples are the design of FCC spectrum auctions (see McMillan [1994], Cramton [1995], McAfee and McMillan [1996], Milgrom [2000]) and the re-design of US hospital-intern market (see Roth and Peranson [1999], Roth [2002]). This approach also has potential to influence policies on other important resource allocation problems. For example, Abdulkadiroğlu and Sönmez [2003] show how ideas in two-sided matching literature can be utilized to improve allocation of students to schools by school choice programs.⁵ This paper, to the best of our knowledge, is the first paper to approach course bidding from a mechanism design perspective.⁶ We believe this approach may be helpful in improving course-bidding mechanisms in practice.

The organization of the rest of the paper is as follows: In Section 2 we introduce the model and in Section 3 we introduce course bidding as well as UMBS course-bidding mechanism. We devote Section 4 to market equilibria and explain why UMBS course-bidding mechanism is not a market mechanism. We devote Section 5 to *Gale-Shapley Pareto-dominant market mechanism*, Section 6 to *interview-bidding*, and conclude in Section 7. Finally in the Appendix we present all proofs and describe some specific versions of UMBS course-bidding mechanism which are currently used at some leading schools.

2 Assignment of Course Seats to Students

There are a number of students each of whom should be assigned seats at a number of courses. Let $I = \{i_1, i_2, \dots, i_n\}$ denote the set of students and $C = \{c_1, c_2, \dots, c_m\}$ denote the set of courses. Each course has a maximum capacity and similarly each student has a maximum number of courses that he can take. Without loss of generality we assume that the maximum number of

⁵In Fall 2003 and in consultation with Alvin Roth, New York City Department of Education decided to use Gale-Shapley student-optimal stable mechanism for allocation of ninth graders to public schools.

⁶Prior to our paper, Brams and Kilgour [2001] study allocation of course seats to students via a mechanism which does not rely on course-bidding.

courses that each student can take is the same.⁷ Let q_I denote the maximum number of courses that can be taken by each student and let q_c denote the capacity of course c . We refer any set of at most q_I courses as a **schedule**, any schedule with q_I courses as a **full schedule**, and any schedule with less than q_I courses as an **incomplete schedule**. Note that \emptyset is also a schedule and we refer it as the **empty schedule**. Each student has strict preferences over all schedules including the empty schedule. We refer a course c to be **desirable** if the singleton $\{c\}$ is preferred to the empty schedule. Let P_i denote the strict preferences of student i over all schedules and R_i denote the induced weak preference relation.

Assigning a schedule to each student is an important task faced by the registrar's office. A **matching** is an assignment of courses to students such that

1. no student is assigned more courses than q_I , and
2. no course is assigned to more students than its capacity.

Equivalently a matching is an assignment of a schedule to each student such that no course is assigned to more students than its capacity. Given a matching μ , let μ_i denote the schedule of student i under μ and let μ_c denote the set of students enrolled in course c under μ . Different registrar's offices rely on different methods to assign course seats to students. However methods based on course bidding is commonly used at business schools and law schools in order to assure that the assignment process is fair and course seats are assigned to students who value them most.

3 Course Bidding

At the beginning of each semester, each student is given a **bid endowment** $B > 0$. In order to keep the notation at a minimum we assume that the bid endowment is the same for each student. Each student is asked to allocate his bid endowment across all courses. Let $b_i = (b_{ic_1}, b_{ic_2}, \dots, b_{ic_m})$ denote the **bid vector** of student i where

1. $b_{ic} \geq 0$ for each course c , and
2. $\sum_{c \in C} b_{ic} = B$.

Course bidding is inspired by the market mechanism and hence student bids are used

⁷It is straightforward to extend the model as well as the results

1. to the more general case where the maximum number of courses that can be taken by different students are possibly different, and
2. to the case where each student can take a maximum number of credits.

- to determine the market clearing bid for each course, and
- to determine a schedule for each student.

More specifically consider the following mechanism which can be used to determine market clearing bids as well as student schedules:

1. Order *all* bids for *all* courses and *all* students from highest to smallest in a *single* list.
2. Consider one bid at a time following the order in the list. When it is the turn of bid b_{ic} , the bid is **successful** if student i has unfilled slots in his schedule and course c has unfilled seats. 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. If the bid is unsuccessful then proceed with the next bid in the list without an assignment.
3. When all bids are handled, no student is assigned more courses than q_I and no course is assigned to more students than its capacity. Hence a matching is obtained. The **market clearing** bid of a course is the lowest successful bid in case the course is full, and zero otherwise.

Variants of this mechanism is used at many schools including University of Michigan Business School, Columbia Business School, Haas School of Business at UC Berkeley, Kellogg School of Management at Northwestern University, Princeton University, and Yale School of Management. The most basic version described above is closest to the version used at University of Michigan Business School and we refer it as **UMBS course-bidding mechanism**. While each of the above schools use their own version, the points we make in this paper carry over. See the Appendix for a description of how these versions differ. We next give a detailed example illustrating the dynamics of the UMBS course-bidding mechanism.

Example 1: There are five students $i_1 - i_5$ each of whom should take two courses and four courses $c_1 - c_4$ each of which has three seats. Each student has 1000 bid points to allocate over courses $c_1 - c_4$ and student bids are given in the following matrix:

b_{ic}	c_1	c_2	c_3	c_4
i_1	600	375	25	0
i_2	475	300	225	0
i_3	450	275	175	100
i_4	200	325	350	125
i_5	400	250	170	180

Positive bids are ordered from highest to smallest as follows:

$$b_{i_1c_1} - b_{i_2c_1} - b_{i_3c_1} - b_{i_5c_1} - b_{i_1c_2} - b_{i_4c_3} - b_{i_4c_2} - b_{i_2c_2} - b_{i_3c_2} - b_{i_5c_2} - b_{i_2c_3} - b_{i_4c_1} - b_{i_5c_4} - b_{i_3c_3} - b_{i_5c_3} - b_{i_4c_4} - b_{i_3c_4} - b_{i_1c_3}$$

We next process each bid, one at a time, starting with the highest bid:

$b_{i_1c_1} = 600$: The bid is successful, student i_1 is assigned a seat at course c_1 .

$b_{i_2c_1} = 475$: The bid is successful, student i_2 is assigned a seat at course c_1 .

$b_{i_3c_1} = 450$: The bid is successful, student i_3 is assigned a seat at course c_1 .

$b_{i_5c_1} = 400$: Course c_1 has no seats left, the bid is unsuccessful.

$b_{i_1c_2} = 375$: The bid is successful, student i_1 is assigned a seat at course c_2 .

$b_{i_4c_3} = 350$: The bid is successful, student i_4 is assigned a seat at course c_3 .

$b_{i_4c_2} = 325$: The bid is successful, student i_4 is assigned a seat at course c_2 .

$b_{i_2c_2} = 300$: The bid is successful, student i_2 is assigned a seat at course c_2 .

$b_{i_3c_2} = 275$: Course c_2 has no seats left, the bid is unsuccessful.

$b_{i_5c_2} = 250$: Course c_2 has no seats left, the bid is unsuccessful.

$b_{i_2c_3} = 225$: Student i_2 has a full schedule, the bid is unsuccessful.

$b_{i_4c_1} = 200$: Student i_4 has a full schedule and course c_1 has a full class, the bid is unsuccessful.

$b_{i_5c_4} = 180$: The bid is successful, student i_5 is assigned a seat at course c_4 .

$b_{i_3c_3} = 175$: The bid is successful, student i_3 is assigned a seat at course c_3 .

$b_{i_5c_3} = 170$: The bid is successful, student i_5 is assigned a seat at course c_3 .

$b_{i_4c_4} = 125$: Student i_4 has a full schedule, the bid is unsuccessful.

$b_{i_3c_4} = 100$: Student i_3 has a full schedule, the bid is unsuccessful.

$b_{i_1c_3} = 25$: Student i_1 has a full schedule and course c_3 has a full class, the bid is unsuccessful.

The outcome of UMBS course-bidding mechanism is

$$\begin{pmatrix} i_1 & i_2 & i_3 & i_4 & i_5 \\ c_1, c_2 & c_1, c_2 & c_1, c_3 & c_2, c_3 & c_3, c_4 \end{pmatrix}$$

with a market-clearing price-vector of $(450, 300, 170, 0)$.

Under the UMBS course-bidding mechanism, there can be two kinds of ties:

1. Bids of two or more students may be the same for a given course, and
2. a student may bid the same for two or more courses.

In practice, the first kind of ties is broken based on a previously determined lottery and the second kind of ties is broken based on the order student submits his bids. Based on this observation and in order to simplify the analysis, throughout the paper we assume that there are no ties. That is:

$$\begin{aligned} &\text{for all distinct } i, j \in I \text{ and } c \in C, && \text{if } b_{ic} = b_{jc} \text{ then } b_{ic} = b_{jc} = 0, \text{ and} \\ &\text{for all distinct } i \in I \text{ and } c, d \in C, && \text{if } b_{ic} = b_{id} \text{ then } b_{ic} = b_{id} = 0. \end{aligned}$$

4 Market Mechanism

Schools that rely on UMBS course-bidding mechanism promote it as a market mechanism. In this section we will explore to what extent this is appropriate.

In the context of course bidding, bids are the means to “buy” course seats. Therefore given a **bid matrix** $b = [b_{ic}]_{i \in I, c \in C}$ and a list of preferences $P = (P_i)_{i \in I}$, a **market equilibrium** can be defined as a pair (μ, p) where

- μ is a matching and it is interpreted as a **market outcome**, and
- $p = (p_c)_{c \in C} \in \mathbb{R}_+^m$ is a **price vector**

such that

1. for any student i and any course $c \in \mu_i$,

$$b_{ic} \geq p_c,$$

2. for any student i and any schedule $s \neq \mu_i$,

$$\text{if } b_{ic} \geq p_c \text{ for all } c \in s, \text{ then } \mu_i P_i s,$$

3. for any course c , if $|\mu_c| < q_c$ then $p_c = 0$.

Here (1) states that student i can afford any course in his schedule, (2) states that his schedule μ_i is better than any other schedule he could **afford**, and (3) states that the market-clearing price of a course is zero if it has empty seats.

We refer a mechanism to be a **market mechanism** if it always selects a market outcome.

4.1 Is UMBS Course-Bidding Mechanism a Market Mechanism?

Given that UMBS course-bidding mechanism is widely used in real-life implementation and given that it is promoted as a market mechanism, it is important to understand whether this mechanism indeed yields a market outcome. There is one major difficulty in this context: While the market equilibrium depends on bids as well as student preferences, UMBS course-bidding mechanism merely depends on bids. Business and law schools which use UMBS course-bidding mechanism implicitly assume that bids carry sufficient information to infer the student preferences and thus it is not necessary to inquire student preferences. Since higher bids are processed before lower bids, they implicitly assume that

1. for any student i and any pair of courses c, d ,

$$b_{ic} > b_{id} \quad \text{if and only if} \quad \{c\} P_i \{d\}, \quad \text{and}$$

2. (a) for any student i , any course c , and any incomplete schedule s with $c \notin s$,

$$\{c\}P_i\emptyset \quad \text{if and only if} \quad (s \cup \{c\})P_i s, \text{ and}$$

- (b) for any student i , any pair of courses c, d , and any incomplete schedule s with $c, d \notin s$,

$$\{c\}P_i\{d\} \quad \text{if and only if} \quad (s \cup \{c\})P_i(s \cup \{d\}).$$

That is,

1. whenever a student bids higher for a course c than another course d , he necessarily prefers a seat at c to a seat at d , and
2. this preference ranking is independent of the rest of his schedule.

The first assumption relates bids to preferences over courses and we refer it as **bid-monotonicity**. The second assumption relates preferences over schedules to preferences over courses and it is known as **responsiveness** (Roth [1985]) in the matching literature. We are now ready to relate UMBS course-bidding mechanism to market equilibria.

Proposition 1 *Suppose the bid matrix b and the preference profile P satisfy bid-monotonicity and responsiveness. Furthermore given b , let μ be the matching and p be the vector of market clearing bids obtained via UMBS course-bidding mechanism. Then*

1. the pair (μ, p) is a market equilibrium of the “economy” (b, P) and
2. the matching μ is the unique market outcome of the “economy” (b, P) .

Therefore if bids are monotonic and preferences are responsive, then not only UMBS course-bidding mechanism is a market mechanism but also it gives the unique market outcome. So a key issue is whether it is appropriate to assume that bids are monotonic and preferences are responsive.

4.2 Are Bids Monotonic?

It turns out that bid-monotonicity is not a realistic assumption. In order to make this point, we shall model how students choose their bids.

Most business and law schools provide data on market-clearing bids of previous years. Based on recent years’ bid-data and possibly some private information, students try to guess which market-clearing bids they will face. Strictly speaking, it is possible that a student can influence the market-clearing bids. However since there are hundreds of students in most applications, this is rather unlikely. Throughout the paper we assume that students are **price-takers** and they

do not try to influence the market clearing bids. Each student rather forms a belief on market-clearing bids based on recent years' bid-data possibly together with some private information and chooses an optimal-bid.

If a student believes that the market-clearing price of a course will be low, it is sub-optimal for him to bid highly for that course regardless of how much he desires to be assigned a seat at this course. Indeed, this point is often made by the registrar's office. This not only violates bid-monotonicity, but more importantly may result in a non-market outcome as well as in efficiency loss. The following example is built on this simple intuition.

Example 2: Consider a student i who shall register up to $q_I = 5$ courses and suppose there are six courses. His 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 his utility for a schedule 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 he 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 he 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.

5 Gale-Shapley Pareto-Dominant Market Mechanism

While UMBS course-bidding mechanism is very intuitive, it makes one crucial mistake: Bids play two possibly conflicting roles under this mechanism:

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

As Example 1 clearly shows, these two roles can easily conflict. Fortunately it is possible to “fix” this deficiency by utilizing the theory on two-sided matching markets developed by David Gale, Lloyd Shapley, Alvin Roth and their followers. The key point is “separating” the two roles of the bids. Under the proposed two-sided matching approach, students are not only asked to allocate their bid endowment over courses but also to indicate their preferences over schedules. In order to simplify the exposition, we initially assume that preferences over schedules are responsive. Recall that under responsiveness students can simply reveal their preferences over individual courses and the empty schedule. Later on we will show to what extent responsiveness can be relaxed.

We are now ready to adopt a highly influential mechanism in two-sided matching literature to course bidding.

Gale-Shapley Pareto-Dominant Market Mechanism:

1. Students are ordered with an even lottery to break ties.
2. Each student strictly ranks the courses in order to indicate his preferences. It is sufficient to rank only desirable courses.
3. Each student chooses a bid-vector.
4. Based on stated preferences, bids, and the tie-breaking lottery, a matching is obtained in several steps via the following **student-proposing deferred acceptance algorithm**.

Step 1: Each student proposes to his top q_I courses based on his stated preferences. 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 case there is a tie, the tie-breaking lottery is used to determine who is rejected and who will be kept on hold.

In general, at

Step t: Each student who is rejected from $k > 0$ courses in Step (t-1) proposes to his best remaining $q_I - k$ courses based on his stated preferences. In case less than $q_I - k$ courses remain, he proposes all remaining courses. Each course c considers the new proposals together with the proposals on hold and rejects all but the highest bidding q_c students. Those who are not rejected are kept on hold. In case there is a tie, the tie-breaking lottery is used to determine who is rejected and who will be kept on hold.

The procedure terminates when no proposal is rejected and at this stage course assignments are finalized.

Let μ^{GS} denote the outcome of Gale-Shapley Pareto-dominant market mechanism and let a price-vector p be determined as follows: For each course c with full capacity, p_c is the lowest successful bid and for each course c with empty seats, $p_c = 0$.

Let $P = (P_i)_{i \in I}$ be the profile of (true) student preferences over schedules. Under responsiveness, for each student i the preference relation P_i induces a strict ranking of all courses. We already assumed that students are price takers and thus they do not try to influence the market-clearing bids. Therefore under price-taking behavior, the stated preferences of students over individual courses are their true preferences.⁸ We are now ready to show that Gale-Shapley Pareto-dominant market mechanism is indeed a market mechanism.

Proposition 2 *Let P denote the list of responsive student preferences over schedules, b denote the bid-matrix, μ^{GS} be the outcome of Gale-Shapley Pareto-dominant market mechanism, and p be the induced price-vector. The pair (μ^{GS}, p) is a market equilibrium of the economy (b, P) provided that students are price-takers.*

It is easy to show that in general there can be several market outcomes. Consider the following simple example.

Example 3: There are two students i_1, i_2 each of whom should take one course and two courses c_1, c_2 each of which has one seat. The bid endowment of each student is 1000 and student preferences as well as bids are as follows:

⁸While natural, price-taking behavior is indeed stronger than what we need to assure that students state their preferences over courses truthfully: Even if students attempt to influence the market-clearing bids, a natural attempt under Gale-Shapley Pareto-dominant market mechanism is via carefully choosing their bids and not by misrepresenting their preferences.

$$\begin{array}{ll}
P_{i_1} : \{c_1\} - \{c_2\} - \emptyset & P_{i_2} : \{c_2\} - \{c_1\} - \emptyset \\
b_{i_1} : (200, 800) & b_{i_2} : (700, 300)
\end{array}$$

Here the pair (μ, p) and the pair (ν, q) are both market equilibria where

$$\mu = \begin{pmatrix} i_1 & i_2 \\ c_1 & c_2 \end{pmatrix}, \quad p = (200, 300) \quad \text{and} \quad \nu = \begin{pmatrix} i_1 & i_2 \\ c_2 & c_1 \end{pmatrix}, \quad q = (800, 700).$$

Nevertheless the outcome of Gale-Shapley Pareto-dominant market mechanism is the “right” one: Thanks to its direct relation with two-sided matching markets, the outcome of this mechanism Pareto dominates any other market outcome.

Proposition 3 *Let P denote the list of responsive student preferences over schedules, b denote the bid-matrix, and μ^{GS} be the outcome of Gale-Shapley Pareto-dominant market mechanism. Matching μ^{GS} Pareto-dominates any other matching μ that is a market outcome of economy (b, P) .*

5.1 Gale-Shapley Pareto-Dominant Market Mechanism and Efficiency

Replacing UMBS course-bidding mechanism with Gale-Shapley Pareto-dominant market mechanism eliminates inefficiencies that result from registrar’s offices using bids as a proxy of the strength of the preferences. However Gale-Shapley Pareto-dominant market mechanism does not eliminate all inefficiencies in general. While this mechanism Pareto dominates any other market mechanism, there may be situations where all market outcomes are Pareto inefficient. The following example, which is inspired by a similar example in Roth [1982], makes this point.⁹

Example 4: There are three students i_1, i_2, i_3 each of whom should take one course and three courses c_1, c_2, c_3 each of which has one seat. The bid endowment of each student is 1000 and student preferences as well as bids are as follows:

$$\begin{array}{lll}
P_{i_1} : \{c_1\} - \{c_2\} - \{c_3\} - \emptyset & P_{i_2} : \{c_2\} - \{c_1\} - \{c_3\} - \emptyset & P_{i_3} : \{c_1\} - \{c_3\} - \{c_2\} - \emptyset \\
b_{i_1} : (300, 500, 200) & b_{i_2} : (400, 350, 250) & b_{i_3} : (360, 310, 330)
\end{array}$$

Let (μ, p) be a market equilibrium of this economy. If $p_{c_1} \leq 300$ then both students i_1, i_3 demand a seat at course c_1 causing excess demand; hence $p_{c_1} > 300$ and the best student i_1 can hope is a seat at his second choice course c_2 . Given this, if $p_{c_2} \leq 350$ then both students i_1, i_2 demand a seat at course c_2 causing excess demand; hence $p_{c_2} > 350$. But then the best student i_2 can hope is a seat at his second choice course c_1 and if $p_{c_1} \leq 360$ then both students i_2, i_3 demand

⁹See also Balinski and Sönmez [1999], Ergin [2002], and Abdulkadiroğlu and Sönmez [2003] for similar examples in the context of school-student matching.

a seat at course c_1 causing excess demand; hence $p_{c_1} > 360$. Given $p_{c_1} > 360$ and $p_{c_2} > 350$, the matching

$$\mu = \begin{pmatrix} i_1 & i_2 & i_3 \\ c_2 & c_1 & c_3 \end{pmatrix}$$

is the only matching which can be supported as a market outcome. This matching can be supported with the price-vector $p = (400, 500, 330)$ among other price-vectors. Nevertheless this unique market outcome is Pareto dominated by the matching

$$\nu = \begin{pmatrix} i_1 & i_2 & i_3 \\ c_1 & c_2 & c_3 \end{pmatrix}.$$

5.2 To What Extent Responsiveness Assumption Can Be Relaxed?

Responsiveness is a very convenient assumption because it simplifies the task of indicating preferences over schedules to the much simpler task of indicating preferences over courses. However in practice it may be violated because of many reasons. For instance:

1. A student may wish to bid for different sections of the same course. More generally a student may bid for two courses he considers to be “substitutes” and may wish to take one or other but not both.
2. There can be additional difficulties due to timing of courses. A student may bid for two courses meeting at the same time and hence it may not be possible to assign him seats in both courses due to scheduling conflicts.

Therefore it is important to understand to what extent responsiveness assumption can be relaxed so that Gale-Shapley Pareto-dominant market mechanism is still well-defined. We need further notation in order to answer this question.

Given a preference relation P_i over schedules (not necessarily responsive) and given a subset of courses $D \subseteq C$, let $Ch_i(D)$ denote the best schedule from D . A preference relation P_i is **substitutable** (Kelso and Crawford [1982]) if for any set of courses $D \subseteq C$ and any pair of courses $c, d \in D$,

$$c, d \in Ch_i(D) \text{ implies } c \in Ch_i(D \setminus \{d\}).$$

Substitutability condition simply states 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 best schedule from the smaller set of available courses. Substitutability is a milder assumption on schedules than responsiveness and complications due to bidding for several alternate courses or courses with conflicting schedules are easily handled under substitutability. That is because, one can easily extend Gale-Shapley Pareto-dominant market mechanism when preferences are substitutable.

Gale-Shapley Pareto-Dominant Market Mechanism under Substitutable Preferences:

1. Students are ordered with an even lottery to break ties.
2. Each student strictly ranks the schedules in order to indicate his substitutable preferences.¹⁰
3. Each student chooses a bid-vector.
4. Based on stated preferences, bids and the tie-breaking lottery a matching is obtained in several steps via the following student-proposing deferred acceptance algorithm.

Step 1: Each student proposes to courses in his best schedule out of all courses. 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 case there is a tie, the tie-breaking lottery is used to determine who is rejected and who will be kept on hold.

In general, at

Step t : Each student who is rejected from one or more courses in Step $(t-1)$ proposes to courses in his best schedule out of those courses which has not rejected him. By substitutability this will include all courses for which he is on hold. Each course c considers the new proposals together with the proposals on hold and rejects all but the highest bidding q_c students. Those who are not rejected are kept on hold. In case there is a tie, the tie-breaking lottery is used to determine who is rejected and who will be kept on hold.

The procedure terminates when no proposal is rejected and at this stage course assignments are finalized.

Thanks to the corresponding results in two-sided matching markets, Proposition 2 and Proposition 3 immediately extends: The outcome of Gale-Shapley Pareto-dominant market mechanism under substitutable preferences is a market outcome and it Pareto dominates any other market outcome. In the Appendix we prove these results for this more general case with substitutable preferences.

What if preferences are not substitutable? For instance what happens if there are complementarities and a student wishes to take two courses together but does not wish to take either one in case he cannot take the other? Recent literature on related models with indivisibilities such as Gul and Stacchetti [1999], Milgrom [2000,2003] suggest that such complementarities might be bad news. Our next result is inspired by the similar negative results in these papers and it shows that the course-bidding approach for individual courses collapses unless preferences are substitutable.

¹⁰If only violation of responsiveness is due to conflicting schedules or bidding for alternate courses, simply indicating preferences over courses and indicating the constraints is sufficient.

More specifically we show that a market equilibrium may not exist unless preferences are substitutable.¹¹ Before we formally state our result, we present a detailed example which illustrates why the lack of substitutability might be bad news.

Example 5: Let $C = \{c_1, c_2, c_3, c_4, c_5\}$ be the set of courses, $q_I = 2$ for each student, and $q_c = 1$ for each course c . Suppose the preferences of student i_1 is such that

$$P_{i_1} : \{c_1, c_2\} - \{c_1, c_3\} - \{c_1, c_4\} - \{c_3, c_4\} - \{c_2, c_3\} - \{c_2, c_4\} - \{c_1\} - \{c_2\} - \{c_3\} - \{c_4\} - \emptyset - s$$

where s is any schedule that includes course c_5 . Preference relation P_{i_1} is not substitutable: $Ch_{i_1}(C) = \{c_1, c_2\}$ and yet $c_2 \notin Ch_{i_1}(C \setminus \{c_1\}) = \{c_3, c_4\}$.

We will construct a student set J , responsive preferences for each $j \in J$ and a bid matrix b such that the resulting economy does not have a market equilibrium.

Let $J = \{i_2\}$ and let P_{i_2} be any responsive preference relation where c_2, c_5, c_1 are the only desirable courses with $c_2 P_{i_2} c_5 P_{i_2} c_1$. Let $I = \{i_1, i_2\}$ be the entire set of students and let the bid matrix $b = [b_{ic}]_{i \in I, c \in C}$ be as follows:

$$b_{i_2 c_1} = b_{i_2 c_5} > b_{i_1 c_1} = b_{i_1 c_2} = b_{i_1 c_3} = b_{i_1 c_4} > b_{i_2 c_2} > b_{i_1 c_5} = b_{i_2 c_3} = b_{i_2 c_4} = 0.$$

We next show that the resulting economy does not have a market equilibrium. Suppose on the contrary, (μ, p) is a market equilibrium of the resulting economy.

First, observe that $c_5 \in \mu_{i_2}$: Since P_{i_2} is responsive, student i_2 prefers any schedule s containing course c_5 to any other schedule s' obtained from s by replacing c_5 by any other course. Moreover he is the highest bidder for course c_5 and therefore he shall be assigned a seat at course c_5 at any market outcome.

Next, observe that $c_3, c_4 \notin \mu_{i_2}$: This is simply because dropping either of these courses will improve his schedule and since he can afford the schedule μ_{i_2} , he can afford any subset of μ_{i_2} as well.

Therefore, we are left with three possibilities: $\mu_{i_2} = \{c_5\}$, or $\mu_{i_2} = \{c_1, c_5\}$, or $\mu_{i_2} = \{c_2, c_5\}$. We will show that none of the three can be the case at a market equilibrium.

Case 1. $\mu_{i_2} = \{c_5\}$: Since μ is a market outcome, $p_{c_5} \leq b_{i_2 c_5}$. If $c_1 \in \mu_{i_1}$ then $p_{c_1} \leq b_{i_1 c_1} < b_{i_2 c_1}$. Otherwise, c_1 has an empty seat and $p_{c_1} = 0$. In either case, $b_{i_2 c_1} \geq p_{c_1}$. By responsiveness we have $\{c_1, c_5\} P_{i_2} \mu_{i_2}$ contradicting (μ, p) is a market equilibrium.

Case 2. $\mu_{i_2} = \{c_1, c_5\}$: Since μ is a market outcome, $p_{c_1} \leq b_{i_2 c_1}$ and $p_{c_5} \leq b_{i_2 c_5}$. Suppose $c_2 \notin \mu_{i_1}$. Then course c_2 has an empty seat, and hence $p_{c_2} = 0$. But then student i_2 prefers replacing course c_1 with course c_2 by responsiveness contradicting that (μ, p) is a market equilibrium. Hence $c_2 \in \mu_{i_1}$. Since $c_1, c_5 \in \mu_{i_2}$, there are three possibilities:

¹¹Intuitively bidding for individual courses is not appropriate when preferences have complementarities and instead one may consider course allocation mechanisms which rely on bidding for schedules (instead of courses). University of Chicago Business School uses one such mechanism. Analysis of schedule-bidding mechanisms is very important but it is beyond the scope of our paper.

- a. $\mu_{i_1} = \{c_2\}$: In this case each of the courses c_3, c_4 has an empty seat and therefore $p_{c_3} = p_{c_4} = 0$.
- b. $\mu_{i_1} = \{c_2, c_3\}$: In this case $p_{c_3} \leq b_{i_1 c_3}$ and since course c_4 has an empty seat, $p_{c_4} = 0$.
- c. $\mu_{i_1} = \{c_2, c_4\}$: In this case $p_{c_4} \leq b_{i_1 c_4}$ and since course c_3 has an empty seat, $p_{c_3} = 0$.

In any of these cases, student i_1 not only affords the schedule $\{c_3, c_4\}$ but also prefers it to his schedule μ_{i_1} contradicting (μ, p) is a market equilibrium.

Case 3. $\mu_{i_2} = \{c_2, c_5\}$: Since μ is a market outcome, $p_{c_2} \leq b_{i_2 c_2}$. Therefore $b_{i_1 c_2} > b_{i_2 c_2}$ implies $b_{i_1 c_2} > p_{c_2}$. Moreover if $c_1 \in \mu_{i_1}$, then $p_{c_1} \leq b_{i_1 c_1}$ and otherwise course c_1 has an empty seat and $p_{c_1} = 0$. In either case student i_1 can afford a seat at course c_1 . But then, student i_1 can afford his top schedule $\{c_1, c_2\}$ contradicting (μ, p) is a market equilibrium.

We next generalize this observation.

Proposition 4 *Let C be the set of courses and suppose there is an agent i whose preferences P_i over schedules is not substitutable. If the number of courses in C is high enough, there exists a bid vector b_i for student i and a set of students J with responsive preferences such that for some bids $(b_j)_{j \in J}$ of these students there is no market equilibrium.*

6 Interview Bidding

In many business schools (such as Michigan, UCLA, Yale, etc.) while part of the interview slots are closed and candidates are invited by companies, the remaining slots are open and candidates are selected through a bidding procedure which is very much like UMBS course-bidding mechanism. There is, however, one important difference: In interview-bidding, students do not have capacity constraints and they can be scheduled as many interviews as their bids allow. A very natural question is whether this **UMBS interview-bidding mechanism** suffers the same difficulties as its course-bidding version. The answer turns out to be negative, provided that no less than q_a students bid for the interview slots of each firm a and student preferences satisfy a minimal **monotonicity** condition: Suppose for any student i , any company a , and any schedule s with $a \notin s$,

$$(s \cup \{a\}) P_i s.$$

This condition simply states that having additional interviews is good news! We are ready to present our final result.

Proposition 5 *Let P be a list of monotonic preferences over schedules and suppose the bid matrix b is such that*

1. for any distinct pair of students i, j and any company a , if $b_{ia} = b_{ja}$ then $b_{ia} = b_{ja} = 0$,¹²
and
2. for any company a , $|\{i \in I : b_{ia} > 0\}| \geq q_a$.

Then the outcome of UMBS interview-bidding mechanism gives the unique market outcome of economy (P, b) .

7 Conclusion

Mechanisms that rely on course bidding are widely used at Business Schools and Law Schools in order to allocate seats at oversubscribed courses. Bids play two important roles under these mechanisms:

1. Bids are used to infer student preferences over schedules, and
2. bids are used to determine who has a bigger claim on each seat.

We have shown that these two roles may easily conflict and the preferences induced from bids may significantly differ from the true preferences. Therefore, while these mechanisms are promoted as market mechanisms, they are not truly market mechanisms. The two conflicting roles of the bids may easily result in efficiency loss due to inadequately using bids as a proxy for the strength of the preferences. We have shown that under a “true” market mechanism the two roles of the bids shall be separated and students should state their preferences in addition to bidding over courses. In this way, registrar’s offices no longer need to “guess” student preferences and they can directly use the stated preferences. This will also give registrar’s offices a more reliable measure of underdemanded courses and in case this measure is used in policy decisions, more solid decisions can be given.¹³

One possible appeal of inferring preferences from bids is that there is a unique market outcome of the induced economy. On the contrary, once students directly submit their preferences in addition to allocating their bids, there may be several market outcomes. Fortunately there exists a market outcome which Pareto dominates any other market outcome and therefore multiplicity

¹²Recall that in practice such ties are broken with a lottery and two identical bids are treated differently. Therefore for practical purposes our assumption is without loss of generality.

¹³For example, the following statement from the Bidding Instructions at Haas School of Business, UC Berkeley shows that low bids may result in cancellation of courses:

Bidding serves three functions. First, it allows us to allocate seats fairly in oversubscribed classes. Second, it allows us to identify and cancel courses with insufficient demand. Third, . . .

of market outcomes is not a serious drawback for our proposal. It is important to emphasize that although relying on the Pareto-dominant market mechanism eliminates inefficiencies based on “miscalculation” of student preferences, it does not eliminate all inefficiencies. There is a potential conflict between Pareto efficiency and market equilibria in the context of course bidding and even the Pareto-dominant market equilibria cannot escape from “market failure.” Furthermore if student preferences do not satisfy a condition known as substitutability, then course bidding loses much of its appeal as a market equilibrium may cease to exist.

A Appendix: Variants of UMBS Course-Bidding Mechanism

University of Michigan Business School: The real-life version slightly differs from the version described in the main body of the paper. In the real-life version, students can bid for multiple sections of the same course or several courses that meet at the same time. In the real-life mechanism, once a bid of a student is successful for a course, the remaining bids of this student for any other course with such scheduling conflicts is dropped.

Yale School of Management: Uses the same mechanism with the University of Michigan Business School except students cannot bid for more than five and less than four courses (where 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 with the Columbia Business School except 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 second year.

- Each quarter there are two rounds of bidding similar to the bidding at Columbia Business School, except
 - 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, his bids for all lower ranked courses in the same cluster are dropped.

B Appendix: Proofs of Propositions

Proof of Proposition 1: Suppose the bid matrix b and the preference profile P satisfy bid-monotonicity and responsiveness. Furthermore given b , let μ be the matching and p be the vector of market clearing bids obtained via UMBS course-bidding mechanism.

1. We first show that the pair (μ, p) is a market equilibrium of the economy (b, P) . Market-clearing bid of a course is the lowest successful bid in case the course is full and zero otherwise. Hence (a) for any student i and any course $c \in \mu_i$, we have $b_{ic} \geq p_c$, and (b) for any course c with $|\mu_c| < q_c$ we have $p_c = 0$. All that remains is showing each student is assigned the best schedule he can afford. Take any student i . By construction of matching μ via UMBS course-bidding mechanism, if there is a course c such that $b_{ic} \geq p_c$ and yet $c \notin \mu_i$ then not only μ_i is a full schedule but also $b_{id} > b_{ic}$ for any course $d \in \mu_i$. Therefore if there is any course c student i affords but not assigned, then he has a full schedule and course c is

worse than any of the courses he is assigned by bid-monotonicity. Therefore responsiveness implies, for any student i and any schedule $s \neq \mu_i$,

$$\text{if } b_{ic} \geq p_c \text{ for all } c \in s, \text{ then } \mu_i P_i s$$

showing that (μ, p) is a market equilibrium.

2. We next show that μ is the only market outcome of the economy (b, P) . Suppose not and let (ν, r) be a market equilibrium where $\nu \neq \mu$. As in UMBS course-bidding mechanism, order all bids for all courses in a single list starting with the highest bid. Let b_{ic} be the highest bid such that $c \in \mu_i \cup \nu_i$ and yet $c \notin \mu_i \cap \nu_i$. Loosely speaking, b_{ic} is the highest bid that is “accommodated” under one of the matchings μ, ν but not both. Since $\nu \neq \mu$, such a bid exists. We have two cases to consider.

Case 1: $c \in \nu_i, c \notin \mu_i$ (i.e. while b_{ic} is an unsuccessful bid under the UMBS course-bidding mechanism, student i is assigned a seat at course c under ν .)

Why was bid b_{ic} unsuccessful under the UMBS course-bidding mechanism at the first place? There are two possibilities.

- (a) Student i filled his capacity under μ : Since b_{ic} is the highest bid such that $c \in \mu_i \cup \nu_i$ and yet $c \notin \mu_i \cap \nu_i$, we have $\mu_i \subseteq \nu_i$. But since $c \in \nu_i, c \notin \mu_i$ we have $|\nu_i| > |\mu_i|$. This contradicts student i filled his capacity by bid b_{ic} .
- (b) Course c filled its capacity under μ : The argument is very similar to argument in part (a). Since b_{ic} is the highest bid such that $c \in \mu_i \cup \nu_i$ and yet $c \notin \mu_i \cap \nu_i$, each student with a successful bid for c under UMBS course-bidding mechanism is assigned a seat at course c under ν as well. But in addition student i is also assigned a seat at course c under ν contradicting course c filled its capacity under μ .

Case 2: $c \notin \nu_i, c \in \mu_i$ (i.e. b_{ic} is a successful bid under the UMBS course-bidding mechanism but student i is not assigned a seat at course c under ν .)

Again there are two possibilities.

- (a) $r_c \leq b_{ic}$: Since bids are monotonic and since b_{ic} is the highest bid such that $c \in \mu_i \cup \nu_i$ and yet $c \notin \mu_i \cap \nu_i$, there is no course d such that $d \in \nu_i, d \notin \mu_i$, and $d P_i c$. Therefore either student i has an incomplete schedule under ν or there is a course e such that $e \in \nu_i, e \notin \mu_i$, and $c P_i e$. Since $r_c \leq b_{ic}$, student i can afford a seat at course c and therefore in either case ν_i cannot be the best schedule by responsiveness: If ν_i is an incomplete schedule then $(\nu_i \cup \{c\}) P_i \nu_i$ and if there is a course e such that $e \in \nu_i, e \notin \mu_i$, and $c P_i e$ then $[(\nu_i \setminus \{e\}) \cup \{c\}] P_i \nu_i$ both contradicting (ν, r) is a market equilibrium.

- (b) $r_c > b_{ic}$: No agent j with $b_{jc} < r_c$ is assigned a seat at course c under ν by definition of a market equilibrium. Furthermore any agent k whose bid b_{kc} for course c is higher than b_{ic} and who is assigned a seat at course c under μ is assigned a seat at course c under ν as well. That is because, b_{ic} is the highest bid such that $c \in \mu_i \cup \nu_i$ and yet $c \notin \mu_i \cap \nu_i$. Therefore $c \notin \nu_i$ implies course c fills strictly more seats under μ than under ν . Hence course c does not have a full class under ν contradicting $r_c > b_{ic} \geq 0$.

Therefore for any market equilibrium (ν, r) , we have $\nu = \mu$ completing the proof. \diamond

Course Bidding and Two-Sided Matching Markets: We next relate course bidding with two-sided matching markets in order to prove Propositions 2 and 3.

Let I be the set of students, C be the set of courses, q_I be the maximum number of courses each student can take, $q_C = (q_c)_{c \in C}$ be the list of course capacities, and $b = [b_{ic}]_{i \in I, c \in C}$ be the bid matrix. Let $P_I = (P_i)_{i \in I}$ be the list of student preferences over schedules and suppose preferences are substitutable. We simply refer each six-tuple (I, C, q_I, q_C, P_I, b) as a **problem**.

Given a problem, construct a **two-sided matching market** as follows: In addition to students who have preferences over schedules (i.e. sets of courses of size at most q_I), pretend as if each course c is also an agent who has strict preferences P_c over groups of students of size at most q_c . Furthermore suppose that preferences of courses are responsive and based on student bids. That is, for each college c ,

1. for any pair of students i, j , $\{i\}P_c\{j\}$ if and only if $b_{ic} > b_{jc}$,
2. for any student i , and any group of students J with $|J| < q_c$, $i \notin J$,

$$(J \cup \{i\})P_c J,$$

3. for any pair of students i, j , and any group of students J with $i, j \notin J$ as well as $|J| < q_c$,

$$(J \cup \{i\})P_c(J \cup \{j\}) \text{ if and only if } \{i\}P_c\{j\}.$$

Let $P_C = (P_c)_{c \in C}$ be the list of course preferences. Given a problem (I, C, q_I, q_C, P_I, b) we refer the six-tuple $(I, C, q_I, q_C, P_I, P_C)$ as an **induced two-sided matching market**.

For a problem, the central concept is a market equilibrium. For a two-sided matching market the central concept is **pairwise stability**: A matching μ is pairwise stable if there is no unmatched student-course pair (i, c) such that

1. (a) student i has an incomplete schedule and $(\mu_i \cup \{c\})P_i\mu_i$ or
- (b) student i has a course d in his schedule such that $[(\mu_i \setminus \{d\}) \cup \{c\}]P_i\mu_i$

and

- (a) course c has an empty slot under μ or
- (b) course c has a student j in its class such that $[(\mu_c \setminus \{j\}) \cup \{i\}]P_c\mu_c$.

The following well-known result is due to Blair [1988].

Proposition 6 *Suppose both students and courses have substitutable preferences over other side of the market. Then*

1. *student-proposing deferred acceptance algorithm yields a pairwise stable matching, and*
2. *this pairwise stable matching is at least as good as any pairwise stable matching for any student.*

Proposition 6 together with the following lemma will be key to prove Propositions 2 and 3.

Lemma 1 *Let (I, C, q_I, q_C, P_I, b) be a problem and $(I, C, q_I, q_C, P_I, P_C)$ be any of its induced two-sided matching markets. A matching μ is a market outcome of the problem (I, C, q_I, q_C, P_I, b) if and only if it is a pairwise stable matching of the two-sided matching market $(I, C, q_I, q_C, P_I, P_C)$.*

Proof of Lemma 1: Let (μ, p) be a market equilibrium of the problem (I, C, q_I, q_C, P_I, b) and suppose μ is not pairwise stable for the two-sided matching market $(I, C, q_I, q_C, P_I, P_C)$. There are four possibilities.

Case 1: There exists an unmatched student-course pair (i, c) such that

- student i has an incomplete schedule and $(\mu_i \cup \{c\})P_i\mu_i$, and
- course c has an empty slot.

Since c has an empty slot, $p_c = 0$. But then whenever student affords schedule μ_i he can afford schedule $s = \mu_i \cup \{c\}$ as well and hence $sP_i\mu_i$ for an affordable schedule s contradicting (μ, p) is a market equilibrium.

Case 2: There exists an unmatched student-course pair (i, c) such that

- student i has a course d in his schedule such that $[(\mu_i \setminus \{d\}) \cup \{c\}]P_i\mu_i$, and
- course c has an empty slot.

Since student i can afford schedule μ_i , he can afford schedule $s = \mu_i \setminus \{d\}$ as well. Moreover since c has an empty slot, $p_c = 0$ and hence he can also afford schedule $s' = s \cup \{c\} = [(\mu_i \setminus \{d\}) \cup \{c\}]$. Therefore $s'P_i\mu_i$ for an affordable schedule s' contradicting (μ, p) is a market equilibrium.

Case 3: There exists an unmatched student-course pair (i, c) such that

- student i has an incomplete schedule and $(\mu_i \cup \{c\})P_i\mu_i$, and
- course c has a student j in its class such that $[(\mu_c \setminus \{j\}) \cup \{i\}]P_c\mu_c$.

Since $|\mu_i| < q_I$, we have $|(\mu_i \cup \{c\})| \leq q_I$ and therefore $s = \mu_i \cup \{c\}$ is a schedule. Moreover (μ, p) being a market outcome with $c \in \mu_j$ and $[(\mu_c \setminus \{j\}) \cup \{i\}]P_c\mu_c$ imply $b_{ic} \geq b_{jc} \geq p_c$ and therefore since student i can afford μ_i , he can afford $s = \mu_i \cup \{c\}$ as well. Hence $sP_i\mu_i$ for an affordable schedule s contradicting (μ, p) is a market equilibrium.

Case 4: There exists an unmatched student-course pair (i, c) such that

- student i has a course d in his schedule such that $[(\mu_i \setminus \{d\}) \cup \{c\}]P_i\mu_i$, and
- course c has a student j in its class such that $[(\mu_c \setminus \{j\}) \cup \{i\}]P_c\mu_c$.

Since (μ, p) is a market outcome with $c \in \mu_j$, $[(\mu_c \setminus \{j\}) \cup \{i\}]P_c\mu_c$ implies $b_{ic} \geq b_{jc} \geq p_c$ and therefore student i can afford a seat at course c . Moreover since he can afford schedule μ_i , he can afford schedule $s = \mu_i \setminus \{d\}$ as well. Therefore he can also afford schedule $s' = s \cup \{c\} = [(\mu_i \setminus \{d\}) \cup \{c\}]$ and hence $s'P_i\mu_i$ for an affordable schedule s' contradicting (μ, p) is a market equilibrium.

These four cases exhaust all possibilities and hence μ shall be pairwise stable for the two-sided matching market $(I, C, q_I, q_C, P_I, P_C)$.

Conversely let μ be a pairwise stable matching for the two-sided matching market $(I, C, q_I, q_C, P_I, P_C)$. Construct the price vector $p = (p_c)_{c \in C}$ as follows:

1. If c has a full class then $p_c = b_{ic}$ where student i is the least desirable student who is assigned a seat at course c under μ .
2. If c has an empty slot then $p_c = 0$.

We will show that (μ, p) is a market equilibrium of the problem (I, C, q_I, q_C, P_I, b) :

1. By construction, $b_{ic} \geq p_c$ for any student i and any course $c \in \mu_i$.
2. Again by construction, if $|\mu_c| < q_c$ then $p_c = 0$.
3. Finally suppose there exists a student i and a schedule $s \neq \mu_i$ that he could afford such that $sR_i\mu_i$. Since preferences are strict, $sP_i\mu_i$ and therefore there is a course c student i could afford such that $c \in s$, $c \notin \mu_i$, and either

- student i has an incomplete schedule μ_i with $(\mu_i \cup \{c\})P_i\mu_i$, or
- there is a course $d \in \mu_i$ such that $[(\mu_i \setminus \{d\}) \cup \{c\}]P_i\mu_i$.

Moreover since student i can afford a seat at course c either

- course c has an empty seat under μ or
- there exists a student $j \in \mu_c$ such that $[(\mu_c \setminus \{j\}) \cup \{i\}]P_c\mu_c$.

Existence of the pair (i, c) contradicts pairwise stability of matching μ and therefore for any schedule $s \neq \mu_i$ student i can afford, $\mu_i P_i s$.

Hence (μ, p) is a market equilibrium. ◇

Proof of Proposition 2 and Proposition 3: We prove the stronger versions of the propositions for substitutable student preferences. Let I be the set of students, C be the set of courses, q_I be the maximum number of courses each student can take, $q_C = (q_c)_{c \in C}$ be the list of course capacities, $b = [b_{ic}]_{i \in I, c \in C}$ be the bid matrix and $P_I = (P_i)_{i \in I}$ be the list of substitutable student preferences. Let μ^{GS} be the outcome of Gale-Shapley Pareto-dominant market mechanism. Since students are price-takers, they will truthfully reveal their preferences. Given the problem (I, C, q_I, q_C, P_I, b) , let $(I, C, q_I, q_C, P_I, P_C)$ be an induced two-sided matching market. By Proposition 6, μ^{GS} is a pairwise stable matching for the two-sided matching market $(I, C, q_I, q_C, P_I, P_C)$ and it is at least as good as any pairwise stable matching for any student. Therefore by Lemma 1, μ^{GS} is a market outcome for the problem (I, C, q_I, q_C, P_I, b) and it Pareto dominates any other market outcome. ◇

Proof of Proposition 4: Let $C = \{c_1, \dots, c_m\}$ be the set of courses, $q_C = (q_{c_1}, q_{c_2}, \dots, q_{c_m})$ be the vector of course capacities and q_I be the maximum number of courses each student can take. Suppose there is a student i whose preferences are not substitutable. Relabel the students so that i_1 is this student. Since P_{i_1} is not substitutable, for some $C' \subseteq C$ there are two distinct courses – without loss of generality – $c_1, c_2 \in Ch_{i_1}(C')$ such that $c_2 \notin Ch_{i_1}(C' \setminus \{c_1\})$. We will construct a set of students J , a bid vector b and a list of responsive preferences $P_J = (P_i)_{i \in J}$ such that the resulting economy has no market equilibrium.

Let $I = J \cup \{i_1\}$ denote the set of all students. For each course $c \in C$, define

$$\begin{aligned} J(c) &= \{i \in I \setminus \{i_1, i_2\} : b_{ic} > \max\{b_{i_1c}, b_{i_2c}\}\} \quad \text{and} \\ K(c) &= \{i \in I \setminus \{i_1, i_2\} : c \in Ch_i(C)\}. \end{aligned}$$

That is, $J(c)$ is the set of students each of whom bids more than students i_1, i_2 for course c , and $K(c)$ is the set of students other than i_1, i_2 each of whom has course c in his best schedule. Also define

$$C^* = Ch_{i_1}(C') \cup Ch_{i_1}(C' \setminus \{c_1\}).$$

Note that

$$c_1, c_2 \in C^* \quad \text{and} \quad Ch_{i_1}(C^*) = Ch_{i_1}(C').$$

Relabel courses so that

$$C^* \cap \{c_3, c_4, \dots, c_{q_I+1}\} = \emptyset.$$

This can be done, provided that the number of courses is high enough. Construct the set of students J , the bid matrix b and the list of responsive preferences $P_J = (P_i)_{i \in J}$ such that:

1. $b_{i_2c} < b_{i_1c}$ for all $c \in C^* \setminus \{c_1\}$,
2. $b_{i_1c} < b_{i_2c}$ for all $c \in \{c_1, c_3, c_4, \dots, c_{q_I+1}\}$,
3. there is no student $i \in J$ such that $b_{i_1c_1} < b_{i_2c_1} < b_{i_2c_2}$ or $b_{i_2c_2} < b_{i_1c_2} < b_{i_1c_1}$,
4. $K(c) = J(c)$ for all $c \in \{c_1, c_2, c_3, \dots, c_{q_I+1}\} \cup C^*$,
5. $|K(c)| = |J(c)| = q_c - 1$ for all $c \in \{c_1, c_2, c_3, \dots, c_{q_I+1}\} \cup C^*$,
6. courses $c_1, c_2, \dots, c_{q_I+1}$ are the only desirable courses for i_2 with

$$\{c_2\} P_{i_2} \{c_3\} P_{i_2} \{c_4\} P_{i_2} \dots P_{i_2} \{c_{q_I+1}\} P_{i_2} \{c_1\}, \text{ and}$$

7. $|J(c) \cap K(c)| \geq q_c$ for all $c \notin \{c_1, c_2, c_3, \dots, c_{q_I+1}\} \cup C^*$.

We will show that there is no market equilibrium of the resulting economy. On the contrary, suppose (μ, p) is a market equilibrium.

Claim 1: For all $c \in \{c_1, c_2, c_3, \dots, c_{q_I+1}\} \cup C^*$ and for all $i \in J(c)$, we have $c \in \mu_i$.

Proof of Claim 1: Suppose that there is a student $i \in J(c)$ such that $c \in \{c_1, c_2, c_3, \dots, c_{q_I+1}\} \cup C^*$ and yet $c \notin \mu_i$. By Condition (4), $i \in K(c)$. There are two possible cases:

Case 1. $p_c \leq b_{i_2c}$: By responsiveness and Condition (4), $c \in Ch_i(\mu_i \cup \{c\})$. Moreover student i can afford the schedule $s = Ch_i(\mu_i \cup \{c\})$ and therefore $s P_i \mu_i$ for an affordable schedule s contradicting (μ, p) is a market equilibrium.

Case 2. $p_c > b_{i_2c}$: Since $i \in J(c)$, we have $b_{i_1c} < b_{i_2c} < p_c$ and $b_{i_2c} < b_{i_1c} < p_c$. Therefore by Condition (5), no more than $q_c - 2$ students can afford a seat at course c and hence course c has an empty seat contradicting $p_c > b_{i_2c}$. ♠

Claim 2: $\{c_3, c_4, \dots, c_{q_I+1}\} \subseteq \mu_{i_2}$.

Proof of Claim 2: Suppose that there is a course $c \in \{c_3, c_4, \dots, c_{q_I+1}\}$ such that $c \notin \mu_{i_2}$. By responsiveness and Condition (6), $c \in Ch_{i_2}(\mu_{i_2} \cup \{c\})$. Therefore since (μ, p) is a market equilibrium, $p_c > b_{i_2c}$. But then the definition of $J(c)$ together with Conditions (2), (5) imply only $q_c - 1$ students can afford a seat at course c and therefore course c has an empty seat contradicting $p_c > b_{i_2c}$. ♠

Claim 3: $\mu_{i_1} \subseteq C^*$.

Proof of Claim 3: Suppose that there is a course $c \in \mu_{i_1}$ such that $c \in (C \setminus C^*)$. There are two possible cases:

Case 1. $c \in \{c_3, c_4, \dots, c_{q_I+1}\}$: By assumption, $c \in \mu_{i_1}$ and by Claim 2, $c \in \mu_{i_2}$. By Conditions (4), (5), there is a student $j \in J(c) \cap K(c)$ such that $c \notin \mu_j$.

Case 2. $c \notin \{c_3, c_4, \dots, c_{q_I+1}\}$: By Condition (7), there is a student $j \in J(c) \cap K(c)$ such that $c \notin \mu_j$.

In either case, (μ, p) being a market equilibrium together with $j \in J(c)$ implies $b_{jc} > b_{i_1c} \geq p_c$, and this together with $j \in K(c)$ and responsiveness of P_j implies $c \in Ch_j(\mu_j \cup \{c\})$ contradicting (μ, p) is a market equilibrium. \spadesuit

We now have the machinery to execute the final part of the proof. Since only courses $c_1, c_2, c_3, \dots, c_{q_I+1}$ are desirable for student i_2 , Claim 2 leaves us with three possibilities: $\mu_{i_2} = \{c_3, c_4, \dots, c_{q_I+1}\}$, or $\mu_{i_2} = \{c_1, c_3, c_4, \dots, c_{q_I+1}\}$, or $\mu_{i_2} = \{c_2, c_3, c_4, \dots, c_{q_I+1}\}$. We will show that none of the three can be the case at a market equilibrium.

Case 1. $\mu_{i_2} = \{c_3, c_4, \dots, c_{q_I+1}\}$: Since (μ, p) is a market equilibrium and since $(\mu_{i_2} \cup \{c_1\})P_{i_2}\mu_{i_2}$ by responsiveness, we have $p_{c_1} > b_{i_2c_1}$. However by Conditions (2), (3), (5), there are only $q_{c_1} - 1$ students whose bids for course c_1 are higher than the bid of student i_2 . Therefore course c_1 has an empty seat under μ contradicting $p_{c_1} > b_{i_2c_1}$.

Case 2. $\mu_{i_2} = \{c_1, c_3, c_4, \dots, c_{q_I+1}\}$: By assumption, $c_1 \in \mu_{i_2}$ and by Claim 1, each one of the $q_{c_1} - 1$ students in $J(c_1)$ is assigned a seat at course c_1 ; therefore

$$c_1 \notin \mu_{i_1}.$$

By Conditions (1), (3), (5), there are exactly q_{c_2} students, including student i_1 , whose bids for course c_2 are higher than the bid of student i_2 . Therefore, since $[(\mu_{i_2} \setminus \{c_1\}) \cup \{c_2\}]P_{i_2}\mu_{i_2}$ by responsiveness, each one of these students should be assigned a seat at course c_2 for otherwise $p_{c_2} = 0$ and student i_2 affords the better schedule $[(\mu_{i_2} \setminus \{c_1\}) \cup \{c_2\}]$. Hence

$$c_2 \in \mu_{i_1}.$$

By Conditions (1), (5) exactly $q_c - 1$ students bid more than student i_1 for each course $c \in Ch_{i_1}(C' \setminus \{c_1\}) \subseteq C^* \setminus \{c_1\}$ and since (μ, p) is a market equilibrium, student i_1 can afford the schedule $Ch_{i_1}(C' \setminus \{c_1\})$. Moreover by Claim 3, $\mu_{i_1} \subseteq C^* \subseteq C'$ and we have already shown that $c_1 \notin \mu_{i_1}$. Therefore $\mu_{i_1} = Ch_{i_1}(C' \setminus \{c_1\})$. However the preferences of student i_1 are not substitutable and in particular $c_2 \notin Ch_{i_1}(C' \setminus \{c_1\})$ and therefore $c_2 \notin \mu_{i_1}$ directly contradicting $c_2 \in \mu_{i_1}$.

Case 3. $\mu_{i_2} = \{c_2, c_3, c_4, \dots, c_{q_I+1}\}$: By assumption, $c_2 \in \mu_{i_2}$ and by Claim 1, each one of the $q_{c_2} - 1$ students in $J(c_2)$ is assigned a seat at course c_2 ; therefore $c_2 \notin \mu_{i_1}$. Since $c_2 \in Ch_{i_1}(C')$,

$$\mu_{i_1} \neq Ch_{i_1}(C').$$

Consider course c_1 . While $b_{i_2c_1} > b_{i_1c_1}$, by assumption $c_1 \notin \mu_{i_2}$ and by Conditions (4), (5), exactly $q_{c_1} - 1$ other students bid higher than student i_1 for course c_1 . Therefore, since (μ, p) is a market equilibrium, student i_1 can afford a seat at course c_1 . Next consider any course $c \in C^* \setminus \{c_1\}$. By Conditions (1) and (5), $q_c - 1$ students bid higher than student i_1 for each such course c . Therefore student i_1 can afford each course in C^* . Moreover by Claim 3, $\mu_{i_1} \subseteq C^*$ and therefore $\mu_{i_1} = Ch_{i_1}(C^*) = Ch_{i_1}(C')$ directly contradicting $\mu_{i_1} \neq Ch_{i_1}(C')$ and completing the proof. \diamond

Proof of Proposition 5: Let P be a list of monotonic preferences over schedules and suppose the bid matrix b is such that

1. for any distinct pair of students i, j and any company a , if $b_{ia} = b_{ja}$ then $b_{ia} = b_{ja} = 0$,¹⁴ and
2. for any company a , $|\{i \in I : b_{ia} > 0\}| \geq q_a$.

Since students do not have capacity constraints in the context of interview bidding, assignment of interview slots for two distinct firms do not interfere under UMBS interview-bidding mechanism and this mechanism simply assigns the interview slots of each company a to highest bidding q_a students. Let μ denote the outcome of UMBS interview-bidding mechanism and let $p_a > 0$ be the lowest successful bid for each company a . Let A denote the set of companies and $p = (p_a)_{a \in A}$.

We first show that (μ, p) is a market equilibrium: By construction, $b_{ia} \geq p_a$ for any student i and any company $a \in \mu_a$. Moreover by assumption $|\{i \in I : b_{ia} > 0\}| \geq q_a$ for any company a . Therefore, since students do not have capacity constraints, $|\mu_a| = q_a$ which in turn implies $p_a > 0$ for any company a . All that remains is showing that under μ each student is assigned the best schedule he can afford given p . By construction of the pair (μ, p) , each student is assigned an interview slot at each company a with $b_{ia} \geq p_a$ and this is the best schedule he can afford by monotonicity.

We next show that μ is the only market outcome: Suppose not and let (ν, r) be a market equilibrium where $\nu \neq \mu$. Since $|\mu_a| = q_a$ for each company a , there is a student i and a company a such that $a \in \mu_i$ and yet $a \notin \nu_i$. Moreover since i is one of the highest bidding q_a students for company a , either

- not all interview slots of company a are filled under ν , or
- there is a student j whose bid for company a is less than the bid of student i and yet who is assigned an interview slot with company a under ν .

In either case we have $b_{ia} > r_a$ and by monotonicity $(\nu_i \cup \{a\})P_i\nu_i$ for an affordable schedule $(\nu_i \cup \{a\})P_i\nu_i$ contradicting (ν, r) is a market equilibrium. \diamond

¹⁴Recall that in practice such ties are broken with a lottery and two identical bids are treated differently. Therefore for practical purposes our assumption is without loss of generality.

References

- [1] A. Abdulkadiroğlu and T. Sönmez, “School Choice: A Mechanism Design Approach.” *American Economic Review*, June 2003.
- [2] M. Balinski and T. Sönmez, “A Tale of Two Mechanisms: Student Placement,” *Journal of Economic Theory* 84, 73-94, 1999.
- [3] C. Blair, “The Lattice Structure of the Set of Stable Matchings with Multiple Partners,” *Mathematics of Operations Research* 13, 619-628, 1988.
- [4] S. J. Brams and D. M. Kilgour, “Competitive Fair Division,” *Journal of Political Economy* 109, 418-443, 2001.
- [5] P. Cramton, “Money Out of Thin Air: The Nationwide Narrowband PCS Auction,” *Journal of Economics and Management Strategy* 4(2), 267-343, Summer 1995.
- [6] H. Ergin, “Efficient Resource Allocation on the Basis of Priorities,” *Econometrica* 70, 2489-2497, 2002.
- [7] D. Gale and L. Shapley, “College Admissions and the Stability of Marriage,” *American Mathematical Monthly* 69, 9-15, 1962.
- [8] F. Gul and E. Stacchetti, “The English Auction with Differentiated Commodities,” *Journal of Economic Theory* 92(1), 66-95, 2000.
- [9] A. S. Kelso and V. P. Crawford, “Job Matching, Coalition Formation, and Gross Substitutes,” *Econometrica* 50, 1483-1504, 1982.
- [10] R. P. McAfee and J. McMillan, “Analyzing the Airwaves Auction,” *Journal of Economic Perspectives* 10(1), 159-175, 1996.
- [11] J. McMillan, “Selling Spectrum Rights,” *Journal of Economic Perspectives*, 8(3), 145-162, 1994.
- [12] P. Milgrom, “Putting Auction Theory to Work: The Simultaneous Ascending Auction,” *Journal of Political Economy* 108(2), 245-272, 2000.
- [13] P. Milgrom, “Matching with Contracts,” Stanford University working paper, 2003.
- [14] A. E. Roth, “The Economics of Matching: Stability and Incentives,” *Mathematics of Operations Research* 7, 617-628, 1982.

- [15] A. E. Roth, "The College Admissions Problem is not Equivalent to the Marriage Problem," *Journal of Economic Theory* 36, 277-288, 1985.
- [16] A. E. Roth, "The Economist as Engineer: Game Theory, Experimentation, and Computation as Tools for Design Economics," *Econometrica* 70, 1341-1378, 2002.
- [17] A. E. Roth and E. Peranson, "The Redesign of the Matching Market for American Physicians: Some Engineering Aspects of Economic Design," *American Economic Review* 89, 748-780, 1999.