

Changing the Course Allocation Mechanism at Wharton*

by Eric Budish[†] and Judd B. Kessler[‡]

This Draft: October 21, 2014

ABSTRACT

This paper reports on an experiment conducted at the Wharton School of the University of Pennsylvania, testing a new mechanism for matching students to schedules of courses. The experiment compared Budish's (2011) approximate competitive equilibrium from equal incomes (CEEI) to the incumbent, a fake-money auction used by Wharton and numerous other professional schools. CEEI outperformed the auction on quantitative measures of efficiency and fairness and qualitative measures of perceived strategic simplicity and student satisfaction. The experiment succeeded in the Roth (1986) sense of "whispering in the ears of princes", persuading the Wharton administration to adopt CEEI and guiding real-world implementation.

Keywords: Market Design, Experiments, Matching, Assignment, Combinatorial Allocation Problems

JEL Codes: D47 (Market Design), C9 (Experiments), C78 (Matching Theory)

* The authors thank Gerard Cachon, without whom this study would never have been possible, and The Wharton School, in particular the Course Allocation Redesign Team and Wharton Computing. We thank Adriaan Ten Kate and Xan Vongsathorn for excellent RA work. We also thank Clayton Featherstone, Emir Kamenica, Stephen Leider, Paul Milgrom, Muriel Niederle, Jesse Shapiro, and Alvin Roth, as well as seminar participants at Boston University, the Stony Brook Workshop on Experimental Game Theory, ESA Santa Cruz, University of Michigan, Stanford, Wharton, and NBER Market Design.

Disclosure: The CEEI mechanism is in the public domain. The software implementation of the CEEI mechanism was funded by and is owned by Wharton, which has branded the mechanism as "Course Match." If Wharton licenses the Course Match software to other institutions and revenues exceed a threshold, then a royalty will accrue to a market design research lab that will be overseen by Eric Budish. Wharton had no right of prior review of the results of the present study.

[†] Budish: Booth School of Business, University of Chicago, eric.budish@chicagobooth.edu

[‡] Kessler: The Wharton School, University of Pennsylvania, judd.kessler@wharton.upenn.edu

I. Introduction

Combinatorial assignment — e.g., assigning students to schedules of classes, or workers to schedules of shifts — is well known to be a difficult market design problem.¹ There are three features of the problem that together make it difficult: the goods are indivisible (e.g., seats in courses), agents have preferences defined over bundles of goods (e.g., schedules of courses), and monetary transfers are prohibited. It may be helpful to think of this problem as similar to a combinatorial auction except without money, and as similar to a school choice problem except that agents require multiple objects rather than a single school. The theory literature on this problem contains numerous impossibility theorems that show that there is no “perfect” mechanism;² the mechanisms actually used in practice have been shown to have critical flaws.³

Perhaps the most prevalent mechanism in practice is the fake-money auction, used since 1996 by the Wharton School of the University of Pennsylvania, and used at a wide variety of other universities including Berkeley (Haas), Chicago (Booth), Columbia (GSB), MIT (Sloan), Michigan (Ross), Northwestern (Kellogg), NYU (Stern), Princeton (undergrad), and Yale (SOM).⁴ The basic idea of the auction is as follows. Students are given equal budgets of an artificial currency, which they use to bid for seats in courses. If a course has q seats, the q highest bidders for that course get to take it, and pay the $q+1^{\text{st}}$ highest bid — a Vickrey-like pricing rule. From a distance, this sounds like a sensible mechanism and a creative application of auction theory — an attempt to get the efficiencies of using a market without using real money. In fact, however, the approach is conceptually flawed: because the bidding currency is fake and has no outside use, the attractive incentive properties of a real-money Vickrey auction do not translate. Students have strong incentive to strategically misreport their preferences, and equilibrium outcomes can be highly unfair and inefficient (Sönmez and Ünver, 2003, 2010).⁵ These

¹ See Roth (2008), Milgrom (2011), and Sönmez and Ünver (2011) for recent overviews of this literature.

² See Pápai (2001), Ehlers and Klaus (2003), Hatfield (2009), and Kojima (2011).

³ See Sönmez and Ünver (2003, 2010), Krishna and Ünver (2008), and Budish and Cantillon (2012).

⁴ See Sönmez and Ünver (2010) for a more complete list of schools using this mechanism and a description of the (minor) design variations across institutions. See Section 2 for more details on Wharton’s variant, which uses a Vickrey auction in an initial allocation round, and then uses double auctions in subsequent rounds.

⁵ More precisely, the auction theory on which the fake-money auction is based assumes that bidders’ preferences are quasi-linear over objects and money. This is a sensible assumption with real money that has an outside use, but not when the currency is artificial. Intuitively, someone who bids 10,000 dollars in a real-money auction and loses to someone who bids 10,001 may be disappointed, but at least they can put their money to some alternative use, whereas a student who bids 10,000 points in a fake-money auction and loses to someone who bids 10,001 may end up graduating with a large budget of useless course-auction currency. In fact this happens surprisingly often in practice, see Section 8.3 of Budish (2010).

theoretical flaws began to manifest in practice at Wharton — most colorfully presented in an article in the student newspaper entitled “Top 10 Reasons that Wharton Students Hate the Auction”.⁶ The Wharton administration grew increasingly concerned about the strategic nature of the auction, that students’ bids were not reflective of their true preferences, and that it was unfair to students who were not adept at “gaming the system”.⁷ As a result, in September 2011, Wharton convened a committee to evaluate their course allocation system.

Budish (2011) proposes a new mechanism for combinatorial assignment, called approximate competitive equilibrium from equal incomes (CEEI). It also uses an artificial currency, but it takes a “fake-money general-equilibrium theory” approach to finding prices and allocations, rather than the “real-money auction theory” approach described above. Roughly, students report their preferences over schedules, and a computer seeks competitive equilibrium prices: prices at which, when each student is assigned their most-preferred affordable bundle of courses, the market clears.⁸ This system is considerably more involved computationally than the fake-money auction — the computer has to find an approximate Kakutani fixed point rather than just calculate the q^{th} highest number in a bunch of lists — but the system also yields considerably more attractive economic properties. First, because students are assigned their most preferred affordable bundle at the realized prices, students now have a dominant strategy (approximately) to report their preferences truthfully; this contrasts with the auction in which students had to invest time and effort in estimating how others would bid and deciding how to best respond, which frequently led to costly optimization mistakes. Second, because the allocation approximates a competitive equilibrium from equal incomes, it satisfies attractive properties of efficiency and fairness. In particular, it is approximately Pareto efficient, and it minimizes “envy” (i.e., where one student prefers another student’s schedule to her own); this contrasts to the auction in which gaming leads to inefficiency (Sönmez and Ünver, 2003, 2010) and students sometimes obtain a very poor schedule (e.g., get zero of the classes they bid on) and so have a lot of envy towards other students. The authors of the present study brought Budish (2011) to the Wharton committee’s attention, and

⁶ Reason #1 was “Even with historical prices, it’s hard to know what to bid.” Reason #10 was “There must be better models out there.”

⁷ This is an issue that has arisen elsewhere, and was an impetus for changing the school choice mechanism in Boston (Pathak and Sönmez 2008).

⁸ Here is the difference between real-money auction prices and fake-money general equilibrium prices expressed mathematically. Let u denote a student’s utility function, p denote a price vector, x denote a schedule, and b denote a budget. In the auction, which assumes quasi-linear utility over goods and (fake) money, the student is allocated the bundle $\max_x u(x) - p \cdot x$. That is, the bundle that maximizes the difference between utility and expenditure. The role of the budget is it constrains what kinds of utility functions the student can express (utility for all courses sums to at most b). In Budish (2011), the student is allocated the bundle $\max_x u(x) : p \cdot x \leq b$. That is, the bundle that maximizes utility subject to the budget constraint.

proposed an experimental test of the mechanism to assess its suitability for use in practice.

Why an experiment? What might we learn from an experiment that we do not already know from theory? We had two specific issues in mind, the first concrete and the second a bit more intangible. The concrete issue was preference reporting. In Budish (2011), a student's "type" is an ordinal preference ranking over all possible schedules of courses, much as in general-equilibrium theory a household's type is an ordinal ranking over all possible consumption bundles. As is standard in mechanism design, agents are assumed to be able to "report their type" to the mechanism. But this assumption clearly strains reality: in a context such as Wharton's, there are hundreds of millions of schedules in a given semester. Can students really report their preferences with sufficient accuracy to produce the results promised by the theory analysis? Will the theoretical benefits of the new mechanism materialize when used by real market participants?

The second and more intangible issue was our sense that a work of mathematical economics was unlikely on its own to persuade Wharton's decision makers (henceforth the "Wharton deans") to change their mechanism. We thought an experimental test would be a critical step on the way from theory to practice; in the words of Roth (1986), experiments "whisper in the ears of princes" as the princes decide what policies to adopt.⁹ There were two related reasons why we thought an experiment would be useful to persuade the Wharton deans. First, we had the sense that decisions makers may want to see the mechanism at work, particularly in an environment similar to their own. While the theory in Budish (2011) promises benefits on efficiency and fairness if agents can report their preferences well enough, the deans may have wanted to see these benefits materialize with their MBA students. By way of analogy, car buyers do not rely on the manufacturing specifications to learn that the car is powerful and drives smoothly; instead, they kick the tires and take the car for a "test drive." Second, we thought decision makers might worry about some *unforeseen* issue beyond our concern about students' ability to report preferences, which the theoretical analysis missed. The first step in the FDA drug approval process is not to test whether the drug is better than alternatives (that is the last step), but simply to ensure that the drug is not harmful to humans for some unforeseen reason. Before adopting a new market design mechanism, a decision maker should worry about unintended "side effects."

These issues dictated the design of our experiment. Most centrally, we used real market participants in an environment closely tailored to the market of interest: Wharton MBA students who were asked to report their real preferences over schedules of real Wharton courses, using a realistic, professionally designed user interface. In addition, we elicited a

⁹ That experiments can be helpful to convince decision makers was suggested by Roth (2008) in his summary of the early lessons of Market Design.

broad range of quantitative and qualitative response data, allowing us to test whether Wharton students could accurately report their preferences and giving comfort to the Wharton administration that we had conducted a wide search for unintended “side effects” of the CEEI mechanism.

The response data we collected allowed us to conduct: (i) an isolated quantitative test of preference-reporting accuracy, (ii) an isolated quantitative test of the efficiency and fairness performance of the CEEI mechanism on its own (i.e. assuming perfect preference reporting), and (iii) a joint quantitative test of the mechanism and the preference-reporting accuracy. We also asked subjects a wide variety of qualitative survey questions regarding their experience of the mechanism on topics ranging from the reporting language, to their perceptions of fairness, to their satisfaction with the schedule they received, to an overall sense of whether they “liked” the mechanism. The survey also gave subjects the chance to provide free-form comments regarding the mechanism.

We briefly summarize our main results. In the joint test of the mechanism along with subjects’ preference-reporting accuracy, CEEI outperformed the Wharton Auction (WA) on each of our quantitative measures of efficiency and fairness. In addition, CEEI outperformed WA on our qualitative measures of strategic simplicity and overall student satisfaction, with most (though not all) differences statistically significant at the 5% level. We also found evidence, however, that subjects had significant difficulty with preference reporting; in addition, the survey revealed a weakness of CEEI that we had not fully anticipated, which is that subjects found it to be too much of a “black box.”

To give a sense of how preference-reporting accuracy impacts the success of CEEI compared to the WA, we highlight our fairness measure of “envy” among subjects in the experiment. Envy is defined as the event where student A prefers student B’s schedule to her own (i.e., student A envies student B). When subjects use the WA, 42% of subjects envy another subject’s schedule. When they use CEEI, the figure is significantly lower, at 31% ($p < 0.05$); this shows that CEEI reduces envy as predicted by the theory. However, if we assume that our experimental subjects’ preference reports perfectly reflect their true preferences, only 4% of subjects have any envy under CEEI. This suggests that if students were able to report their preferences perfectly under CEEI, envy would nearly be eliminated.

In the isolated test of preference-reporting accuracy, we find evidence consistent with two sources of preference-reporting mistakes: limitations of the preference-reporting language and difficulty using the preference-reporting language. In the first category, we find evidence that there were certain kinds of non-additive preferences that were important to students but that they were mechanically unable to express using the language provided. This suggests potential scope for enhancing the preference-reporting language. In the second category, we find — consistent with a popular intuition in the

matching market design literature — that students more accurately reported ordinal preference information (I like course A better than B) than cardinal preference information (*how much* I like A better than B). This result suggests that additional training on how to use the language, and more time to submit their reports (subjects had only 10 minutes in the experiment), might be valuable for practical implementation.

The experiment was successful in the Roth sense of “whispering in the ears of princes” in two regards. First, the experimental test persuaded the Wharton deans to adopt CEEI beginning in Fall 2013.¹⁰ It is worth noting that there was no singular result that convinced the Wharton deans to switch to CEEI. As will be shown in the following sections, when CEEI outperforms the WA, it often does so by a modest margin. Our impression is that it was the complete set of experimental results, as well as the fact that the experimental results were all broadly consistent with the theory, that convinced the deans to adopt CEEI over WA. Second, the experiment yielded lessons about how to improve the performance of CEEI that have guided its implementation in practice. In particular, the nature of preference-reporting mistakes and the qualitative responses about CEEI being a “black box” helped guide how the CEEI system is explained to students and how they are taught to report their preferences.¹¹

Finally, results from the first two semesters of using CEEI at Wharton demonstrate it has been a success in practice. While data limitations prevented a full empirical before-and-after welfare comparison, the quantitative data that are available show that CEEI increased equity in both total expenditure and the distribution of popular courses. Administration survey data show that CEEI increased students’ satisfaction with their assigned schedule, students’ perceptions of fairness, and students’ overall satisfaction with the course allocation system.

The remainder of this paper is organized as follows. Section 2 describes the experimental design. Section 3 presents our main results on fairness and efficiency. Section 4 analyzes preference-reporting mistakes and their lessons for implementation and future work. Section 5 reports on the qualitative response questions and the search for “side effects” of the CEEI mechanism. Section 6 reports on evidence of the success of CEEI in practice and concludes.

¹⁰ They have branded the CEEI mechanism and interface as “Course Match” and are planning to license the system to other academic institutions.

¹¹ For example, clearing prices are reported to students so the mechanism seems more transparent. Gerard Cachon, the chair of Wharton’s Course Allocation Redesign Team, writes in personal correspondence: “I have heard that this makes a difference – some students say ‘when I saw the prices, I understood why I got what I got.’”

II. Experimental Design

The experiment placed groups of subjects into a market for Wharton courses in a hypothetical semester. In the face of scarcity and competition from others in a session, subjects were asked to construct the best schedule of five courses that they could, first under CEEI or the WA and then under the other mechanism (see Appendix for study instructions). Under each mechanism, subjects were instructed to “construct your most preferred schedule” from the twenty-five course sections available (see Appendix for list of course offerings) imagining that it was their last semester at Wharton.¹²

Subjects constructed a schedule under the first mechanism. Subjects then constructed a schedule under the second mechanism, starting from scratch (i.e. starting with zero courses, just as in the first mechanism). After using each mechanism, subjects answered survey questions about the mechanism and the schedule they received. After using both mechanisms, subjects made up to nineteen binary comparisons between schedules. That is, we showed subjects two schedules and asked the subjects to pick which of the two they preferred. Subjects then answered additional survey questions about both mechanisms and wrote free-response comments.

As described in the Introduction, we had two main goals for the experimental design. First, we wanted to test whether market participants would be able to use a feasible preference-reporting language accurately enough to realize the theoretical benefits promised by Budish (2011). Second, we wanted to address potential concerns from decision makers that the theory might have missed something important for practice; that is, we wanted to look for “side effects” of the CEEI mechanism. These goals informed a number of important design decisions.

First, we used real market participants (i.e. Wharton MBA students) as subjects. 132 MBA students participated in one of eight sessions (in groups of fourteen to nineteen) in a computer lab at The Wharton School.¹³ These students were recruited with an email

¹² Gerard Cachon, the chair of Wharton’s Course Allocation Redesign Team, chose the list of courses to be representative of Spring semester offerings with a tilt towards popular classes. The twenty-five course sections and their descriptions were kept the same in all sessions, although the number of seats in each class varied by the number of subjects who showed up to each session. The twenty-five course sections were comprised of twenty-one different courses, of which four courses had two sections each. At Wharton there are around 200 course sections (from 100 to 150 courses) in a typical Spring semester.

¹³ Three pilot sessions were run with MBA students in the week preceding the experiment. During these sessions, a number of bugs in the CEEI code were identified and fixed such that the experiment would successfully solve for CEEI prices (which did not happen in the first pilot session and which took too long in the other two sessions for the experiment to be completed). After these pilot sessions, the experimental instructions were finalized (i.e., changed for clarity and length) for the eight sessions conducted the week of November 28, 2011.

sent by the Wharton Dean's office, which offered \$250 to two randomly selected subjects per session (see Appendix for recruitment email).¹⁴

Second, we had subjects use their real preferences over schedules of classes (i.e., rather than endowing them with preferences). Getting to use real market participants, and having them use the mechanisms with their real preferences, comes at a cost. Namely, we do not easily observe subjects' *true* preferences over schedules, which we need to evaluate their ability to report their preferences and to judge the success of the market design mechanisms. To solve this problem, we had subjects report preferences for schedules using binary comparisons at the end of the study. By showing subjects pairs of schedules and asking them to choose which of the two schedules they preferred, we were able to construct a dataset that we treat as subjects' true preferences over schedules. The assumption underlying this design choice is that reporting a full utility function over schedules is hard, but making binary comparisons between two schedules is easy.

The alternative to using the real preferences of Wharton MBA students would have been to use an "endowed preferences" approach in which subjects are offered monetary incentives for achieving certain outcomes in the experiment. While this technique is employed successfully in market design experiments with goals different from ours,¹⁵ it was not ideal for our setting. We were explicitly interested in testing subjects' ability to report preferences in an environment where reasonably sized problems have too many possible outcomes to endow subjects with complete rank order preferences (in our experiment there were approximately 50,000 possible schedules). To use the endowed preference approach, we would have needed to endow subjects with preferences using some shorthand and then introduced another, different shorthand to have subjects report those preferences back to us.¹⁶

¹⁴ Recruitment did not explicitly mention course allocation, and we asked students who participated not to discuss the study with others to avoid potential selection issues in later sessions. In particular, we wanted to attract student subjects who were generally representative of the Wharton MBA student body and to avoid attracting students who were disproportionately happy or unhappy with the current course auction. Subjects were representative of the Wharton student population on every dimension except race and importantly were representative with regard to attitudes toward the Wharton Auction (see Table A1 in the Appendix).

¹⁵ See a growing experimental literature in market design, which uses laboratory experiments to test between market designs or uncover behavioral deviations from theory, including Kagel and Roth 2000; Chen and Sönmez 2006; Pais and Pinter 2008; Calsamiglia, Haeringer and Klijn 2010; Kagel, Lien and Milgrom 2010; Kessler and Roth 2012; Featherstone and Niederle 2014; Featherstone and Mayefsky 2014; Fragiadakis and Troyan 2014; and Kessler and Roth 2014.

¹⁶ In addition, even if we had been able to find two sufficiently different ways to communicate rank order preferences, one might worry that our results were a function of the preferences we had chosen to give to subjects and be concerned that our results would fail to generalize to the kind of preferences real Wharton MBA students have over Wharton schedules.

Third, subjects reported preferences using a realistic user interface, akin to one that might have been used in actual implementation. The interface was designed specifically for this experiment by information technology professionals at Wharton. This allowed us to test whether market participants could accurately report their preferences using the type of preference-reporting language and interface they would face in practice. It also allowed us to look for benefits of the new mechanism in a close approximation to the setting that the deans cared about, using the particular subjects the deans cared about (i.e., letting the deans take the new mechanism for a “test drive”).

Fourth, we collected a broad range of additional non-choice data (see Appendix for the list of qualitative questions). Asking a broad range of questions allowed us to explore a number of possible “side effects” associated with CEEI as a course allocation mechanism. In particular, we asked subjects about their preferences between mechanisms, their perceptions of simplicity and ease of use, and the extent to which they had control over the outcome in the mechanism. These survey responses were supplemented by administrative data on subjects. Because we used Wharton students as subjects and had them log in using their Wharton ID number, we had demographic information allowing us to show the deans side effects results separately for particular subgroups they might be concerned about (e.g., men and women).

2.1 Wharton Bidding Points Auction

At the time of our study, the Wharton Auction (WA) that Wharton students used to select courses involved nine rounds spaced over a period of a few months with some rounds designated for second-year students only.¹⁷ To make the WA simple enough to play during the experimental session, we implemented a simplified version of the Wharton Auction for our experiment. Our version lasted four rounds instead of nine, all students could participate in all rounds, and each student was allocated the same budget of 5,000 points of artificial currency. Our WA also minimized the inequity that arises in the Wharton Auction for reasons other than the auction design per se.¹⁸

¹⁷ For instance, the auction for Fall classes began in April with rising second-years bidding for courses and ended in September, during the first week of classes.

¹⁸ Giving students equal budgets eliminated a source of inequality — and envy of others’ schedules — that arises due to variance in initial budgets at the start of a semester. In addition, limiting the auction to four rounds and allowing all students to bid in each round minimizes scope for speculation in our version of the auction. In particular, the two main sources of speculation in the WA are (1) second-year students buying up sections popular with first-year students during the rounds in which only second-year students can bid and then selling them at a premium to first-year students in later rounds and (2) students buying moderately popular sections in the early rounds and selling them at a premium in later rounds to students who did not get their first choice sections. Allowing all students to bid in the first round directly eliminates the first source of

The instructions for the WA (see Appendix for WA instructions) were relatively familiar to subjects, since all had used the real Wharton Auction to pick their courses (i.e. subjects were in the midst of using it for either their second or fourth time to select Spring semester courses). Subjects used the standard web interface of the real WA so that it would be as familiar as possible.

In the first round of the simplified WA, subjects bid points for course sections. They had five minutes to select their bids, and their bids could not sum to more than their 5,000-point budget. For a section with k seats, the k highest bidders had the course section added to their schedule and paid the $k+1^{\text{th}}$ highest bid. In rounds 2-4, subjects had two-and-a-half minutes each to select bids and asks. After each round, the market would clear such that buyers and sellers would transact course sections at the larger of the first losing bid and the highest winning ask.

2.2 Approximate Competitive Equilibrium from Equal Incomes (CEEI)

Unlike the WA, CEEI was unfamiliar to the subjects in the experiment. Consequently, instructions (see Appendix for CEEI instructions) took about 10 minutes to read, as compared to 5 minutes for the WA instructions.

CEEI has four steps: (i) students report their preferences, (ii) each student is assigned an equal budget (5,000 points in the experiment) plus a small random amount (used to break ties),¹⁹ (iii) the computer finds (approximate) market-clearing prices, (iv) each student is allocated her most preferred affordable bundle — the affordable bundle she likes best given her report in step (i) based on her budget set in step (ii) and the prices found in step (iii).²⁰

It was explained to subjects that their only responsibility in using the CEEI mechanism was to tell the computer their preferences over schedules and that the computer would compute market-clearing prices and buy them the best schedule they could afford at those

speculation. The small size of the markets and the short duration of the game make behavior consistent with the second source of speculation riskier.

¹⁹ Budish's (2011) result that prices exist for CEEI that (approximately) clear the market requires that students have non-identical budgets. The budgets can be arbitrarily close to equal but cannot be exactly equal. The intuition is that the budget inequality helps break ties. For example, suppose students A and B both place an extremely high value on course X, which has 1 available seat. If A's budget is 5000 and B's budget is 5001, then setting the price of course X to 5001 clears the market because B can afford it while A cannot. The WA breaks ties in the auction itself rather than in the budgets. If both A and B bid 5000 points for course X, then the computer randomly selects one student to transact.

²⁰ See Budish (2011) for a more complete description of how CEEI works. See Budish, Othman, and Sandholm (2010) for how to calculate the market-clearing prices in step (iii).

prices. It was explained that they did not need to think about other subjects' preferences or market clearing prices when reporting their own preferences and that they could do no better than reporting their true preferences as best they could. In particular, the instructions explained that CEEI "is designed so you do not need to think about the prices of the courses or the values that other students assign to courses. You get the best schedule possible simply by telling the computer your true values for courses".

2.3 Preference-Reporting Language

The way preferences over schedules are reported by agents in CEEI is one of the key elements in the practical implementation of the mechanism.²¹ As noted above, the CEEI mechanism requires an ordinal ranking over all feasible schedules. In practice, there are likely to be too many schedules to rank (in the experiment with twenty-five course sections there are approximately 50,000 feasible schedules; in practice at Wharton with around 200 course sections each semester there are on the order of hundreds of millions of feasible schedules). Consequently, market participants need to use a shorthand to report their preferences. In particular, we solicit a vector of information that can be translated into a rank order preference list over all possible schedules.

There were two key components of the reporting language that we implemented. First, subjects could report cardinal item values (i.e., a value for each course section) on a scale of 0 to 100. Second, subjects could report "adjustments" for any pair of course sections. Adjustments assigned an additional value, either positive or negative, to schedules that had both course sections together. Adjustments are a simple way for students to express certain kinds of substitutabilities and complementarities.²²

To calculate the rank order list of all possible permissible schedules, the CEEI system summed the cardinal values and any positive or negative adjustments. Schedules with higher sums were taken as preferred to those with lower sums.²³

²¹ See, e.g., Milgrom (2011) on the importance of preference-reporting languages in practical market design.

²² Students did not need to report schedule constraints, which were already known by the system. For example, the system prevented a student from buying two course sections that met at the same time or two different sections of the same course.

²³ Observe that cardinal item values over individual courses induce ordinal preferences over bundles of courses. For example, if a student's values for courses $\{A,B,C,D\}$ are $\{100,80,50,10\}$, and the student needs at most two courses, then her ordinal preferences over bundles are $\{A,B\} > \{A,C\} > \{B,C\} > \{A,D\} > \{A\} > \{B,D\} > \dots$. If the student had an adjustment that getting $\{B\}$ and $\{D\}$ together had an additional value of 80, then the new order would be $\{A,B\} > \{B,D\} > \{A,C\} > \{B,C\} > \{A,D\} > \{A\} > \dots$

The interface (see Figure 1) showed a list of the twenty-five course sections available to subjects. Each section had a slider, defaulted to 0, that the student could use to indicate her value for that course section. As turning a rank-order list of schedules into a vector of cardinal item values and cardinal adjustments is a non-trivial task, we gave subjects explicit instructions on how they might report these cardinal utilities. Subjects were instructed to submit a value of 100 for their favorite course section and instructed to submit a relative value (between 1 and 100) for any other course section they had interest in taking. We recommended reporting a positive value for at least twelve course sections to ensure that subjects received a complete schedule. Subjects had ten minutes to choose these values and enter any adjustments.

Figure 1: Screenshot of CEEI User Interface

Wharton UNIVERSITY of PENNSYLVANIA COURSE REGISTRATION COURSE MATCHING SYSTEM

Assign Values | My Top 10 Schedules | First Survey | Second Survey | Third Survey | Fourth Survey

MY ADJUSTMENTS

Courses	Individual Value	Combined Adjustment	Combined Value
ACCT742003, ACCT897402	64 + 27 = 91	-91	0 delete

MY VALUES

Course	Title	Instructor	Meeting	Credit	Open	Value	Apply Adjustment
ACCT742003	PROBLEMS IN FIN REPORTIN	LAMBERT R	MW 1:30 PM-3:00 PM	1.00	5	64	<input type="checkbox"/>
ACCT897402	TAXES AND BUS STRATEGY	BLOUIN J	MW 12:00 PM-1:30 PM	1.00	4	27	<input type="checkbox"/>
FNCE726003	ADVANCED CORP FINANCE	VAN WESEPE	TR 12:00 PM-1:30 PM	1.00	5	39	<input type="checkbox"/>
FNCE728003	CORPORATE VALUATION	CICHELO M	MW 3:00 PM-4:30 PM	1.00	4	27	<input type="checkbox"/>
FNCE750001	VENT CAP & FNCE INNOVAT	WESSELS D	MW 1:30 PM-3:00 PM	1.00	4	55	<input type="checkbox"/>
FNCE750002	VENT CAP & FNCE INNOVAT	WESSELS D	MW 3:00 PM-4:30 PM	1.00	4	32	<input type="checkbox"/>
FNCE891001	Corporate Restructuring	JENKINS M	TR 1:30 PM-3:00 PM	1.00	4	45	<input type="checkbox"/>
LGST806407	NEGOTIATIONS	BRANDT A	W 3:00 PM-6:00 PM	1.00	3	34	<input type="checkbox"/>
LGST806409	NEGOTIATIONS	DIAMOND S	R 3:00 PM-6:00 PM	1.00	3	0	<input type="checkbox"/>

Figure 1 is a screenshot of the top of the user interface. Of the nine course sections that are visible, the subject has reported positive values for the first eight. To make adjustments, subjects clicked two checkboxes in the final column of the interface and were prompted to enter the adjustment. Any previously entered adjustments were listed at the top of the interface. The subject has made one adjustment of -91, which tells the mechanism that getting the two accounting classes (i.e., the first two courses visible) together in his schedule together is worth 0, effectively reporting that the subject wants one or the other, but not both, accounting courses. The subject does not need to input scheduling constraints, which are known to the system.

To help subjects report their preferences over schedules correctly, subjects were provided with a “top-ten widget” (see Figure 2), which allowed them to see what CEEI thought their ten favorite schedules were, in order, with the accompanying sum of the cardinal utilities and adjustments (listed as “Schedule Value”). They could use the widget at any time while reporting their values and could then go back to make corrections to their reported values (e.g., if they realized the ten schedules listed were not their favorites or were in the wrong order).

Figure 2: Screenshot of Top Ten Widget

MY TOP 10 SCHEDULES

Given the values you reported, your agent thinks these are your 10 favorite schedules. Your agent will try to buy you these schedules, in this order. Note that depending on the market clearing prices, the schedule you get may not appear on this list, but your agent will buy you the best schedule that you can afford.

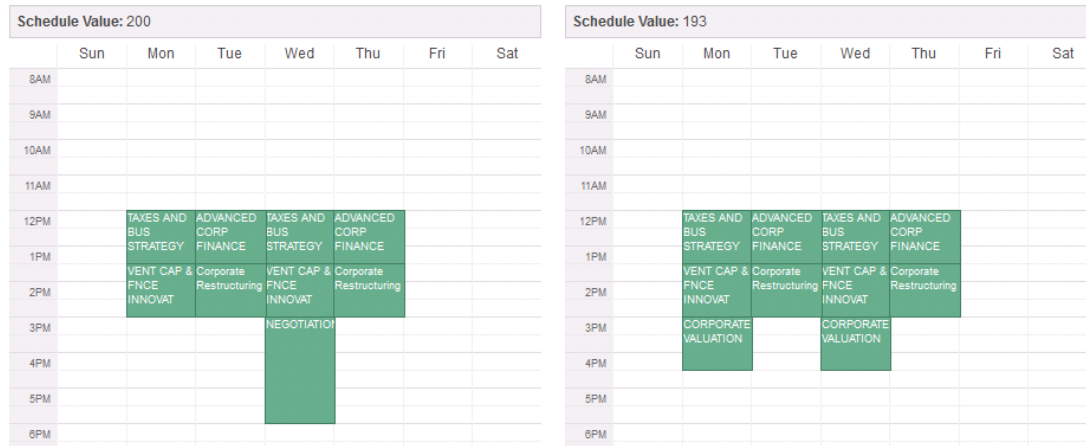


Figure 2 is a screenshot of the top of the “top ten widget”. It shows two feasible schedules of 5 courses each (e.g. “Taxes and Business Strategy meets from 12:00-1:30 on Monday and Wednesday in both schedules) and the sum of the cardinal reports, listed as “Schedule Value.” The rest of the top ten schedules were shown below these, and subjects could scroll down the screen to see all ten.

2.4 Binary comparisons

After using both mechanisms, subjects were asked to evaluate binary comparisons of schedules, reporting which of two schedules they preferred and whether they “slightly prefer,” “prefer,” or “strongly prefer” one schedule over the other (see Figure 3).²⁴ The logic behind using binary comparisons was that while reporting ordinal preferences over every possible schedule is complex, making binary comparisons between two schedules is simple.

²⁴ Subjects were not provided with an option of reporting indifference.

Figure 3: Screenshot of Binary Comparison Question

What is your preference between the schedules below?

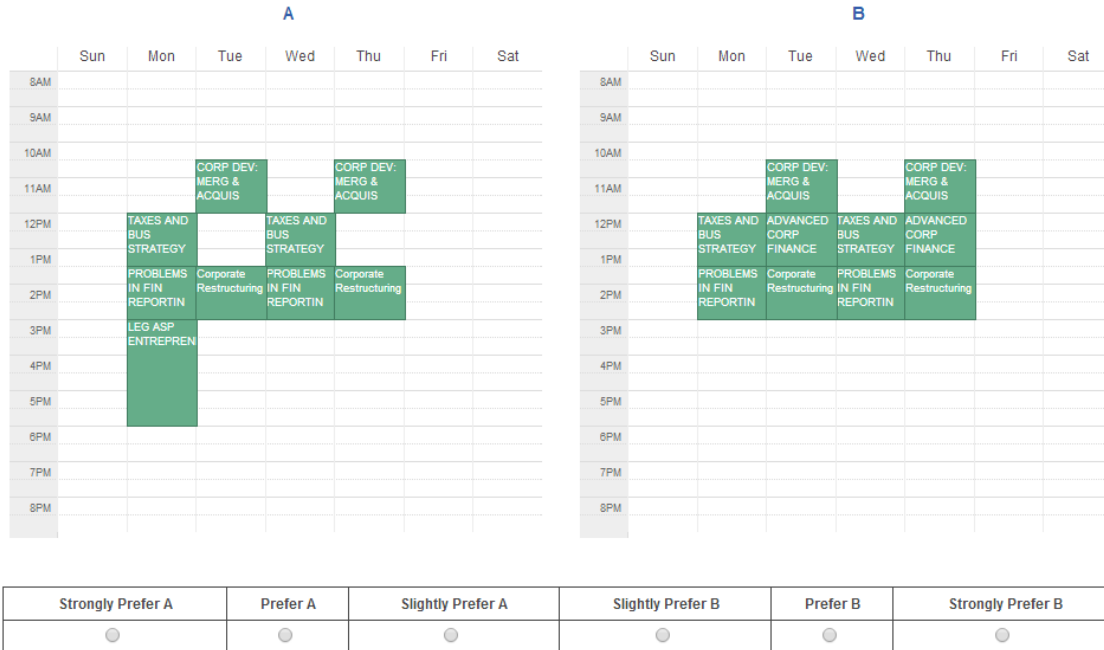


Figure 3 is a screenshot of a binary comparison. It shows two schedules and asks the subject to pick which of the two she prefers. Subjects made up to nineteen binary comparisons.

We designed the set of binary comparisons to allow us to perform a number of tests within each mechanism and between mechanisms. In particular, the first and last binary comparisons were between the schedule the subject received under CEEI and the schedule she received under the WA. This comparison was asked twice, once at the beginning and once at the end, with the schedules reversed (i.e., the schedule shown on the left in the first question was shown on the right in the last question) and is used to construct a social welfare comparison between the two systems.

In addition, up to twelve binary comparisons per subject were asked so that we could construct a measure of envy for each course allocation mechanism. Envy occurs when an individual prefers someone else’s schedule to his own schedule. For each allocation mechanism, each subject was asked to compare his schedule from that mechanism to up to six schedules other subjects in his session received from the mechanism (e.g., he compared his CEEI schedule to others’ CEEI schedules and his WA schedule to others’ WA schedules).²⁵

²⁵ The others’ schedules were chosen to be somewhat desirable to the subject (so that these comparisons would be non-trivial). In particular, each subject was randomly shown up to six

Finally, we included five binary comparisons in which we showed subjects CEEI schedules they would have received if their budget had been 10% or 30% higher or 10% or 30% lower than it actually was. These comparisons were designed to check whether subjects reported their preferences accurately, as subjects should always prefer the schedule they obtain with a larger budget. In practice, however, all comparisons subjects made between two legal schedules are tests of the reporting language since CEEI preference reports create a rank order list that predicts either an indifference or a strict preference between any two schedules. Consequently, we can analyze the preference-reporting language by investigating how often subjects validate or reject this predicted preference in their binary choice data.²⁶ Table 1 shows the types of binary comparisons we asked of subjects.

schedules from the set that generated at least 50% of the utility (based on the reports to the CEEI system) as their realized schedule under the CEEI mechanism.

²⁶ Rarely, subjects saw schedules that were not “legal” by Wharton standards: schedules from the WA could have had more than or fewer than five courses and schedules from CEEI could have fewer than five courses (the system never bought more than five courses but would buy fewer than five if a bundle of fewer than five courses maximized utility). When subjects were making their binary choices, they were asked to imagine that they would have to make their schedules “legal” (i.e. drop a class if they had an extra or pick up a class with an empty seat if they had too few). When we construct the CEEI value associated with schedules with fewer than five courses, we assume that a value of 0 for a course that is picked up, and for schedules with more than five courses, we use the highest possible CEEI value that can be constructed using a subset of five courses from the courses available.

Table 1: Binary Comparison Questions

Question	Number	Rule
CEEI vs. Auction	2	Subjects compared their realized schedule from the auction and their realized schedule from CEEI with the order of schedules randomized. Subjects saw this as the first comparison and again as the last comparison, with the order of schedules reversed.
CEEI Envy	Up to 6	Subjects compared their realized schedule from CEEI with CEEI schedules of other subjects, randomly selected from the pool of CEEI schedules that had CEEI cardinal utility of at least 50% of the cardinal utility of their CEEI schedule.
Auction Envy	Up to 6	Subjects compared their realized schedule from the Auction with Auction schedules of other subjects, randomly selected from the pool of Auction schedules that had CEEI cardinal utility of at least 50% of the cardinal utility of their CEEI schedule.
CEEI Budget Questions	Up to 5	Subjects compared their realized schedule from CEEI with the schedule they would have gotten from CEEI if their budget had been 30% smaller, 10% smaller, 10% larger, and 30% larger. Subjects also compared the schedule they would have gotten with a 10% larger budget to the schedule they would have gotten with a 10% smaller budget. Subjects were not shown these comparisons if they would have received the same schedule under both budgets.

Table 1 describes the type of binary comparison questions subjects were asked and the responses available to the subjects. The scale was: “Strongly Prefer A”, “Prefer A”, “Slightly Prefer A”, “Slightly Prefer B”, “Prefer B”, “Strongly Prefer B”.

2.5 Incentives

Before we delve into the discussion of the results, it is worth reiterating that decisions in our experiment are not incentivized. Since we did not give subjects preferences, we do not know which schedules they actually prefer and so cannot pay them more for getting better schedules under either of the mechanisms. This raises two potential concerns. One concern is that subjects might purposefully misreport their preferences, or purposefully lie in the answers to binary comparison questions and qualitative questions. The second concern is that subjects might not work as hard in the experiment as if they were being paid.

Regarding the first concern, while it is impossible for us to rule out the possibility that subjects would purposefully misrepresent themselves, we have no reason to believe that they would do so. We find the second concern, that subjects might not exert as much effort as they would in an incentivized environment, more realistic. But, to the extent that the experiment being unincentivized led subjects to work less hard than they would otherwise in creating their schedules, or led them to report bids randomly in the auction or values randomly in CEEI, we would interpret this as noise in our data that would not systematically favor one market design to the other. Such noise would just make it harder to identify any effects between the systems.

III. Results on Fairness and Efficiency

In this section, we present our results on fairness and allocative efficiency. We present two sets of results: the first set of results is based on the binary comparisons and so is a joint test of the CEEI mechanism with the reporting language as compared to the WA. The second set of results is based on the reported preferences under CEEI, and represents an isolated test of CEEI versus WA under the assumption that subjects report their preferences perfectly to CEEI.

We think of the first set of results as quantifying how well CEEI actually did on measures of fairness and efficiency and the second set as providing an upper bound for how well CEEI might do in practice if students were able to report their preferences more accurately. Section 4 will explore imperfect preference reporting in much greater detail.

3.1 Results on Fairness

We begin with our results on fairness, which provide the most direct test of the theory in Budish (2011).

Student A is said to *envy* Student B if A prefers B's schedule to her own. An allocation is called *envy-free* if no student envies another. Envy-freeness, introduced by Foley (1967), is arguably the most important criterion of outcome fairness in the economics literature on distributive justice (Moulin 1995; Arnsperger 1994). Unfortunately, in an indivisible goods problem such as course allocation, it is impossible to eliminate envy altogether. If there is some star professor whose course all students want to take (and whose course they value over any other bundle of courses), the students who do not get that course will envy the students who do. Budish (2011) shows, however, that CEEI approximately eliminates envy. More precisely, envy only occurs because of the small randomness in students' budgets. A student with a budget of 5001 might envy a student with a budget of

5002, if there is some schedule that costs 5002 that the first student wants but cannot afford, but the student with a budget of 5001 will never envy a student with a budget of 5000. In addition, any envy that does occur is bounded in magnitude. In the WA, by contrast, there is no such guarantee.

Our binary comparisons directly tested for the presence of envy and its magnitude. Specifically, each subject was presented with a set of binary comparisons for each mechanism asking which schedule they preferred between their own schedule from that mechanism and another randomly chosen subject's schedule from that mechanism, as well as the intensity of that preference ("slightly prefer", "prefer", "strongly prefer"). In total, 117 students completed binary comparisons looking for envy in CEEI and 119 completed binary comparisons looking for envy in the WA.²⁷ Table 4 shows that CEEI generated less envy than WA, measured either by the percentage of subjects who display any envy in Panel (a) or by the percentage of binary comparisons across the entire experiment that resulted in envy in Panel (b).²⁸ We report one-sided tests as the theory in Budish (2011) predicts CEEI will have less envy than the WA. The difference is especially significant when we restrict attention to what we call "large" envy, which drops envy caused by a mild preference for the other schedule.

Table 2 shows that there is significantly less envy in CEEI than in the WA, as expected given the theory. Next, we look at the envy comparison under the assumption that preference reporting under CEEI is perfectly accurate. Under this assumption, we can look for envy by directly comparing a subject's utility from their own schedule to their utility from another subject's schedule. While in principle we can do this for all pairs of subjects in a session, we restrict attention to the Subject A-Subject B pairs for which there were binary comparison tests as well to facilitate comparison with the results in the previous table. Table 3 displays the results.

²⁷ We do not have data from all 132 subjects for two reasons. First, a bug in the code for the first three sessions prevented getting binary comparison data from the six subjects who got the same schedule under both CEEI and the WA. Of the 126 subjects who were left, nine had no other subject in their session with a CEEI schedule that had utility greater than or equal to 50% of their own CEEI utility and seven had no other subject in their session with a WA schedule that had utility greater than or equal to 50% of their own CEEI utility.

²⁸ The 119 subjects who answered at least one Wharton Auction envy question answered 4.2 envy questions on average; the 117 subjects who answered at least one CEEI envy question answered 4.1 envy questions on average. The difference in number of envy questions is not statistically significant (t-test; $p=0.529$)

Table 2: Envy under CEEI and WA — Joint test of the mechanism and the reporting mechanism, using Binary Comparisons

	WA	CEEI	Probability Ratio Test (one-sided)
Panel (a) – By Subject (WA: N = 119, CE: N = 117)			
% of subjects who display any envy of another subject’s schedule	42%	31%	$p = 0.036$
% of subjects who display any large envy (“prefer” or “strongly prefer”) of another subject’s schedule	34%	21%	$p = 0.008$
Panel (b) – By Comparison (WA: N = 499, CE: N = 475)			
% of comparisons in which the subject displays any envy of the other subject’s schedule	19%	12%	$p = 0.002$
% of comparisons in which the subject displays any large envy (“prefer” or “strongly prefer”) of the other subject’s schedule	14%	8%	$p = 0.002$

Table 2 reports the envy results based on binary comparisons. Panel (a) shows the percentage of subjects who experienced any envy in the WA and in CEEI (top row) and the percent of subjects who display “large envy,” which we define as reporting that they “prefer” or “strongly prefer” another subject’s schedule (second row). Panel (b) repeats the analysis but shows the percentage of all binary comparisons subjects faced in the WA and in CEEI. Signrank Tests between the percentage in the WA and in CEEI are reported in the last column. These are one-sided tests based on the theory in Budish (2011), which suggests that CEEI should have less envy.

Table 3 shows that, under the assumption of perfect preference reporting, envy under CEEI is nearly eliminated. This can be interpreted as both a positive result and a negative result for CEEI. The negative interpretation is that subjects had substantial difficulty reporting their preferences in the CEEI mechanism: any subject who had no envy under the assumption of perfect reporting but had some envy based on binary comparisons must have failed to report their preferences accurately, since their preference based on reported utility contradicted their subsequent binary choice. The positive interpretation, on the other hand, is that it gives a sense of the upper bound of performance of CEEI in idealized conditions, in which students are able to perfectly report their utility functions (e.g., with more training, more time, and improvements to the preference-reporting language).

Table 3: Envy under CEEI and WA — Isolated test of the mechanism, assuming perfect reporting of CEEI preferences

	WA	CEEI	Probability Ratio Test (one-sided)
Panel (a) – By Subject (WA: N = 119; CE: N = 117)			
% of subjects who envy another subject’s schedule according to CEEI	29%	4%	$p < 0.001$
Panel (b) – By Comparisons (WA: N = 499; CE: N = 475)			
% of comparisons in which one subject envies the other subject’s schedule according to CEEI	15%	2%	$p < 0.001$

Table 3 reports the envy results based on the utility measure constructed from the CEEI preference reports. We analyze the Student A-Student B pairs for which there was a binary comparison. When necessary, we made non-legal WA schedules legal by maximizing CEEI utility among courses in the WA schedule subject to scheduling constraints. Panel (a) shows the percentage of subjects who experienced any envy in the WA and in CEEI. Panel (b) repeats the analysis but shows the percentage of all comparisons subjects faced in the WA and in CEEI. Signrank Tests between the percentage in the WA and in CEEI are reported in the last column. These are one-sided tests based on the theory in Budish (2011), which suggests that CEEI should have less envy.

3.2 Results on Allocative Efficiency

We now turn to our results on allocative efficiency. The theory in Budish (2011) focuses on ex-post Pareto efficiency, that is, whether or not an allocation leaves Pareto-improving trades on the table (trades that make all students weakly better off with at least one strictly better off). Unfortunately, the binary comparisons data cannot be used to measure ex-post Pareto inefficiency, since they only contain data on subjects’ preferences between a small number of pairs of *schedules*, rather than data about preferences over individual courses that could help us determine, e.g., if subjects A and B should swap course X for course Y.²⁹ Instead, we use a simple measure of allocative efficiency that was important to the Wharton deans, namely, how many students prefer their CEEI schedule to their

²⁹ An exception is that our binary comparison data do allow us to detect Pareto-improving trades in the unusual event of an envy cycle, meaning that student A envies student B’s schedule, who envies student C’s schedule, ... who envies student A’s schedule. Over all of our sessions, we detected two envy cycles under CEEI and one envy cycle under the WA, with each of these three envy cycles consisting of two students.

WA schedule, both in aggregate and in each economy (i.e., each session of subjects who were competing for the same seats in courses).

Recall that we asked subjects to compare their CEEI schedule to their WA schedule twice, with many other binary comparisons in between, and with the position of CEEI and WA schedules flipped in the second binary comparison. We say that a subject prefers CEEI (WA) if they report that they prefer their CEEI (WA) schedule both times they were asked; we say that the subject is indifferent if their preference between the two reverses between the two binary comparisons. Subjects who had identical schedules did not see these binary comparison questions. Table 4 reports the allocative efficiency results.

Table 4: Allocative Efficiency — Results from Binary Comparisons of CEEI Schedule and WA Schedule

Session	Students in the Session	Prefer CEEI	Prefer WA	Indifferent	Identical	Majority Voting Result
1	18	8	5	4	1	CE
2	15	5	5	2	3	Tie
3	19	10	4	3	2	CE
4	16	5	4	3	4	CE
5	18	8	8	2	0	Tie
6	14	6	5	1	2	CE
7	18	8	6	2	2	CE
8	14	6	5	0	3	CE
All	132	56	42	17	17	6-0-2

Table 4 shows allocative efficiency results from each session, reported in the order in which they were conducted. *Prefer CEEI* reports the number of subjects in the session who reported they preferred the CEEI schedule in both binary comparisons. *Prefer WA* reports the number of subjects in the session who reported they preferred the WA schedule in both binary comparisons. *Indifferent* reports the number of subjects whose preference for schedules reverses in the two binary comparison questions. *Identical* reports the number of subjects who received identical schedules under the two mechanisms and so did not see the binary comparison questions. *Majority Voting Result* asks which of the two mechanisms would be preferred by the majority of subjects if they were to vote for the mechanisms ex-post based on their preference over the schedules they received from each mechanism.

As can be seen from Table 4, subjects preferred CEEI to WA by a margin of 56-42 (one-sided binomial probability test against the hypothesis that the ratio is 50%, $p = 0.094$), with seventeen students indifferent between the two schedules and seventeen students receiving exactly the same schedule under each.³⁰ At the session level, the majority of students preferred CEEI to WA by a margin of 6 to 0 (one-sided binomial probability test that ratio is 50%, $p = 0.016$), with 2 ties. The session-level aggregation makes sense to the extent that we think of each session as its own market and view majority rule as a social welfare criterion. Both the individual-level and session-level measures suggest that CEEI outperforms the Wharton Auction on this measure of allocative efficiency.³¹

We now turn to efficiency results based on the reported preferences. We report three different analyses.

First, we can repeat the allocative efficiency and majority voting exercises from above but using the reported preferences instead of the binary comparisons. At the individual level, 69% of students prefer their CEEI schedule to their WA schedule based on the reported preferences (one-sided binomial probability test, $p < 0.001$). At the session level, the majority of students prefer CEEI to WA in seven sessions and there is one tie (one-sided binomial probability test, $p < 0.01$).

Second, with reported preferences, we can directly assess the level of ex-post Pareto inefficiency. Specifically, we formulate an integer program that solves for the maximum number of Pareto-improving trades among the subjects in each section, given subjects' reported preferences and the initial allocation arrived at in the experiment. We consider two different ways to account for unutilized capacity. Our preferred method creates an additional fictitious player called the "registrar" who holds all unutilized capacity and has zero utility from each course. We also include a measure that ignores excess supply and seeks trades just among the subjects (i.e., ignores the possibility of Pareto-improving trades involving the registrar). Note that for this latter measure, it is theoretically guaranteed that CEEI will have zero possible trades. For the former measure, CEEI will have Pareto-improving trades because of the small amount of market-clearing error that is

³⁰ There was no statistically discernible difference in the intensity of preference between the students who preferred CEEI to WA and the students who preferred WA to CEEI. On a 3-point scale, with 3 = "strongly prefer," 2 = "prefer," and 1 = "slightly prefer," students who preferred CEEI had an average intensity of 2.27 and students who preferred WA had an average intensity of 2.42 (these are not statistically significantly different, t-test $p = 0.199$). Overall, on a 6-point scale ranging from 6 = strongly prefer CEEI to 1 = strongly prefer WA, the average score among the 115 subjects who faced the binary comparisons was 3.65.

³¹ We did a similar measure of envy by session and find that CEEI wins (i.e. has fewer subjects displaying envy) in 6 sessions and CEEI loses in 2 sessions. This win-loss record is nearly significantly different from 0.5 at the 10% level (one-sided binomial probability test $p = 0.145$).

sometimes necessary to run CEEI.³² Still the theory predicts fewer Pareto-improving trades under CEEI than under the WA, which can have Pareto-improving trades under either measure.

Table 5 reports the results of this exercise. As predicted by Budish (2011), there is substantially less scope for Pareto-improving trades under CEEI than under the WA. For instance, under the WA, 74% of students are able to engage in at least one Pareto-improving trade, versus 17% of students under CEEI.

Table 5: Results on Pareto Efficiency: Reported Preferences

	WA	CEEI	Probability Ratio Test (one-sided)
Panel (a): includes Pareto-improving trades with “registrar” who holds excess supply			
# of Pareto-improving trades detected (% of course seats)	260 (32.8%)	44 (5.6%)	$p < 0.001$
# of students involved in at least one trade (% of students)	98 (74.2%)	22 (16.7%)	$p < 0.001$
Panel (b): ignores excess supply, trades only among students excluding “registrar”			
# of Pareto-improving trades detected (% of course seats)	235 (29.7%)	0 (0%)	$p < 0.001$
# of students involved in at least one trade (% of students)	89 (67.4%)	0 (0%)	$p < 0.001$

Table 5 shows results analyzing Pareto Efficiency based on preferences reported to CEEI. For each session and each mechanism, we used an integer program to detect Pareto-improving trades given the allocation and students’ reported preferences. The integer program maximizes the number of Pareto-improving trades subject to feasibility constraints. We restrict attention to trades in which each student involved in the trade gives and gets a single course seat. A single student may engage in an unlimited number of such trades, and a single trade may involve arbitrarily many students. In Panel (a), excess supply after the initial allocation is held by a “registrar” who is willing to engage in any feasible one-for-one trade. In Panel (b), excess supply after the initial allocation is discarded.

³² Budish (2011) shows that there need not exist prices that exactly clear the market, but guarantees existence of prices that clear the market to within a small amount of approximation error. In Budish (2011), error is defined as the square root of the sum of squares of excess demand errors (too many students assigned to a class) and excess supply errors (empty seats in a desired class). The Wharton course allocation committee viewed excess demand errors as more costly than excess supply errors, and tuned the CEEI software accordingly for the experiment. Over the eight sessions, there were ten total seats of excess supply (median: one seat per session) and two total seats of excess demand (median: zero seats per session). The Pareto-improving trades exercise reported in the text treats the registrar as owning the ten seats of excess supply and ignores the two seats of excess demand.

Third, if we assume interpersonal comparability of utilities, we can look directly at the magnitudes of subjects' utility changes between mechanisms. We do this in two ways. First, we look at each subject and calculate the percentage difference in utility between their schedule from the WA and their schedule from CEEI. This measure is depicted below as Figure 4:

Figure 4: Distribution of change in utility going from WA to CEEI

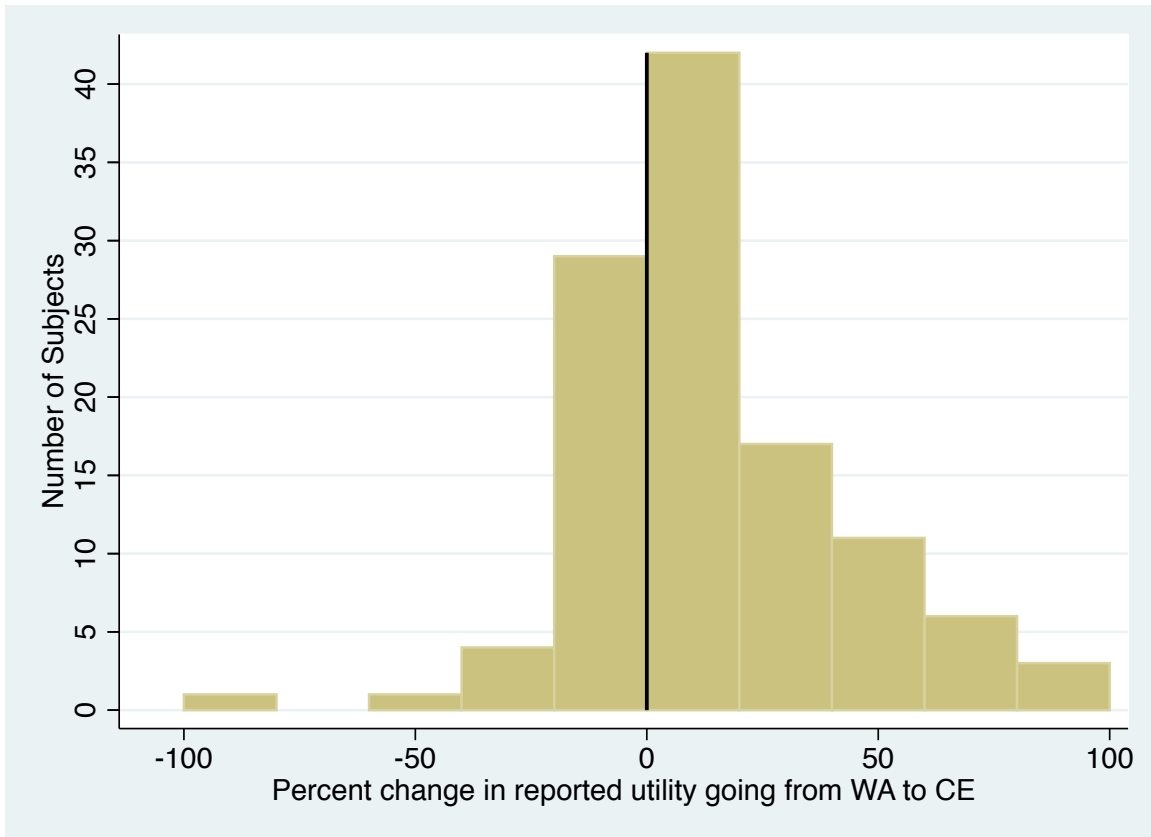


Figure 4 shows a histogram of the percentage change in legal utility going from the schedule received in the WA to the schedule received in CEEI. Bins are 20 percentage points wide and the graph excludes the 18 subjects who got the same utility from both schedules. One observation had a utility increase of over 100% and is included in the highest percentage increase bar.

The majority of mass is to the right of 0 in Figure 4, demonstrating that more subjects prefer their CEEI schedule to their WA schedule according to their CEEI reports (seventy-eight prefer CEEI and thirty-five prefer the WA). Moreover, the winners win more than the losers lose: thirty-seven students have at least a 20% utility improvement when switching from the WA to CEEI, whereas only six students have at least 20% utility harm from switching from the WA to CEEI.

Second, we plot the distribution of utilities from schedules coming from the WA and coming from CEEI and compare the two plots (Figure 5). The distribution of utilities under CEEI second-order stochastically dominates the distribution under the WA. This implies that a utilitarian social planner prefers the distribution of outcomes under CEEI to that under WA, so long as the planner has a weak preference for equality (the social welfare analogue of risk-aversion). Moreover, the comparison is nearly one of first-order stochastic dominance, which would be an even stronger result. However, there are a few subjects who do extraordinarily well in the WA, i.e., the right tail of outcomes under the WA is better than that under CEEI, so we do not obtain first-order stochastic dominance.

Figure 5: Distribution of Utility under CEEI and the WA, based on Reported Preferences

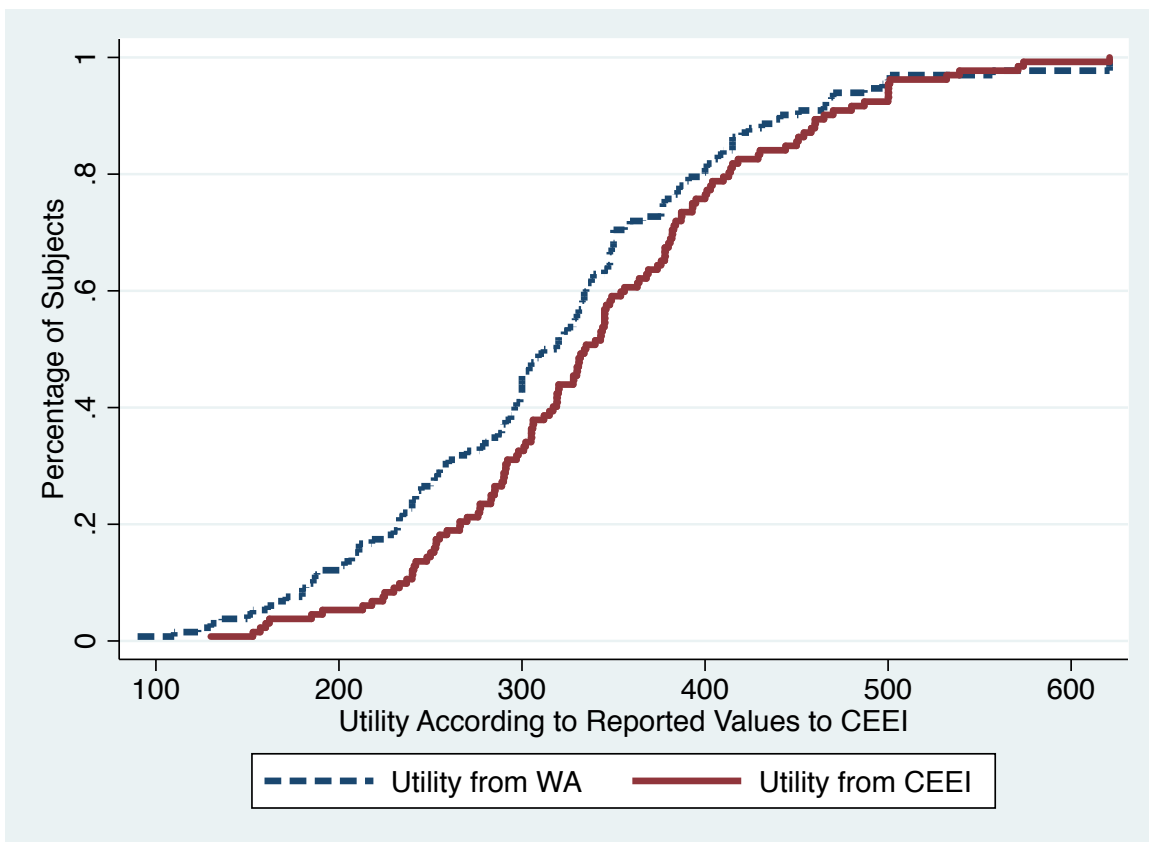


Figure 5 plots the CDF of utility according to reported values to CEEI for both the WA and CEEI. Three utilities (two in the WA and one in CEEI) are above 2,000 and have been Winsorized at 621, the next-highest utility value. The distribution of utilities under CEEI second-order stochastically dominates the distribution under the WA. The distribution of utilities under CEEI almost first-order stochastically dominates the distribution under the WA except for some subjects who received very high utilities under the WA.

3.3 Slack Analysis

In consultation with the Wharton Course Allocation Redesign Team, we recruited twenty-five students per session with the aim of obtaining twenty subjects per session. Our turnout was worse than forecast, especially in the evening sessions. As a result, the number of subjects per session ranged from fourteen to nineteen. This variation in attendance inadvertently generated variation in what we term “slack,” defined as $100 \times \left(\frac{\#seats\ supplied}{\#seats\ demanded} - 1 \right)$. “Slack” is the percentage of excess capacity in the available courses times 100, and it ranged from 12.6 to 29.3 in our sessions.³³

If there is too much slack in a market, then the allocation problem is trivial: under virtually any course allocation system, all students will get exactly what they want (or something close to it). Thus, in sessions with a relatively large amount of slack, we might expect that CEEI and WA would do equally well, whereas in sessions with a relatively small amount of slack, we may expect the relative advantages of CEEI to be more pronounced.

Table 6 presents regression results that analyze how the level of slack in our sessions affects whether subjects were more likely to display envy under CEEI or the WA and whether subjects preferred their schedule from CEEI to their schedule from the WA. This analysis was meaningful to the decision makers at Wharton because the slack in their real-world problem is in the range of 15 to 20, slightly less than the average amount of slack in our experiment.

The significant negative effect of slack on the likelihood that subjects show less envy under CEEI and on the preference of the CEEI schedule demonstrates that as slack decreases (i.e., the allocation problems get harder), the benefits of CEEI over the WA in terms of fairness and allocative efficiency become more pronounced. The results suggest that for each point of slack (i.e., percentage point of excess seats), the percentage of subjects who experience less envy under CEEI than the WA decreases by 0.696 percentage points. Similarly, for each additional point of slack, the percentage of subjects who prefer their CEEI schedule to their WA schedule decreases by 0.918 points. Essentially, CEEI outperforms the WA more dramatically on the harder problems where there is low slack.

³³ We adjusted the number of seats available in the experimental session based on the number of subjects who participated, but in a coarse way. For example, if fewer than eighteen students showed up to a session, then all courses with five seats would be adjusted to have four seats. This allowed us to implement and communicate changes to the number of seats even though subjects had printed instructions with the number of seats for each course.

Table 6: Effect of Slack on Relative Performance of CEEI vs. WA

	<i>Fairness</i>		<i>Allocative Efficiency</i>
	Less Envy CEEI	Less Strong Envy CEEI	Prefer CEEI
	(1)	(2)	(3)
Constant	55.47 (3.54)***	56.50 (2.99)***	55.30 (1.21)***
<i>Slack – Mean(Slack)</i>	-0.696 (0.235)**	-1.35 (0.218)***	-0.918 (0.305)**
Observations	115	115	132
Clusters	8	8	8

Table 6 reports OLS regressions. *Slack* is defined as $100 \times \left(\frac{\#seats\ supplied}{\#seats\ demanded} - 1 \right)$. It ranges from 12.6 to 29.3 across the eight sessions, and $Mean(Slack) = 20$. *Less Envy CEEI* = 100 if the subject displayed envy under the WA but not under CEEI; *Less Envy CEEI* = 50 if the subject displayed envy under both or no envy under both systems; *Less Envy CEEI* = 0 if the subject displayed envy under CEEI but not under the WA (this allows us to treat the constant as the percentage of subjects who have less envy under CEEI under the average amount of slack). *Less Strong Envy CEEI* is the same as Less Envy but only counting subjects that report they “prefer” or “strongly prefer” another subject’s realized schedule. The envy results are restricted to the 115 subjects who saw at least one envy binary comparison under each mechanism. *Prefer CEEI* = 100 if the subject preferred their CEEI schedule to their WA schedule both times they were asked; *Prefer CEEI* = 50 if the subject is indifferent (i.e., switched preferences between schedules) or got an identical schedule under both systems; *Prefer CEEI* = 0 if the subject preferred their WA schedule to their CEEI schedule both times they were asked (this allows us to treat the constant as the percentage of subjects who would prefer CEEI under the average amount of slack). Robust standard errors, clustered by session, are in parentheses. *, **, and *** indicate significance at 0.1, 0.05, and 0.01 respectively.

3.4 Discussion of Results on Fairness and Efficiency

The results presented in this section suggest that CEEI consistently outperforms the WA on measures of both fairness and allocative efficiency. CEEI outperforms on every measure considered, with most comparisons statistically significant. The improvements generated by CEEI are especially large on the “hard” allocation problems with low slack.

However, CEEI wins by much less based on the binary comparisons (i.e. our joint test of CEEI together with its reporting language) than based on the CEEI preference reports (i.e., our isolated test of CEEI, assuming perfect preference reporting). This tells us that

difficulty in preference reporting was an important factor in mechanism performance. We turn to this subject in the next section.

IV. Preference Reporting

The theory in Budish (2011) assumes that students report ordinal preferences over all possible schedules of courses. The assumption plays an analogous role in Budish (2011) as the assumption of complete preferences plays in consumer theory, in that it ensures the mechanism knows the agent's most-preferred choice from any possible choice set. This guarantees that the mechanism will deliver the agent his most-preferred schedule out of all permissible schedules he can afford at the realized prices.

In practice, however, the mechanism cannot realistically ask agents to report a rank order preference over all feasible schedules (e.g., there are hundreds of millions of possible schedules in a typical semester at Wharton). It will instead rely on a shorthand (i.e., a vector of information an agent can feasibly report) and will use that information to create a rank order preference list for the agent. This raises the question of whether agents are able to report their preferences using such a shorthand with sufficient accuracy to realize the theoretical benefits of CEEI. This question was one of the main motivations for running the experiment.

Results from Section 3 suggest that agents are able to report their preferences “well enough” for CEEI to outperform the WA on measures of fairness and allocative efficiency, but preference reporting is still an important concern for CEEI. While CEEI outperforms the WA dramatically under the assumption that CEEI reported preferences are perfectly accurate (see Table 3), under the binary comparison analysis, which does not presume accurate reporting, the advantage of CEEI over the WA is much smaller in magnitude (see Table 2). This suggests that if preference reports had been perfect reflections of their preferences, subjects would have reaped significantly more gains from CEEI.

In this section, we investigate why subjects failed to report their preferences perfectly. We do this both because understanding the sources of mistakes may help guide future implementation of the mechanism, and, more broadly, to inform researchers who encounter similar language design issues in other settings.

Conceptually, there are two possible reasons why the preference reports subjects made to the mechanism might not reflect their true preferences. First, the preference-reporting language we designed for the experiment does not allow students to express all possible preferences. If students have preferences that they are unable to express using the language, this can create a wedge between students' reported preferences and their true

preferences. Second, even for preferences that students are in principle able to express, students may nevertheless have difficulty reporting with perfect accuracy. We investigate each of these two sources of mistakes in turn, in Sections 4.2 and 4.3, respectively. Section 4.1 precedes those sections and presents summary statistics about the use of the mechanism and the preference-reporting mistakes. Section 4.4 discusses the lessons for practical implementation of CEEI and for language design more broadly.

4.1 Summary Statistics

Table 7 presents summary statistics describing how subjects used the preference-reporting language. The data suggest that subjects generally followed the instructions we provided. We advised subjects to report positive cardinal values for at least twelve courses. The median number of courses assigned positive values was 12 and the vast majority of subjects (76.5%) reported positive values for 11 or more courses. In addition, we advised subjects to assign their favorite course a value of 100 and to assign all other courses a relative value. Again, the vast majority of subjects (75.0%) reported a value for 100 for one and only one course. Generally speaking, subjects spread their values of courses evenly from 0 to 100. The last three rows suggest that most subjects chose not to use any adjustments (the median subject used 0 adjustments) and the average number of adjustments across all subjects was slightly more than 1. Of those adjustments that were made, they were evenly split between positive and negative adjustments.

Table 7: Use of the Preference-Reporting System

	Mean	Min	25 th Pct.	Median	75 th Pct.	Max
# courses valued $v > 0$	12.45	7	11	12	14	24
# courses valued $v = 100$	1.40	0	1	1	1	8
# courses valued $50 \leq v \leq 99$	4.87	0	3	5	7	10
# courses valued $0 < v < 50$	6.17	0	4	6	8	17
# adjustments	1.08	0	0	0	2	10
# adjustments > 0 (complements)	0.55	0	0	0	1	10
# adjustments < 0 (substitutes)	0.53	0	0	0	1	6

Table 7 reports use of the preference-reporting system for the 132 subjects in the experiment. v is the cardinal value assigned to a particular course section.

Table 8 presents summary statistics on the prevalence and magnitude of preference-reporting “contradictions.” For each binary comparison a subject makes, we compare the predicted preference given the subject’s report to the preference-reporting language of CEEI (i.e., summing their cardinal values and adjustments) against the student’s response in the binary comparison. If the CEEI preference report suggests that the subject prefers schedule A to schedule B, then we call that *accurate* if the subject reflects that preference in the binary comparison, but if the subject says he prefers B to A in the binary comparison, we call that a *contradiction*. Since we think that the binary comparisons reflect subjects’ actual preferences between pairs of schedules, a higher percentage of contradictions suggest that the preference-reporting language was less able to capture subjects’ true preferences.

Table 8: Prevalence and Magnitude of Preference-Reporting Contradictions

Utility Difference Between Schedules	# Comparisons with this Utility Difference	Accurate	Contradictions			
			All	Weak Preference	Preference	Strong Preference
All	1,662	84.42%	15.58%	5.17%	6.92%	3.49%
1-9	123	70.73%	29.27%	9.76%	11.38%	8.13%
10-49	516	77.13%	22.87%	6.78%	11.24%	4.84%
50-99	427	85.25%	14.75%	6.32%	5.62%	2.81%
100+	596	92.95%	7.05%	2.01%	3.19%	1.85%

Table 8 shows the percentage of binary comparisons that were contradictions. For each binary comparison, the *Utility Difference Between Schedules* is the utility of the schedule with the higher utility minus the utility of the schedule with the lower utility, as determined by the student’s preference reports under CEEI. The table shows all 1,662 comparisons where this number is greater than 0 and so the preference reports suggest one schedule is preferred to the other. The *Accurate* column reports the percentage of these comparisons where the binary comparison choice confirms the preference report prediction. The *Contradictions* columns report the percentage of binary comparisons that contradicted the CEEI preference reports overall and at each level of preference.

A few observations can be made from the pattern of data in Table 8. First, there are a significant number of contradictions. Across all our data, the binary comparison response contradicted the preference reports 15.58% of the time. Second, there are very few contradictions in the bottom right of the table (i.e., when preference reports assign a big utility difference between the two schedules and the binary comparison indicates that the schedule with the lower utility is “strongly preferred”), suggesting that there are few “big” contradictions. In general, as we move down rows in Table 8, the data shift to the

left, meaning that the preference reports are more likely to pick the preferred schedule and contradictions are more likely to be associated with a weak preference. Of the 596 comparisons when the utility difference is 100 or more, the preference reports contradict the binary comparison responses only 7.05% of the time (and only 1.85% of cases are contradictions in which subjects report a strong preference for the disfavored schedule). In contrast, in the 123 comparisons in which the utility difference between the schedules based on the preference reports is less than 10, 29.27% of cases are contradictions.

4.2 Limitations of the Preference-Reporting Language

The preference-reporting language we used in the experiment was not fully expressive (as defined, e.g., in Nisan 2006). That is, there exist ordinal preferences over schedules that subjects would be mechanically unable to express using the language that was provided. While students could report values for individual courses and positive or negative adjustments for pairs of courses, to be expressive — at least in principle — the language would have had to allow students to report adjustments not just for pairs of courses but for arbitrary sets of courses.³⁴ To allow agents to feasibly express complex preferences in practice, the language would also need to make it easy for students to report the non-additive preferences that are most important to them.³⁵

The set of non-expressible preferences is vast, and we do not have a disciplined way of exploring all such possibilities as a source of preference-reporting contradictions.³⁶ Instead, we look at two specific sources of non-additive preferences that the Wharton committee suggested to us would be the most important, both of which arise from scheduling considerations per se rather than the contents of the classes within the schedule.³⁷

³⁴ In theory, adjustments of sets of five courses would be sufficient for the language to be fully expressive as a complete Wharton schedule is five courses, and an agent could offer an adjustment for every feasible schedule.

³⁵ For example, suppose a student wants to express that they want at most one out of a set of k classes. They could express this in principle using just pairwise adjustments, but it would take $k(k-1)/2$ such adjustments (reporting that any two of the k courses together have negative total value). A simpler way to convey the same preferences would be to report a constraint of the form “at most one out of these k ,” were the ability to do so provided. There are numerous analogous examples.

³⁶ With roughly 50,000 possible schedules in the lab, there are 50,000! possible ordinal preferences over schedules, or roughly $10^{12,499}$. As such, the up to nineteen binary comparisons we ask of subjects does not provide enough data to identify patterns in such a large set without prior guidance on where to look.

³⁷ In the lab environment there were twenty-five representative Spring 2012 classes out of over 150 course sections in practice. With more courses, it is likely that there would be common sources of non-additive preferences arising from curricular considerations as well. For instance, a

The first is whether the student's schedule is balanced — at least one class on each day Monday through Thursday (there are essentially no classes on Friday). The second is whether the schedule is contiguous — every day on which a student has class he has at most one 1.5-hour gap between the start of the first class and the end of that last one. According to the Wharton committee, each of these characteristics makes a schedule “elegant,” which some students value highly. However, subjects are not able to express a value for either characteristic using the supplied preference language. We therefore investigate whether there are more contradictions when the schedule we expect a subject likes less based on the preference reports to CEEI has one of these elegant features (and thus should get a utility bump that is unreported).

Table 11 is broken up into two panels, one for each of the two features: whether the schedule is balanced (Panel A), and whether the schedule is contiguous (Panel B). The table summarizes the prevalence and magnitude of preference-reporting accuracies and contradictions as a function of which of the schedules in the binary comparison had the elegant schedule feature.

Results from Table 11, Panel A show that subjects' binary comparison responses are more likely to contradict their CEEI preferences when the schedule that CEEI predicts they disfavor is balanced. Subjects are more likely to make a contradiction when the schedule CEEI predicts they disfavor is balanced and the other is not (29.41% are contradictions) than when both are balanced (15.05% are contradictions; $p < 0.01$) or when the schedule CEEI predicts they favor is balanced and the other one is not (13.6% are contradictions; $p = 0.036$).

Table 11, Panel B depicts a similar pattern when looking at whether a schedule is contiguous. Subjects are directionally, but not significantly, more likely to make a contradiction when the schedule CEEI predicts they disfavor is contiguous and the other is not (19.23% are contradictions) as compared to when both are contiguous (15.83% are contradictions, $p = 0.401$) or when the schedule CEEI predicts they favor is contiguous and the other is not (12.56% are contradictions, $p = 0.120$).

student might wish to take at most three Finance classes out of the twelve offered in Spring 2012, but in the lab there were only four Finance classes in total so this constraint would be unlikely to bind. See Milgrom (2009) for an example of a preference-reporting language that allows agents to express preferences of this form – at most k out of set S .

Table 11: Prevalence and magnitude of preference-reporting contradictions for comparisons with and without elegant schedules

Type of Comparison	# Comparisons with this Utility Difference	Accurate	Contradictions			
			All	Weak Preference	Preference	Strong Preference
Panel A: Balanced Schedule						
Neither has it	23	73.91%	26.09%	13.04%	13.04%	0%
Only higher rated has it	66	86.40%	13.60%	3.03%	6.06%	4.55%
Only lower rated has it	51	70.59%	29.41%	7.84%	15.69%	5.88%
Both have it	1,522	84.95%	15.05%	5.06%	6.57%	3.42%
Panel B: Contiguous Schedule						
Neither has it	52	88.46%	11.56%	9.63%	0%	1.92%
Only higher rated has it	199	87.44%	12.56%	4.52%	4.02%	4.02%
Only lower rated has it	192	81.77%	18.23%	7.29%	7.81%	3.12%
Both have it	1,219	84.17%	15.83%	4.76%	7.55%	3.53%

Table 11 shows the 1,662 comparisons and splits them based on whether each of the schedules in the comparison was balanced, defined as at least one class on each day Monday through Thursday (Panel A), and whether each of the schedules in the comparison was contiguous, defined as not having a break of more than 1.5 hours between classes on a given day (Panel B). The *Type of Comparison* indicates which schedules in the comparison had the elegant feature. “Higher rated” and “lower rated” refer to the schedule predicted to be favored by the CEEI preference reports and the schedule predicted to be disfavored, respectively. The *Accurate* column reports the percentage of these comparisons where the binary comparison choice confirms the preference report prediction. The *Contradiction* columns report the percentage of binary comparisons that contradicted the CEEI preference reports overall and at each level of preference.

That subjects are more likely to make a contradiction when CEEI predicts they favor a schedule that is not balanced or not contiguous suggests that at least some of the contradictions are due to the preference-reporting language failing to provide a way for agents to report important features of their preferences. An important caveat is that each of these specific types of non-expressible preferences accounts for only a small number

of contradictions each; however, there are likely other non-expressible preferences that we do not quantify here.

4.3 Difficulty Using the Preference-Reporting Language

The previous section investigated preference-reporting mistakes arising from limitations of the reporting language, that is, types of preferences that mechanically cannot be expressed using the tools provided. The other potential source of preference-reporting mistakes is cases where the language was in principle sufficient for the subject to express their preferences, but he nevertheless failed to do so accurately. Our reporting language had two components — cardinal values to express preferences for individual courses and pairwise adjustments to express certain kinds of substitutabilities and complementarities for pairs of courses — which we explore in turn.

A common intuition in the market design literature is that agents find it easier to report ordinal preferences over items (e.g., I like w better than x better than y better than z) than cardinal preferences over items (e.g., *how much* I like w better than x better than y better than z). We therefore examine whether the comparisons for which cardinal preference information was pivotal were more likely to generate contradictions than comparisons for which ordinal information was sufficient for CEEI to determine the subject's preference.

We define a comparison between schedules A and B as an *ordinal comparison* if the student's preference report generates a clear preference between A and B based on ordinal information alone. For example, if A consists of the student's $\{1^{\text{st}}, 3^{\text{rd}}, 5^{\text{th}}, 7^{\text{th}}, 9^{\text{th}}\}$ favorite courses, B consists of the student's $\{2^{\text{nd}}, 4^{\text{th}}, 6^{\text{th}}, 8^{\text{th}}, 10^{\text{th}}\}$ favorite courses, and neither schedule triggers adjustments, then A can be determined to be preferred to B without knowing the specific cardinal utilities the student assigned to each course. When one schedule can be determined to be preferred to the other based on ordinal information alone, we say that schedule "rank dominates" the other schedule; e.g., in the example above we say that schedule A rank dominates schedule B.

We define a comparison between schedules A and B as a *cardinal comparison* if neither schedule triggers an adjustment and neither schedule rank dominates the other. For example, if schedule A consists of a student's $\{1^{\text{st}}, 2^{\text{nd}}, 8^{\text{th}}, 9^{\text{th}}, 10^{\text{th}}\}$ favorite courses and schedule B consists of a student's $\{3^{\text{rd}}, 4^{\text{th}}, 5^{\text{th}}, 6^{\text{th}}, 7^{\text{th}}\}$ favorite courses, ordinal information alone is insufficient to determine which is preferred. These are the comparisons for which students' ability to report cardinal preference information accurately is put to the test.

Table 9 summarizes the prevalence and magnitude of preference-reporting accuracies and contradictions as a function of whether the comparison in question is an *ordinal*

comparison or a *cardinal comparison*. The table shows that contradictions are much more common in the case of cardinal comparisons (31.72%) than in ordinal comparisons (10.94%), an increase in the likelihood of contradiction of 189%. This difference is highly statistically significant (Probability Ratio Test, two-sided test, $p < 0.001$) and is robust to controls for the difference in utility between the two schedules.³⁸

Table 9: Prevalence and magnitude of preference-reporting contradictions for ordinal and cardinal comparisons

Type of Comparison	# Comparisons with this Utility Difference	Accurate	Contradictions			
			All	Weak Preference	Preference	Strong Preference
All	1,580	84.18%	15.82%	5.32%	6.96%	3.54%
Ordinal	1,207	89.06%	10.94%	4.06%	4.39%	2.49%
Cardinal	373	68.36%	31.64%	9.38%	15.28%	6.97%

Table 9 shows all 1,580 comparisons in which neither schedule triggered an adjustment. *Ordinal* indicates the CEEI reports predicted which schedule was preferred based on ordinal information alone. *Cardinal* indicates the CEEI reports predicted which schedule was preferred based on the cardinal information associated with each course. *Accurate* reports the percentage of these comparisons where the binary comparison choice confirms the CEEI preference report prediction. The *Contradictions* columns report the percentage of binary comparisons that contradicted the CEEI preference reports overall and at each level of preference.

These results are strongly consistent with the intuition from the market design literature that ordinal preference information is easier to report than cardinal preference information (e.g., Bogomolnaia and Moulin 2001), and suggest that difficulty reporting cardinal preference information was an important source of mistakes in the lab.

Next, we explore subjects' use of adjustments. Table 7 above shows that pairwise adjustments were not used as widely as one might have expected — the median subject used no adjustments. We ask whether, in the cases where adjustments were used, they enhanced or detracted from reporting accuracy.

³⁸ The rank dominant comparisons have slightly larger utility differences than the cardinal comparisons (99.4 versus 82.7), but controlling for the difference in utility, we still observe that comparisons are significantly more likely to be a contradiction if they rely on cardinal information rather than just ordinal information. A regression that controls non-parametrically for the utility difference shows that the cardinal comparisons are 16.1 percentage points (i.e., 147%) more likely to be a contradiction ($p < 0.001$).

Table 10 summarizes the prevalence and magnitude of preference-reporting accuracies and contradictions as a function of whether the comparison activated an adjustment. Due to the relatively limited use of adjustments, only eighty-two of the binary comparisons involved a schedule in which one or more adjustments were activated for the subject. That said, in these eighty-two cases, only 10.98% yielded preference-reporting contradictions versus 15.82% for the comparisons that did not involve an adjustment (Probability Ratio Test, two-sided test, $p=0.239$). The relatively low rate of contradictions in the eighty-two cases when adjustments are activated suggests that adjustments did not detract from preference-reporting accuracy, and may have slightly enhanced it (though the difference is not statistically significant).³⁹

Table 10: Prevalence and magnitude of preference-reporting contradictions for comparisons with and without adjustments

Type of Comparison	# Comparisons with this Utility Difference	Accurate	Contradictions			
			All	Weak Preference	Preference	Strong Preference
All	1,662	84.42%	15.58%	5.17%	6.92%	3.49%
No-Adjustment	1,580	84.18%	15.82%	5.32%	6.96%	3.54%
Adjustment	82	89.02%	10.98%	2.42%	6.10%	2.42%

Table 10 shows all 1,662 comparisons. *Adjustment* indicates that one of the schedules in the comparison activated an adjustment in the CEEI preference reports. *No-Adjustment* indicates that neither schedule activated an adjustment. Accurate reports the percentage of these comparisons where the binary comparison choice confirms the CEEI preference report prediction. The *Contradiction* columns report the percentage of binary comparisons that contradicted the CEEI preference reports at each level of preference.

4.4 Discussion

Reporting preferences to the CEEI mechanism was a non-trivial task. Subjects were asked to report a vector of cardinal values — one per course and one per pairwise adjustment — that could be summed to generate an overall value for a schedule with the understanding that CEEI would treat schedules with larger sums as more preferred. Subjects were asked to do this in ten minutes with limited instruction and training. It is

³⁹ The success of those using adjustments could be driven by selection into using the tool, although we find no difference in the rate of contradictions between those subjects who report adjustments to CEEI and those who do not. See Table A2 in the Appendix.

thus unsurprising that there are contradictions between the rank order list created by CEEI and the binary comparisons. Nevertheless, understanding the nature of these contradictions may be useful to help guide implementation, and may also be of independent interest for researchers engaged with the design of preference-reporting languages in other market design settings.

Our analysis yields two lessons for practical implementation. The first is to give students more training to understand the preference-reporting language and more time to complete their reports. The training could focus specifically on how to report cardinal preference information, since it was the comparisons where cardinal preference information was pivotal that were most prone to mistakes, and on the use of adjustments, since this tool was rarely used but where used seemed to enhance accuracy. The second is to explore enhancements to the reporting language to allow students to express common forms of non-additive preferences, such as the elegant schedule preferences we explored in Section 4.2.

In year one of implementation, Wharton acted on the first of these lessons. Students were provided with extensive training sessions regarding the reporting language, with some of the training materials specifically focused on the cardinal preferences issue. The top-ten widget in the preference-reporting user interface, which provides interactive feedback to students about their reported preferences, was enhanced for year one of implementation to allow students to see substantially more than ten schedules. This allowed students to assess whether they have reported accurately not just for their very most preferred schedules (which may be unattainable if the student likes mostly popular courses) but further down their overall ranking as well. Students were given several weeks to think about and then report their preferences, as opposed to the ten minutes they had in the lab. These improvements should help close the gap in the rate of contradictions between the cardinal and ordinal comparisons and push the contradiction rate down overall. For year one, however, Wharton elected not to modify the language itself; there is some discussion about doing so in subsequent years.

Our analysis in this section also yields two takeaways for market design researchers interested in language design more broadly. First, our results provide further evidence of the importance of non-additive preferences in practical market design. To some extent, the importance of non-additive preferences is obvious (e.g., from bidding behavior in spectrum auctions), but empirical evidence on the effect of non-additive preference-reporting languages on mechanism performance is still comparatively scarce (Cantillon and Pesendorfer 2006 and Reguant 2014 are two notable examples). Second, our results provide empirical support for the common intuition in the market design literature that it is easier for participants in a market to report ordinal preference information than cardinal preference information.

V. Qualitative Analysis

As we noted at the beginning of the paper, one of the motivations for running an experiment with real Wharton MBA students was to allow the Wharton deans to see the benefits materialize in their own specific environment (i.e., to take the mechanism for a “test drive”) and to allow them to observe whether there were any issues not captured by the theory that undermined its potential benefits (i.e., to look for “side effects”). For example, a mechanism might have great theoretical properties, but if agents find it frustrating or confusing, they may rebel against using it. Fear that a new market design would lead to disgruntled market participants was particularly prevalent at Wharton, where student satisfaction is a top priority — not just in the allocation of courses but also in the process that leads to those schedules.

To address these concerns, our experimental design collected a wide array of qualitative response data on subjects’ attitudes toward the new mechanism to look for such “side effects.” In this section, we present results from the qualitative questions that subjects answered in response to using each of the two course allocation mechanisms. These results went a long way to convincing the Wharton deans about the advantages of switching to CEEI as a new course allocation mechanism. In addition, the results suggested some improvements for implementation.

Table 12 reports average responses to Likert-scale questions that we asked of subjects about each mechanism. It shows the seven questions in which we found statistically significant differences between responses about CEEI and about the WA (the other eight qualitative Likert-scale questions, listed in Table A3 in the Appendix, yielded no significant differences). Note that higher numbers indicate stronger agreement with the statement.

Table 12: Qualitative Responses about CEEI and WA

Questions about each System	CEEI Average	WA Average	Signrank p-value
I understand how this course allocation system works.	4.83	5.92	$p=0.0000$
I enjoyed participating in this course allocation system.	4.72	4.37	$p=0.0953$
I like this course allocation system.	4.55	4.18	$p=0.0947$
I felt like I had control over my schedule in this course allocation system.	3.95	4.45	$p=0.0730$
This course allocation system is simple.	4.45	3.73	$p=0.0012$
I had to think strategically about what other students would do in this course allocation system.	2.93	6.42	$p=0.0000$
Someone with perfect knowledge of the historical supply and demand for courses could have had an advantage over me in this system.	3.67	6.04	$p=0.0000$

Table 12 shows the seven questions that resulted in statistically significant differences between the two mechanisms. The other eight questions yielded no significant differences at $p>0.1$ for all questions. Questions were rated on a scale of: 1=“Strongly Disagree,” 2=“Disagree,” 3=“Somewhat Disagree,” 4=“Neither Agree or Disagree,” 5=“Somewhat Agree,” 6=“Agree,” 7=“Strongly Agree.” *CEEI Average* and *WA Average* take the mean of the response values across all 132 subjects in the experiment. Higher average numbers display higher agreement with the statement. Since each subject gave an answer for each of the two mechanisms, we can use a non-parametric Signrank Test that responses are equal across the two mechanisms. Since the theory in Budish (2011) does not make a prediction about these qualitative responses, the tests are two-sided.

As can be seen from Table 12, on four of the fifteen questions subjects answered, we see significant differences between the mechanisms with $p<0.01$ with three additional questions that were significantly different at the $p<0.1$ level. The responses to the qualitative response questions tell a story about the benefits of CEEI relative to the WA and the potential improvements in implementation that could be adopted within the CEEI system. We address the relative benefits of CEEI in subsection 5.1 and then address the potential improvements in subsection 5.2. In subsection 5.3 we investigate how preference for CEEI interacted with subject demographics.

5.1 Strategic Simplicity and Student Satisfaction

One of the features of CEEI that appealed to the Wharton deans was the strategic simplicity from the perspective of agents in the system. While in the WA subjects must consider the bids of other agents in the system, CEEI is designed to be (approximately) strategy proof so agents simply have to report their preferences truthfully and can ignore their beliefs about the preferences of other agents. The survey results suggest that subjects understood that the CEEI mechanism did not require strategizing.

The average response to the question that asked subjects' agreement with the statement "I had to think strategically about what other students would do in this course allocation system" was 2.93 for CEEI (i.e., close to "Somewhat Disagree," a 3 on the Likert scale) and was 6.42 for WA (i.e., close to the midpoint between "Agree" and "Strongly Agree"). This difference is highly statistically significant. Similarly, the average response to the question: "Someone with perfect knowledge of the historical supply and demand for courses could have had an advantage over me in this system" was 3.67 for CEEI (i.e., between "Somewhat Disagree" and "Neither Agree or Disagree") and was 6.04 for WA (i.e., close to "Agree"). Again, this difference is highly statistically significant. These were the two questions that elicited the largest difference in response between WA and CEEI.⁴⁰

These sentiments worked their way into subjects' free responses. One subject wrote, "Multiple rounds and the historical price research introduced in the auction make it quite time-consuming and therefore kind of annoying." Another wrote, "Really like the idea of the new system as it removes the inherent 'gaming' aspect of the auction – I'm a believer in free markets but the auction is a disadvantage to those that don't have the time or skill required to fully research and intelligently participate in it."

This strategic simplicity spilled over into subjects' thoughts on the overall simplicity of the course allocation system. In response to question 10, subjects agreed more with the statement "This course allocation system is simple" for CEEI than for the WA (4.45 versus 3.72, $p=0.0012$). One subject wrote, "I loved the matching system! I wound up with the same schedule for both but the matching system took less time + stress. It's more fair, and easier. Definitely a welcome change!!"

⁴⁰ One might be somewhat surprised that the difference between CEEI and WA on these measures is not even larger. One explanation is that at least some of our subjects were reluctant to accept that the CEEI mechanism was not "gameable" like the WA was. One lesson for implementation that came out of these survey responses was to do a more thorough job of explaining this fact to students, since understanding that historical information and strategizing was not necessary for CEEI was positively correlated with other measures of satisfaction with CEEI.

In addition, subjects also appeared somewhat more satisfied with CEEI than with the WA overall, although these results are slightly more suggestive given the results are only significant at the level of $p < 0.1$. Subjects were slightly more likely to agree with “I like this course allocation system” for CEEI than for the WA (4.55 vs. 4.18, $p = 0.0947$) and slightly more likely to agree with “I enjoyed participating in this course allocation system” for CEEI than WA (4.72 vs. 4.37, $p = 0.0953$). These results helped convince the Wharton deans that there was nothing unexpected about the CEEI mechanism that led subjects to dislike the system; that is, there was no unanticipated “side effect” that made CEEI unappealing to our Wharton student subjects.

5.2 CEEI is a Black Box

Despite CEEI scoring higher marks for strategic simplicity and student satisfaction, the qualitative survey results did not unanimously favor CEEI. In particular, subjects reported that they did not understand how CEEI worked and that they felt as though they had less control over their schedule in CEEI than they did in the WA. In particular, subjects reported much higher agreement with “I understand how this course allocation system works” for the WA than for CEEI (5.92 vs. 4.83, difference significant at $p < 0.01$). Similarly, subjects were slightly more likely to agree with “I felt like I had control over my schedule in this course allocation system” for WA than for the CEEI (4.45 vs. 3.95, $p = 0.0730$).

Our interpretation of these results is that subjects felt that the CEEI system was a bit of a “black box” in that subjects submitted preferences and then the mechanism spit out a schedule without any explanation. This sentiment was also reflected in the subjects’ free responses. One subject wrote: “I like the idea of getting the best schedule I could afford, but didn’t like feeling like I wasn’t in control. I would feel helpless if I got a schedule that wasn’t close to what I preferred.” Another wrote: “The course matching system is just a black box where there’s one round and we rely on the computer to make judgments for us.”

This finding has helped guide implementation at Wharton. In particular, students have significantly more training on the CEEI system than in the experiment. In addition, while subjects in the experiment were not shown the market clearing prices of the courses in the market — which prevented them from understanding why they got the specific bundle of courses they received from the mechanism (and why, for example, they failed to get a particular course they might have wanted) — in Wharton’s implementation students are shown the market clearing prices so that they can develop an ex-post understanding for why CEEI assigned them the courses it did.

5.3 Demographics

One additional set of results that arose from the qualitative response data was on the relative preferences between CEEI and the WA for men and women. This result turned out to be important for the Wharton deans who were facing evidence that women at Wharton did not particularly like the WA. A Wharton survey of all second-year students in the year of our experiment found that women reported lower ratings for the effectiveness of the real Wharton Auction than men did (7-point scale of effectiveness, 4.95 for men vs. 4.28 for women, $p < 0.001$).

We found a similar pattern in our data with regard to attitudes toward the auction. In our Likert-scale questions, female subjects reported liking the WA significantly less than male subjects reported liking it (4.51 for men vs. 3.81 for women; t-test, $p = 0.021$). In addition, women reported liking CEEI significantly more than the WA (4.53 for CEEI vs. 3.81 for the WA; paired t-test, $p = 0.027$). Finally, when we compare liking of CEEI, we see that the gender effect is gone (4.56 for men vs. 4.54 for women; t-test, $p = 0.931$).⁴¹

That in our data CEEI was liked more than the WA (directionally for men and significantly for women) and that CEEI eliminated the gender gap were important for the Wharton deans in deciding to switch to CEEI.

VI. Conclusion

Wharton adopted CEEI for use in practice, under the name “Course Match”, beginning in Fall 2013. The following is the excerpt of the “Course Match User Manual” describing the role of the experiment:

In the Fall of 2011, Wharton faculty and staff joined with 132 MBA students and put the Course Match theories to the test. In eight separate sessions students were presented with a list of 25 classes and given an introduction to Course Match. Each student then built two schedules, one using Wharton Auctions, the previous system for course selection, and another using Course Match. With their two schedules complete, the students answered a series of questions based both on their own results and those of their peers. The results were clear. Students were more satisfied with their Course Match schedules than with those generated by the Auction. They were less envious of their peers’ schedules and they found Course Match easier to use even though they received only minimal training on the new system.

⁴¹ There is not a significant interaction going from WA to CEEI between men and women because men also like CEEI slightly more than they like the WA.

The experiment thus worked in the Roth (1986) sense of “whispering in the ears of princes”. Looking back, it is clear that the use of a wide range of both quantitative and qualitative data was a good design decision, especially because no one result was overwhelming, but the results in gestalt strongly supported the claims made by the theory. It was also a good design decision to use Wharton students as our experimental subjects: using an experiment to whisper in the ears of princes is easier when the princes’ subjects are the experimental subjects.

As of the present writing, CEEI has been in use at Wharton for three semesters. Unfortunately, it was not possible to obtain the data that would have been necessary to do a full empirical before-and-after comparison of the two mechanisms.⁴² However, some data that are available are suggestive that CEEI has improved fairness. Specifically, Gerard Cachon, the chair of Wharton’s Course Allocation Redesign Team, computed two measures of equity for WA and CEEI. The first was a Gini index of the distribution of the price of the student’s final schedule: CEEI reduced this inequality measure from 0.53 to 0.31. The second was the distribution of the twenty most popular courses in each year. Under the WA, 31% of students got zero of the top twenty courses, and 6% got three or more, versus 13% and 0%, respectively, under CEEI. That is, under CEEI fewer students got none of the most popular courses, and fewer (zero!) got three or more. While for any one student we cannot tell whether their lack of top twenty courses reflects their preferences or is a source of disappointment, the aggregate distribution suggests that CEEI improved equity.

Evidence from Wharton’s annual student survey is also consistent with claims made by the theory and experiment. At our urging, the student survey added a few questions about course allocation in the last year of the auction’s use. The questions were written in such a way that they could be used again in the first year of CEEI by simply replacing the name of the allocation mechanism. The percentage of students responding either Agree or Strongly Agree to the statement “I was satisfied with my schedule from {the course auction system / course match}” increased from 45% in 2013 (the last year of the WA) to 64% in 2014

⁴² Ideally, we would have used a school-wide survey to obtain true preference from students during the last year of the WA; this would have allowed us to compare student outcomes from actual play of the WA to counterfactual play of CEEI, analogously to the study conducted by Budish and Cantillon (2012). Unfortunately, the Wharton administration did not want to conduct such a survey, fearing that a survey of students’ “true preferences” at the time they were participating in the WA would have been confusing — especially given that a school-wide announcement had been made concerning the adoption of the new, truthful mechanism. Due to the complexity of the equilibrium of the WA — it is not really understood — it is an open question whether it is possible to infer true preferences from strategic play in the absence of such a survey.

(the first year of CEEI), suggesting increased allocative efficiency. The percentage responding either Agree or Strongly Agree for the statement “{The course auction, Course match} allows for a fair allocation of classes” increased from 28% to 65%, suggesting CEEI increased fairness. The percentage of students responding either Effective or Very Effective to the question “Please rate the effectiveness of the {course auction, course match} system” increased from 24% to 53%, suggesting an increase in overall student satisfaction.

The Wharton administration has also expressed to us that they place a value on getting to see the distribution of students’ true preferences for courses; this is data they intend to use to refine curricular offerings. We are unable to quantify this benefit experimentally, but we mention it because it seemed important to the administrators and is a benefit of using a truthful mechanism that has been noted in other market design contexts (see Roth 2008).

We conclude on a broader methodological note. The market design literature has had two main sets of success stories. First is the design of combinatorial auctions for contexts such as spectrum auctions, procurement, and power markets (see Milgrom 2004, 2011; Cramton *et al.* 2006). Second is the design of matching algorithms for contexts such as entry-level labor markets, school choice, and kidney exchange (see Roth 2002, 2008). A distinguishing feature of the present study is that the *participants* in the market are like the typical participants in a matching market design context — ordinary individuals, playing the game a single time — but the *preferences* they are asked to report are the kinds of complex preferences over packages that are more familiar in a combinatorial auction context. Our paper thus represents a new kind of market design success story, and shows that market designs with combinatorial aspects can be successful even when the participants are ordinary individuals playing the game a single time.

VII. References

- Arnsperger, Christian. 1994. "Envy-Freeness and Distributive Justice." *Journal of Economic Surveys*, 8 (2): 155-186.
- Bogomolnaia, Anna, and Herve' Moulin. 2001. "A New Solution to the Random Assignment Problem." *Journal of Economic Theory*, 100: 295-328.
- Budish, Eric. 2011. "The Combinatorial Assignment Problem: Approximate Competitive Equilibrium from Equal Incomes." *Journal of Political Economy*, 119 (6): 1061-1103.
- Budish, Eric and Estelle Cantillon. 2012. "The Multi-Unit Assignment Problem: Theory and Evidence from Course Allocation at Harvard." *American Economic Review*, 102 (5): 2237-2271.
- Budish, Eric, Abe Othman and Tuomas Sandholm. 2010. "Finding Approximate Competitive Equilibria: Efficient and Fair Course Allocation." Proc. of 9th Int. Conf. on Autonomous Agents and Multiagent Systems.
- Calsamiglia, Caterina, Guillaume Haeringer, and Flip Klijn. 2010. "Constrained school choice: An experimental study." *American Economic Review*, 100(4): 1860-1874.
- Cantillon, Estelle and Martin Pesendorfer. 2006. "Combination Bidding in Multi-Unit Auctions." Working Paper.
- Chen, Yan and Tayfun Sönmez. 2006. "School Choice: An Experimental Study." *Journal of Economic Theory*, 127(1), 202-231.
- Cramton, Peter, Yoav Shoham and Richard Steinberg (eds.). 2006. *Combinatorial Auctions*. Cambridge, MA: MIT Press.
- Ehlers, Lars, and Bettina Klaus. 2003. "Coalitional Strategy-Proof and Resource-Monotonic Solutions for Multiple Assignment Problems." *Social Choice and Welfare*, 21: 265–280.
- Featherstone, Clayton and Eric Mayefsky. 2014. "Why Do Some Clearinghouses Yield Stable Outcomes? Experimental Evidence on Out-of-Equilibrium Truth-Telling." Working Paper.
- Featherstone, Clayton and Muriel Niederle. 2014. "Improving on Strategy-proof School Choice Mechanisms: An Experimental Investigation." Working Paper.
- Foley, Duncan. 1967. "Resource Allocation and the Public Sector." *Yale Econ. Essays*, 7: 45-98.

- Fragiadakis, Daniel E. and Peter Troyan. 2014. "Improving Welfare in Assignment Problems: an Experimental Investigation." Working Paper.
- Hatfield, John William. 2009. "Strategy-Proof, Efficient, and Nonbossy Quota Allocations." *Social Choice and Welfare*, 33 (3): 505–515.
- Kagel, John H., Yuanchuan Lien, and Paul Milgrom. 2010. "Ascending prices and package bidding: A theoretical and experimental analysis." *American Economic Journal: Microeconomics*, 2(3): 160-185.
- Kagel, John H., and Alvin E. Roth. 2000. "The dynamics of reorganization in matching markets: A laboratory experiment motivated by a natural experiment." *Quarterly Journal of Economics*, 115(1): 201-235.
- Kessler, Judd B., and Alvin E. Roth. 2012. "Organ Allocation Policy and the Decision to Donate." *American Economic Review*, 102 (5): 2018–2047.
- Kessler, Judd B., and Alvin E. Roth. 2014. "Loopholes Undermine Donation: An Experiment Motivated by an Organ Donation Priority Loophole in Israel." *Journal of Public Economics*, 114 (2014) 19–28.
- Kojima, Fuhito. 2009. "Random assignment of multiple indivisible objects." *Mathematical Social Sciences*, 57 (1): 134-142.
- Krishna, Aradhna and Utku Ünver. 2008. "Improving the Efficiency of Course Bidding at Business Schools: Field and Laboratory Studies." *Management Science*, 27 (2): 262–282.
- Milgrom, Paul. 2004. *Putting Auction Theory to Work*. Cambridge, UK: Cambridge University Press.
- Milgrom, Paul. 2009. "Assignment Messages and Exchanges." *American Economic Journal: Microeconomics*, 1(2): 95-113.
- Milgrom, Paul. 2011. "Critical Issues in the Practice of Market Design." *Economic Inquiry*, 49: 311–320.
- Moulin, Hervé. 1995. *Cooperative Microeconomics*. London: Prentice Hall.
- Nisan, Noam. 2006. "Bidding Languages for Combinatorial Auctions." In Peter Cramton *et al.*, eds. *Combinatorial Auctions*, MIT Press, 215-232.
- Pápai, Szilvia. 2001. "Strategyproof and Nonbossy Multiple Assignments." *Journal of Public Economic Theory*, 3 (3): 257–271.

- Pathak, Parag A. and Tayfun Sönmez. 2008. "Leveling the Playing Field: Sincere and Sophisticated Players in the Boston Mechanism." *American Economic Review*, 98 (4): 1636-1652.
- Pais, Joana, and Ágnes Pintér. 2008. "School choice and information: An experimental study on matching mechanisms." *Games and Economic Behavior*, 64(1): 303-328.
- Reguant, Mar, 2014. "Complementary Bidding Mechanisms and Startup Costs in Electricity Markets. *Review of Economic Studies*, forthcoming.
- Roth, Alvin E. 1986. "Laboratory Experimentation in Economics." *Economics and Philosophy*, 2 (2): 245-273.
- Roth, Alvin E. 2002. "The Economist as Engineer." *Econometrica*, 70(4), 1341-1378.
- Roth, Alvin E. 2008. "What Have We Learned from Market Design?" *The Economic Journal*, 118: 285–310.
- Sönmez, Tayfun and Utku Ünver. 2003. "Course Bidding at Business Schools." Working paper. <http://ideas.repec.org/p/wpa/wuwpga/0306001.html>.
- . 2010. "Course Bidding at Business Schools." *International Economic Review*, 51 (1):99–123.
- . 2011. "Matching, Allocation and Exchange of Discrete Resources." In *Handbook of Social Economics*, vol. 1A, edited by Jess Benhabib, Alberto Bisin, and Matthew O. Jackson, 781–852. Amsterdam: North-Holland.
- Wharton "Course Match User Manual"
https://spike.wharton.upenn.edu/mbaprogram/course_match/Course%20Match%20Manual%20-%20v1_4.pdf

APPENDIX

Appendix A: Study Instructions

Study Instructions

Thank you for participating in this study.

If you have a question about the study at any time, please raise your hand.

In this study you will be constructing hypothetical class schedules for the spring semester of your second year at Wharton.

You will construct a schedule twice, once under each of two different course allocation systems.

One course allocation system is a simplified version of Wharton's current MBA "Course Auction". The other is an alternative course allocation system for Wharton MBA courses called the "Course Matching System".

Half the sessions in the study will use the "Course Auction" first and half will use the "Course Matching System" first.

After you construct a schedule under each system, you will answer a series of questions about the schedule you have constructed and about the system that you used.

After you have constructed schedules under both systems, you will be asked to compare around 15 to 20 pairs of schedules. For each pair of schedules you will be asked which of the two you prefer.

While using each system, please imagine that it is the spring term of your second year at Wharton, so this will be your last chance to take Wharton classes. Please try to construct your most preferred schedule given the courses that are available.

We are using a subset of 25 spring semester course sections. These course sections were selected to be representative in terms of scheduling, department, and popularity level.

There may be some courses that you would be interested in taking that are not included on this list. There is a limited set of courses because there are only approximately 18 students in the study today and so we cannot replicate the entire course offerings of a normal spring semester. (Note that the actual roster for this spring may differ in terms of which courses are offered, the professors teaching them, and their meeting times.)

We ask you to imagine that these are the only courses available in the spring semester of your second year at Wharton, and to construct your most preferred schedule given these

courses. Since this is your last semester, any budget points that you do not use are worthless.

Please imagine that you do not need to take any particular courses for your major or any other graduation requirements, but that you do need to take 5 credit units. If you have already taken one of the courses in the sample, then you should assume that you cannot take the course again in the spring semester. On the other hand, you should assume that you can take any course in the sample that you have not already taken, that is, ignore any prerequisite requirements. Notice that all of the courses are semester length and worth one credit unit.

Imagine that this is the schedule you would construct the week before classes begin. Once classes start you would be able to drop a course, but you would have to replace it with a course that had an open seat.

In real life, we know you take these decisions very seriously. We ask that you take the decisions in this session seriously as well. We will provide you with time to think carefully while using each system.

Note: Neither the schedules you construct nor the decisions you make in this experiment will have any impact on your actual spring semester courses or your point budget in the actual Wharton MBA Course Auction.

The course sections that are available are listed in the packet that has been given to you. Please take five minutes to look through the packet of courses that are available. Think about how interested you are in each of the courses and what would be your ideal schedule or schedules. We will begin with the first system in five minutes.

Instructions for the Course Auction

This procedure is a simplified version of Wharton's current MBA Course Auction. It is similar to the Course Auction that you have already used during your time at Wharton, but with a few differences:

- Every student starts with the same number of budget points (5,000)
- There are 4 rounds of auction activity
- All students are considered second-year students bidding on courses for their last semester
- All students need 5 credit units (CUs)

You are given a budget of 5,000 points. There are then 4 rounds of the auction, all of which we will play today. In the first round you can bid on as many courses as you would like so long as the sum of your bids is less than or equal to your budget. In the next three rounds, you can buy and sell courses with other students.

Instructions for Round 1

Submitting Bids

In the first round, you can submit bids for as many different course sections as you like. The sum of your bids cannot exceed your budget of 5,000 points.

How are prices calculated?

Prices are calculated the same way as in the current Wharton Course Auction. The price of a section is set at the highest losing bid or 100 points, whichever is higher. For example, if a section has 5 seats, the price for the section is set equal to the sixth highest bid for it, if that bid is at least 100 points, otherwise the price is 100. For example, if the sixth highest bid is 120, then the five highest bidders would each get a seat and be charged 120 points. If fewer students bid for a section than it has seats, then the price of the section is set to 100.

What sections do I get?

You get any section for which your bid is greater than or equal to the price. In the event of a tie, where two or more students submit exactly the same bid and there is not enough space for all of them, the computer randomly assigns the available seats to students who bid that amount.

What happens to my budget?

For each section that you receive, your budget will be decreased by the price of the section. For example, if you bid 1000 for the only section of Course A and its price is 400, then you will receive a seat in Course A, and your budget will be decreased by 400 points. If you do not get a seat in the course then you will not give up those 400 points.

Instructions for Rounds 2, 3, and 4

Submitting Bids and Asks

In Rounds 2 through 4, you can submit bids for as many different sections as you like, just as in Round 1. You can also submit asks, which are offers to sell, for any section that you currently have. The sum of your bids cannot exceed your current budget. You can ask whatever amount you like.

How are prices calculated?

For any section where there are both bids and asks, a trading price is set if there is at least one bid higher than the lowest ask. When this is the case, the computer sets a price to make as many trades as possible. This involves finding a price such that the number of bids higher than that price is the same as the number of asks lower than that price.

Suppose the following bids and asks are submitted for a section during a round.

Bids: 101, 323, 143, 103, 187, 280, 156, and 152.

Asks: 225, 64, 298, 171, and 0.

To see which bids and asks are successful and what the clearing price is, first arrange all the bids in descending order and the asks in ascending order as shown in the table below:

Only three trades can take place at a uniform price. →

Bids	Asks
323	0
280	64
187	171
156	225
152	298
143	
103	
101	

Since only the top three bids are higher than the three lowest asks (and the fourth highest bid is lower than the fourth lowest ask), only three trades can go through. The clearing price is determined as the larger of the first losing bid and the highest winning ask; in this case, the first losing bid is 156, and highest winning ask is 171 — hence the clearing price is 171. The clearing price amount is transferred from each of the successful bidders

to each successful seller (the accounts of unsuccessful bidders and sellers remain unaffected).

If there are extra seats in a section, for example if a section does not reach capacity in Round 1, then those seats are treated as if they are being offered for an ask of 100 points.

You can always be guaranteed to drop a section by submitting an ask of “0”.

What should my schedule look like at the end of Round 4?

At the end of Round 4 you should have: (1) no more than 5 credit units in your schedule; (2) no sections that have a time conflict with each other; and (3) no more than one section in each course.

Is my schedule after Round 4 my final schedule?

Not necessarily. Recall, you should imagine that this is the schedule you would construct the week before classes begin. Once classes start you would be able to drop a course, but you would have to replace it with a course that had an open seat.

If you have any questions, please raise your hand.

Instructions for Between Systems

You have just constructed a schedule under the first system and answered some questions about the schedule and the system. You will now construct a schedule under the other system.

You are constructing a schedule in this system starting “from scratch” such that the decisions you and the other students in this session made while using the first system do not affect anything about activity in this system.

You should again construct the best schedule you can for your spring term of your second year at Wharton. The same course sections are available for this system as were available for the last one.

Instructions for the Course Matching System

The Course Matching System is different from the Wharton Course Auction with which you may be familiar.

The Course Matching System works differently from an auction in that you do not directly bid for course sections. Instead, the computer acts as your agent to buy the best schedule of courses you can afford.

Your job is to tell the computer how much you value individual course sections and whether you assign extra value (or negative value) to having certain course sections together. This process will be explained in detail below.

Since you can tell the computer how much you like every course or pair of courses that might be in your schedule, the Course Matching System only needs one round. In that round, the computer will use your preferences to buy you the best schedule you can afford.

Since the computer is going to optimally buy courses for you, your job is to provide the computer with all the information it needs about how much you value the courses. This is obviously very important, since the computer is going to buy the optimal schedule for you given only what it knows about how you value courses.

The way to communicate your values to the computer is as follows:

1) You tell the computer how much you value each course section that you have any interest in taking.

- First, you pick a favorite course section and assign it a value of 100.
- Second, you assign all other course sections that you have any interest in taking a value between 1 and 100.

The reason that you assign your favorite course section a value of 100 and all other sections a number between 1 and 100 is that all values are relative.

For example, if you value every course at 100 then you are telling the computer that you value all courses equally. If you value one course at 100 and another course at 50, you are telling the computer you value the course at 100 twice as much as the course at 50.

Unlike using other course allocation systems, when using the Course Matching System, you do not need to think about what other people are doing. All you need to do is communicate how you value course sections to the computer so it knows how to make tradeoffs for you.

How does assigning value to courses work?

Suppose that among the many course sections you assign a positive value, you tell the computer the following values for the single section courses A through E:

Course A = 100

Course B = 80

Course C = 60

Course D = 15

Course E = 10

This tells the computer that you are particularly interested in Courses A, B and C, and somewhat interested in Courses D and E. In particular, it tells the computer that you prefer getting Courses A, B, and C ($100 + 80 + 60 = 240$) than getting Courses A, D, and E ($100 + 15 + 10 = 125$).

It also tells the computer that you prefer getting Courses B and C ($80 + 60 = 140$) than Courses A, D, and E, which only sum to 125. For any two schedules, the computer thinks you prefer whichever schedule has a larger sum.

For simplicity, this example valued only 5 course sections. You should list a positive value for as many courses that you have any interest in taking. We recommend that you assign a positive value to at least 12 course sections. This way the computer can distinguish between a section that has low positive value to you and a section that has zero value to you.

Can I assign values for multiple sections of the same course?

Yes, and you will probably want to do this. To explain, suppose three sections of a course are offered, all on Mondays and Wednesdays. Professor Smith teaches the 10:30-12:00 and 12:00-1:30 sections while Professor Jones teaches the 3:00-4:30 section. You may assign values of 90, 80 and 15 to these three sections, respectively, to signify that you greatly prefer Professor Smith to Professor Jones, and slightly prefer 10:30 to 12:00. Because you can only take one section of a course, you will be assigned at most one of these three course sections, even though you entered values for all three.

Again, there is no limit to the number of course sections that you may assign a positive value.

2) You tell the computer if you assign extra (or negative) value to certain pairs of classes.

To do this, you check the boxes next to any two sections and indicate an extra positive or negative value to having both sections together. These “adjustments” are shown at the top of the page of your valuations.

Why might I assign extra value to two courses together?

Some students might get extra value from having two courses that are back-to-back in their schedule (e.g. they do not like breaks between classes).

Some students might get extra value from having two courses that are related in their schedule (e.g. they might get extra value from taking two courses from the same department if each one becomes more useful with the other).

You can think of these courses as complements, i.e. the combination of the two courses together is greater in value than the sum of their values.

How does assigning extra value work?

Suppose you specify the following values for single section courses A through C:

Course A = 40

Course B = 30

Course C = 85

And suppose you assign an extra value of 20 for getting Course A and Course B together.

Then you are telling the computer that getting Course A and Course B together in your schedule has a value of 90 (90 = 40 for Course A + 30 for Course B + 20 for getting both together).

This means that the computer would try to get you Course A and Course B together before trying to get you Course C. If you had not assigned the extra value to Courses A and B together, the computer would have tried to get you Course C before trying to get you Courses A and B.

Why might I assign negative value to two courses together?

Some students might get negative value from having two courses that are back-to-back in their schedule (e.g. they prefer to take breaks between classes).

Some students might get negative value from having two courses that are related in their schedule (e.g. they might decide that they only want to take one class from a certain department).

You can think of these courses as substitutes, i.e. the second course is worth less when you already have the first.

How does assigning negative value work?

Suppose you specify the following values for single section courses A through C:

Course A = 40

Course B = 30

Course C = 55

And suppose you assign a negative value of -20 for getting Course A and Course B together.

Then you are telling the computer that getting Course A and Course B together in your schedule has a value of 50 ($50 = 40$ for Course A + 30 for Course B - 20 for getting both together).

This means that the computer would try to get you Course C before getting you Course A and B together. If you had not assigned the negative value to Courses A and B together, the computer would have tried to get you Courses A and B before trying to get you Course C.

You can also use an adjustment to tell the computer “I want to take at most one of these two courses”. Using the example above, suppose you want to take either Course A or Course B, but you absolutely do not want to take both. Then you should assign a negative value of -70 for Course A and B together. That negative adjustment tells the computer that the combination has value 0 to you ($0 = 40$ for Course A + 30 for Course B - 70 for getting both together). Therefore, you may get Course A or Course B, but the computer will never get both for you.

When do I not need to enter in an adjustment?

You do not need to enter an adjustment when two sections are from the same course or two sections are offered at the same time. The computer already knows that you cannot take these sections together. For example, if Professor Baker teaches two sections of the same course, one from 9:00-10:30 and the other from 10:30-12:00, then you can assign a positive value for each of them, but you don't need to assign a positive or negative adjustment for the combination.

Once the computer knows how much you value each course section, it will buy the best schedule you can afford.

How do I know that I am reporting my values right?

To help make sure you are reporting your values right, you can click a button on the navigation bar to see your top 10 schedules. Given the values you reported, the computer thinks that these are your 10 favorite schedules, ranked in order. This means that the computer will try to buy you these schedules in this order. If the order of these schedules does not look right to you, go back and adjust your values until they appear in the right order.

What is my budget that the computer will use to buy courses for me?

Each student is given a budget of 5,000 points.

How are prices determined?

The Course Matching System sets prices based on demand for the courses so that demand equals supply. Courses that are more highly demanded get higher prices and courses that are less popular get lower prices or prices of zero.

One way to think about how prices are set is that each student's computer asks for the best possible schedule for its student. When everyone has their best possible schedule, some courses will have too many students. The price of those courses will rise. Then, given the new set of prices, each student's computer asks again for the best possible schedule for its student at the new set of prices. Some courses will be undersubscribed or oversubscribed and prices will adjust again. This process repeats until there is a set of prices where all popular courses are full and every student gets their best possible schedule given those prices.

Given the set of prices, it may be necessary to break a tie between two or more students who want a course section. These potential ties are broken by assigning a randomly selected small budget increase to each student.

Shouldn't the values I report to the computer depend on the prices of courses or other student's values?

No! The Course Matching System is designed so you do not need to think about the prices of the courses or the values that other students assign to courses. You get the best schedule possible simply by telling the computer your true values for courses.

To see this, notice that if your favorite course, to which you assign a value of 100, is a course whose demand is less than the number of available seats, then it will have a price of zero and you will get that course without using any of your budget. The computer can

then use the remainder of your budget to try to get the other course sections that you value highly.

Another way to think about reporting your values to the computer is to imagine you are sending the computer to the supermarket with your food budget and a list of your preferences for ingredients for dinner. You want to report your true values so that the computer can make the right tradeoffs for you when it gets to the supermarket and observes the actual prices for each ingredient.

Are my values equivalent to “bids”?

No! As mentioned above your values are only compared to each other and never compared with other students’ values.

Is the schedule I receive after I report my values my final schedule?

Not necessarily. Recall, you should imagine that this is the schedule you would construct the week before classes begin. Once classes start you would be able to drop a course, but you would have to replace it with a course that had an open seat.

If you have any questions, please raise your hand.

Please use this page to write any additional comments about your experience during this session. These are anonymous comments, so please do not include your name.

Appendix B: List of Course Sections Available in Experiment

Course	Title	Instructor	Day Code	Start Time	Stop Time	Available Seats
ACCT742	PROBLEMS IN FIN REPORTIN	LAMBERT R	MW	0130PM	0300PM	5
ACCT897	TAXES AND BUS STRATEGY	BLOUIN J	MW	1200PM	0130PM	4
FNCE726	ADVANCED CORP FINANCE	VAN WESEP,E	TR	1200PM	0130PM	5
FNCE728	CORPORATE VALUATION	CICHELO M	MW	0300PM	0430PM	4
FNCE750	VENT CAP & FNCE INNOVAT	WESSELS D	MW	0130PM	0300PM	4
FNCE750	VENT CAP & FNCE INNOVAT	WESSELS D	MW	0300PM	0430PM	4
FNCE891	CORPORATE RESTRUCTURING	JENKINS M	TR	0130PM	0300PM	4
LGST806	NEGOTIATIONS	DIAMOND S	R	0300PM	0600PM	3
LGST806	NEGOTIATIONS	BRANDT A	W	0300PM	0600PM	3
LGST809	SPORTS BUSINESS MGMT	ROSNER S	TR	0300PM	0430PM	5
LGST813	LEG ASP ENTREPRENRSH	BORGHESE R	M	0300PM	0600PM	5
MGMT691	NEGOTIATIONS	MUELLER J	TR	1030AM	1200PM	3
MGMT721	CORP DEV: MERG & ACQUIS	CHAUDHURI S	TR	0900AM	1030AM	4
MGMT721	CORP DEV: MERG & ACQUIS	CHAUDHURI S	TR	1030AM	1200PM	4
MGMT782	STRATEGIC IMPLEMENTATION	MURMANN J	TR	1200PM	0130PM	5
MGMT833	STRAT & PRAC OF FAMILY	ALEXANDER W	TR	0130PM	0300PM	4
MKTG756	MARKETING RESEARCH	IYENGAR R	MW	1030AM	1200PM	5
MKTG773	CUSTOMER BEHAVIOR	REED A	TR	1030AM	1200PM	5
MKTG776	APPL PROB MODELS MKTG	FADER P	W	0300PM	0600PM	5
MKTG778	STRATEGIC BRAND MGMT	MOGILNER C	TR	0130PM	0300PM	5
OPIM690	MANAG DECSN MAKING	MILKMAN K	MW	0130PM	0300PM	5
OPIM692	ADV TOPICS NEGOTIATION	SCHWEITZER M	TR	0130PM	0300PM	4
REAL721	REAL ESTATE INVESTMENTS	FERREIRA F	MW	0130PM	0300PM	4
REAL721	REAL ESTATE INVESTMENTS	WONG M	TR	0130PM	0300PM	4
REAL821	REAL ESTATE DEVELOPMENT	NAKAHARA A	W	0300PM	0600PM	4

[Note: The available seats were selected to create scarcity with twenty students participating in a session. When more than fifteen and fewer than eighteen subjects showed up to a session, we turned the five-seat courses into four-seat courses. When fewer than sixteen subjects showed up to a session, we turned the five-seat courses into four-seat courses and the four-seat courses into three-seat courses. This allowed there to be scarcity even when there were fewer than twenty subjects in a session.]

ACCT742: PROBLEMS IN FIN REPORTIN - LAMBERT R

Financial statements are a primary means for firms to communicate information about their performance and strategy to investors and other groups. In the wake of numerous accounting scandals and the recent financial meltdown (which accounting both helped and hindered), it is more important than ever for managers and investors to understand (i) the financial reporting process, (ii) what financial statements do and do not contain, and (iii) the types of discretion managers have in presenting transactions they have undertaken. This course is designed to help you become a more informed user of accounting numbers by increasing your ability to extract, interpret, and analyze information in financial statements.

While this is not a course in equity valuation per se, equity valuation is one of the most common uses of financial statement data. Accordingly, we will examine the relation between Accounting 742 -stock prices and financial statement information. We will also study the use of financial ratios and forecasted financial statement data in models of distress prediction.

ACCT897: TAXES AND BUS STRATEGY - BLOUIN J

Traditional finance and strategy courses do not consider the role of taxes. Similarly, traditional tax courses often ignore the richness of the decision context in which tax factors operate. The objective of this course is to develop a framework for understanding how taxes affect business decisions.

Part of being financially literate is a having a basic understanding of how taxation affects business decisions that companies typically face: forming the business and raising capital, operating the firm, distributing cash to shareholders through dividends and share repurchases, expanding through acquisition, divesting lines of business, and expanding internationally. Taxes have a direct impact on cash flow and often divert 40% to 50% of the firm's pretax cash flow to the government. Having an understanding of taxation and how firms plan accordingly is important whether you will be running the firm (e.g., executive in large company, entrepreneur, or running a family owned business) or assessing it from the outside (e.g., financial analyst, venture capitalist, or investment banker). Taxes are everywhere and it pays to have some understanding of them.

FNCE726: ADVANCED CORP FINANCE - VAN WESEP,E

The objective of this course is to teach students how to apply modern financial theory to the investment and financing decisions of a modern corporation. The course is designed for finance majors who will pursue careers in major corporations, the financial sector, and consulting firms. The core theory introduced in the Financial Analysis course is extended for applications to the strategic financial decision areas of a firm.

The theme of this course is value-based management. Financial theory explains the real world using abstract and simplified models. Such conceptual models are often not sufficiently rich for

dealing with all the complexities of the real world. Financial decisions based on rigorous theory and models are superior to ad hoc alternatives. Concepts and techniques introduced in this course should help you express key decisions in terms of their impact on firm value.

The first four sections of the course focus on a range of financial issues that confront managers in their ordinary course of doing business; mainly financial planning, capital budgeting, and the interaction between investment and financing decisions. The options approach to investment decisions is in Section IV. Section V deals with dividend policy. The last two sections of the course focus on financial distress and corporate restructuring, mergers, and acquisitions.

FNCE728: CORPORATE VALUATION - CICHELO M

The objective of this course is to teach students about the analysis and valuation of equity securities. In learning the primary valuation techniques used to estimate market values for equity securities, we will pay special attention to financial statement analysis. Additionally, the course will highlight the importance of identifying and focusing on key value drivers. The analytical framework and valuation techniques, as well as the practical market information students learn in this class will be useful for careers in corporate finance, asset management, research, sales, trading, financial market regulation or financial journalism.

The course will be segmented into four major sections:

- (1) During the first weeks of class we discuss the drivers of corporate value, specifically return on investment and organic revenue growth. We next examine how to build an ROIC-based valuation model and how this differs from and complements the traditional discounted cash flow model.
- (2) The second section covers financial analysis using data from the annual report. We start with the traditional competitive benchmarking and next move to current metrics such as return on invested capital (ROIC) and economic profit. Our primary goal will be to build a true understanding of operating performance across business units and for the entire company.
- (3) In the third section, we build an integrated valuation model using discounted cash flow. The section starts with the fundamentals of forecasting, how to determine the appropriate forecast period, and issues related to continuing value. We derive the weighted average cost of capital, focusing on how to estimate the inputs.
- (4) In the final section, we discuss alternatives to DCF valuation, comparables analysis and options. We use multiples analysis to triangulate our DCF valuation and options analysis to handle uncertainty.

FNCE750: VENT CAP & FNCE INNOVAT - WESSELS D

This course will focus on the primary activities performed by venture capital professionals, including how they raise capital, how they structure their funds, as well as how they select, fund, and exit high-growth privately-held companies.

- We start by outlining how venture capital funds are organized, how and from whom they raise capital, and in what type of firms they typically invest. A good portion of the section

will examine the risk return profile of venture capital and whether venture capital “beats the market,” whether it lowers risk of the limited partner’s portfolio, and how the great uncertainty associated with growth firms should be priced.

- The second section of the course will examine corporate valuation and value drivers. Given the incredible uncertainty associated with high-growth companies, alternative methods such as key value driver models and comparable transactions must be employed to triangulate results.
- Special attention will be given to the valuation process for small, illiquid, high-growth companies versus mature companies. For instance, how do you bound reasonable estimates of revenue growth, operating margins, and capital productivity when little historical data is available?
- The third section of the course will examine valuation techniques necessary to value complex securities associated with venture capital and high growth companies. Preferred stock held by venture capitalist has conversion features that resemble a combination of debt and equity. Therefore, options models must be employed to determine their economic (versus fully-diluted) value.

FNCE891: CORPORATE RESTRUCTURING - JENKINS M

The objective of this course is to familiarize students with the financial, legal, and strategic issues associated with the corporate restructuring process. The main focus of the course will be on restructuring financially distressed firms. We’ll begin by reviewing the financial instruments commonly used by risky firms (leveraged loans and high-yield bonds) and learn to interpret the contracts that govern them (credit agreements and bond indentures). We’ll then survey a variety of restructuring methods (out-of-court workouts, exchange offers, prepackaged and pre-negotiated bankruptcies, distressed asset sales, and Chapter 11 reorganizations) available to troubled firms and study the dynamics of the restructuring process through a number of historical and current case studies. Finally, we’ll consider distressed debt as an asset class and develop techniques for investing in distressed securities.

The course will provide students with tools to value distressed companies, understand the legal framework governing bankruptcy and reorganization in the U.S. and other countries, and navigate the key strategic issues facing managers and investors in distressed companies. It will also provide students with a specialized vocabulary and important facts about the restructuring industry, distress investing, and leveraged financial markets.

LGST806: NEGOTIATIONS - BRANDT A

Course Objectives: The aim of the course is, using case study and my practical background in negotiating in the sports world for 25 years, to educate, assist and, through practice, enable you to become a more effective negotiator and conflict resolver. This effectiveness in negotiating and conflict resolution requires many things, including:

- The understanding that you should not change and be something you are not; simply use tools effectively as who you are;

- The creativity to execute deals that others might overlook;
- Knowing when to walk away;
- The insight to recognize ethical traps – and the wisdom to avoid them;
- Understanding the importance of relationships;
- The ability to work with people whose backgrounds, expectations, culture, and values differ from your own;
- The ability to resolve conflicts; and
- The capacity to reflect and learn from your experience.

This course links both the science and art of negotiation and conflict resolution, but it is more art than science. It will give you the opportunity to identify your strengths as a negotiator and to work on your relative weaknesses. More fundamentally, the course will provide both a conceptual framework to diagnose problems and promote agreement in a range of settings from your professional to your personal life.

LGST806: NEGOTIATIONS - DIAMOND S

- Class 1: Ratings War Case - Trust
- Class 2: Pink Cadillac Case
- Class 3: The Diva Case - Interest Based Negotiation
- Class 4: Pheasant Egg - Hard Bargainers
- Class 5: The Warranty
- Class 6: Sell Phones Case – Coalition Building
- Class 7: Problem Solving Session
- Class 8: Breath & Taxes Case
- Class 9: Mediation Case
- Class 10: El Camino Real Case – Group Processes
- Class 11: Family Business Case – Real World M&A
- Class 12: Alpha Beta Case – Cross Cultural Issues
- Class 13: Course Wrap Up

LGST809: SPORTS BUSINESS MGMT - ROSNER S

This course examines various business disciplines as they apply to the sports industry. The course provides the student with an overview of the business of the intercollegiate, Olympic and professional sports enterprises. In addition, the course investigates the business related issues encountered by managers of sport organizations and covers how business principles can be applied to effectively address these issues.

LGST813: LEG ASP ENTREPRENRSHP - BORGHESE R

Legal and Transactional Aspects of Entrepreneurship is a practical and intensive course that examines the critical legal issues confronting start-up and emerging growth companies. The

course provides perspective on how to use the law strategically to manage risk, deploy resources and maximize shareholder value. Topics include the enforceability of confidentiality, non-competition and other restrictive covenants in employment agreements, choice of business form including the legal, financial and tax advantages and disadvantages of general partnerships, limited partnerships, corporations and limited liability companies, tax and securities law aspects of raising capital, structuring venture capital and private equity financings, letters of intent and mergers and acquisitions, employment law, and intellectual property law including trade secrets, copyrights, patents and trademarks.

MGMT691: NEGOTIATIONS - MUELLER J

We negotiate daily with potential employers, co-workers, bosses, landlords, merchants, service providers, partners, parents, children, friends, roommates, and many others. Our negotiation skills affect what price we will pay, the amount of our salary and compensation, what movie we watch, and who will clean up the kitchen. Why do we sometimes get our way, while at other times walk away frustrated by our inability to achieve the agreement and resolution we want?

Negotiation is the art and science of securing agreements and resolving disputes between two or more interdependent parties. The purpose of this course is to help you develop expertise in managing negotiations that occur in a variety of business settings. It is designed to be relevant to a broad spectrum of problems faced by managers. As a manager, you not only need analytical skills to discover optimal solutions to problems, but also good negotiation skills to get these solutions accepted and implemented.

The learning method is experiential. You will prepare for and engage in a variety of negotiation exercises (individually, and as a team). The objective is to explore your talents, skills, shortcomings, and strengths as a negotiator in a safe setting, to learn about yourself and how you respond in specific situations. If you discover a tendency that you think needs to be changed, this is the place to try something new. The course is sequenced so that cumulative knowledge can be applied and practiced. The skill set you develop here will serve you in both your personal and professional life.

MGMT721: CORP DEV: MERG & ACQUIS - CHAUDHURI S

As product and factor markets globalize, technology rapidly evolves, and competition intensifies, companies worldwide are fundamentally changing their structures and processes to keep pace and take advantage of new opportunities. We are witnessing the emergence of the disaggregated and distributed global firm that leverages internal and external capabilities around the world in real time, blurring traditional organizational boundaries and leading to the creation of virtual enterprises.

To catalyze this transformation and stay ahead in competency, cost, and time to market, companies are utilizing an array of powerful but often risky inorganic strategies, in the hopes that

they can “plug and play” with local and global entities, and gain the needed resources to compete effectively.

This course explores the various modes of corporate development available to managers to drive firm growth and change, including alliances, outsourcing, corporate venturing, and particularly mergers and acquisitions. The objectives are three-fold: (1) to arm you with a set of tools to facilitate the selection of the appropriate growth strategy in a given situation; (2) to provide you with insights as to how to manage partnerships like alliances, outsourcing, and corporate venturing; and, (3) to develop a comprehensive framework for executing M&As, from initiation to implementation.

The emphasis is on strategic and operational aspects of these transactions, rather than financial considerations. While we will cover deals from a variety of industries, a number of them are from technology-based sectors. This is not only due to the recent prevalence and continued importance of external growth strategies in these sectors, but also because the fast pace provides early assessments of outcomes and management lessons. As we will see, insights from these settings are generalizable to many other contexts.

MGMT782: STRATEGIC IMPLEMENTATION - MURMANN J

As the difficult process of strategic planning evolves and moves toward definition, the even more difficult process of strategy implementation or execution comes into play. The two – strategy formulation and execution – are separate, but highly interdependent, and both are critical to strategic success.

This course focuses on strategy implementation or execution, with emphasis on the decisions, actions, and conditions that facilitate the successful attainment of strategic objectives. The need for this course is explained by the fact that much more is known about strategy formulation and planning than the implementation or execution of strategy. Valid, logical strategic plans often fail to achieve their potential because of implementation problems and related shortcomings. Hence, this course focuses on implementation or execution as a valid academic and practical concern of managers.

Strategy implementation or execution is often misunderstood. It’s “getting your hands dirty” and getting things done, but it’s more than this. It involves macro issues as well as micro issues. It involves conceptual thinking as well as routine elements. It integrates long- and short-term decisions and actions. In short, implementation or execution covers a great deal and integrates much of what you’ll handle in future managerial jobs.

MGMT833: STRAT & PRAC OF FAMILY - ALEXANDER W

Family-controlled private and public companies are the dominant form of enterprise worldwide, comprising more than 90% of all businesses. They are currently undergoing intense competitive transition in form and function and more than three trillion dollars of assets will change generational management during the next ten years. This course is designed for those persons who desire to understand the distinct strategies and practices of family-controlled companies and family wealth management. It will focus on shareholder decision-making; financial and market-driven options for long-run competitiveness, organizational structures, and management team issues; strategic planning from a resource-based perspective; transition planning for the corporate entity, family dynamics and communication issues; and leadership empowerment. The course is intended for those who plan to consult or provide professional services to family-controlled companies and for those contemplating a career in a family firm. It will present both a theoretical framework for understanding the family form of business organization and a practice perspective on consulting to family firms and/or working as a family member in the family business.

MKTG756: MARKETING RESEARCH - IYENGAR R

The goal of the course is to familiarize students with the fundamentals of Marketing Research. Marketing Research involves developing research questions, collecting data, analyzing it and drawing inferences, with a view to making better business decisions. To this end the course is organized into two basic parts: (1) Data Collection and Research Design, and (2) Tools and Applications of Market Research. In essence, this is an Applied Statistics course where we focus on inference from Marketing Research data.

MKTG773: CUSTOMER BEHAVIOR - REED A

THIS CLASS SEEKS TO ANSWER 8 CRITICAL CONSUMER RELATED QUESTIONS:

1. How do I find out about my Market, Product Offering and Customer?
2. How do I make my Persuasive Communication work?
3. How do I know how customers perceive my product?
4. How do I relate Demographic Customer Characteristics to Behavior?
5. How do I construct key Psychographic Segments?
6. How do I relate Customer Psychographic tendencies to Reactions?
7. How do I estimate effectiveness of my Marketing Efforts?
8. How do I respond to uncontrollable factors in the Market Place?

Students will learn HOW TO answer each question using a specific managerial framework and an analytical tool. Note: Some of these analytics (e.g., Logistic Regression, Cluster analysis, ANOVA, Chi-Square analysis) partially overlap with statistical content from other marketing courses (e.g., Marketing 756), but our emphasis will be on how to use them to understand customers better. This is a drill down course that builds on concepts from MKTG 621 and 622.

To help provide additional insight into each question, at least one guest speaker from Industry will come into the class to discuss specific real world insights/applications related to that

MKTG776: APPL PROB MODELS MKTG - FADER P

Over the past five decades, statisticians have developed a number of models that have proven to be highly effective in their ability to explain and predict empirical patterns within many areas in business and the social sciences. These models use some basic “building blocks” from probability theory to offer behaviorally plausible perspectives on different types of timing, counting, and choice processes. Researchers in marketing have actively contributed to (and benefited from) these models for a wide variety of applications, such as new product sales forecasting, analyses of media usage, and targeted marketing programs. Other disciplines have seen equally broad utilization of these techniques.

As new forms of information technology provide increasingly rich descriptions of individual level shopping/purchasing behavior, these models offer great value to practicing managers, particularly those interested in pursuing CRM (“customer relationship management”) activities.

Furthermore, as more managers become comfortable with non-linear optimization techniques (using, for example, the “Solver” feature within Microsoft Excel), the specification and interpretation of these models can become a regular part of the sophisticated manager’s toolkit.

Taken as a whole, the methodological approaches covered in this course are well-suited to address the types of questions that are being asked with increasing frequency and interest by investors and managers of today’s data-intensive businesses.

MKTG778: STRATEGIC BRAND MGMT - MOGILNER C

Which brands make you happy? Apple? Starbucks? The Daily Show? Google? What draws you into these brands? How do companies create compelling brand experiences? How could you cultivate a brand that makes consumers happy? This course explores such questions with the goal of identifying the ingredients for building and managing an inspired brand, where brand is defined as “a reputation” – departing from traditional perspectives of brand.

The class will involve a broad ecosystem of contributors. Leaders from the world of brand—both small entrepreneurial companies and large, global market-leaders—will be incorporated into the class to offer first-hand perspectives about the challenges and lessons along their varied paths to success. This approach is intended to make the walls between the classroom and the world outside a little more porous.

The course has been created for individuals interested in building their own brands and/or immersing themselves in the enhancement of an existing brand. The course will interweave lectures, guest speakers, case discussions, in and out of class exercises—all of which will culminate in a Brand Audit group project that students will present in the final days of class.

The course will provide students with an appreciation of the role of branding and (taking a consumer-centric approach) will augment students’ ability to think creatively and critically about

the strategies and tactics involved in building, leveraging, defending, and sustaining inspired brands.

OPIM690: MANAG DECSN MAKING - MILKMAN K

Over the last 30 years, psychologists and economists have joined forces to study how people process information and actually make decisions, rather than how they would make decision if they were fully rational and selfish. This research program (dubbed behavioral economics) has provided an understanding of how people's decisions deviate from "optimal" choices as well as the consequences of such deviations. This course is devoted to understanding the nature, causes and implications of these limitations. The first two thirds of the course will focus on when individuals make decisions that deviate from the predictions of economics, and the final third of the course will focus on implications of these systematic decision biases for managers and policy makers.

The course has two main objectives. The first is improving the ability of the student (as a future manager) to influence the behavior of others, be they consumers, employees or people outside of a business relationship altogether. This will be accomplished by building on the toolbox that standard economics provides for influencing behavior (namely, incentives and information) with the insights from the aforementioned stream of research in behavioral economics.

The second objective is to improve the quality of students' own managerial decisions, primarily by enhancing the students' intuitive empirical abilities but also by improving their understanding of project evaluation.

OPIM692: ADV TOPICS NEGOTIATION - SCHWEITZER M

This course is designed to teach negotiation principles and to enable student to develop their negotiation skills. This course builds upon and assumes familiarity with the negotiation concepts covered in the prerequisite for this course: "Negotiations."

In this course, we extend the study and practice of negotiations, and we develop a deeper understanding for how specific aspects of the negotiation process (e.g., emotions, deadlines, trust violations) impact outcomes. Through course lectures, readings, and exercises, students will develop a rich framework for thinking about the negotiation process and acquire tools for guiding the negotiation process.

REAL721: REAL ESTATE INVESTMENTS - FERREIRA F

This course provides a broad introduction to real estate markets. Value of land, real estate prices, basic project evaluation, financing strategies, and capital markets issues related to real estate are covered. No prior knowledge of the industry is required, but students are expected to rapidly acquire a working knowledge of real estate markets.

Classes are conducted in a standard lecture format with discussion encouraged. The course contains cases that help students evaluate the impact of more complex financing and capital markets tools used in real estate.

REAL721: REAL ESTATE INVESTMENTS - WONG M

There are two primary goals of this class:

- 1) To expose you to the terms, issues, and topics in commercial real estate;
- 2) To give you the basic skills and intuition you need to begin to evaluate a variety of real estate investments.

Real estate is a multi-faceted field, encompassing both an operating industry and a broad category of investments. It has its own institutional features, jargon, and investment structures. As the survey course in the Real Estate Department, this class aims to provide a broad overview of the real estate field, rather than a narrow focus on any particular topic. We delve more deeply into a handful of aspects of the real estate field when I believe they are particularly relevant or when the example provides a more general insight. Higher-level classes in the Real Estate Department examine in more detail many of the topics from this class.

The presumption in this class is that you have no prior real estate experience, and no pre-existing knowledge of the real estate industry is necessary to do well in this class. However, if you have prior experience in the real estate industry, some topics might be familiar to you already.

REAL821: REAL ESTATE DEVELOPMENT - NAKAHARA A

Course Objectives. Four objectives will drive this course, helping you:

1. Become better decision-makers and real estate industry leaders.
2. Understand and assess the risks in real estate development and investments.
3. Be more productive in your first job.
4. Familiarize yourself with the real estate development process.

Course Topics. This course focuses on “ground-up” development as well as re-hab, redevelopment, and acquisition investments. We will examine the similarities and differences of traditional real estate product types including office, R&D, retail, warehouses, lodging, single family and multi-family residential, mixed use, and land. We will also analyze “specialty” uses like golf courses, resorts, and senior assisted living, and concepts like New Urbanism, sustainability, and timeshares. You will learn the development process from market analysis, site acquisition, zoning, entitlements, approvals, site planning, building design, construction, financing, and leasing to ongoing management and disposition. Special topics - workouts, leadership, and running a development company - will be discussed. Throughout the course, we will focus on risk management. In a business filled with uncertainties, minimizing risk results in maximizing long run profits and net worth accumulation.

Appendix C: Recruitment Materials

From: Kaufold, Howard
Sent: Thursday, November 17, 2011 3:09 PM
To: whg12; whg13
Subject: Do Wharton Research Study, Get Free Food, and Earn Your Chance at Cash Prize!

Dear Students,

We would like to ask for your help in a research study that is recruiting current Wharton MBA students. The research, conducted by a Wharton faculty member along with one of our curricular committees of faculty, department chairs and students, is attempting to understand the decisions of Wharton MBA students as they relate to pending changes in the MBA program. Through this study we will learn valuable information that we will use to improve the experience of Wharton students for years to come.

We want to emphasize that your participation is strictly voluntary. However, as a token of our appreciation, at the end of each session we will randomly choose two students and each one will receive \$250. (Each session will have approximately 20 students.) In addition, we will provide you with lunch (noon sessions) or dinner (6pm sessions). Your help will also be greatly appreciated as we want to ensure that we understand as best as possible the preferences of our MBA students with respect to these important design changes in the MBA program.

The study will last 90 minutes and take place in either Room F80 or F375 of Jon M. Huntsman Hall. Sessions will begin at 12 **noon** and **6pm** on

Monday 11/21 – F375 JMHH

Monday 11/28 – F80 JMHH

Tuesday 11/29 – F80 JMHH

Wednesday 11/30 – F80 JMHH

Thursday 12/1 – F80 JMHH

Please click <http://mktgweb.wharton.upenn.edu/mba-bhlab/> to sign up for any available time slot on one of the days listed above. (You need only participate in one session.)

We understand that this a busy time of the year for all students, but we do very much hope you will be able to help us with this valuable research study for our MBA program. Thanks in advance.

Yours,

[SIGNATURE]

Thomas S. Robertson, *Dean*

[SIGNATURE]

Howard Kaufold, *Vice Dean*

Appendix C. Representativeness

Subjects were representative of all Wharton MBA students on demographics as well as attitudes towards, and behavior in, the Wharton auction. Using data provided by the Wharton Dean's Office, Table 1 shows the demographics of our 132 subjects as well as the universe of Wharton MBA students in the 2011-2012 academic year. The final column reports the p-value of either a *test of proportions* or a *t-test* comparing our subjects to the universe of students. We see that based on demographics, our subjects are representative of the Wharton student body with $p > 0.1$ for each variable except "no race reported" ($p = 0.07$). Importantly for our purposes, our subjects look identical to the student body with regard to Auction behavior: namely, the number of points they had at the start of the Spring Auction (which began before the study took place) and the number of points they had when our study took place (points in the fourth round of the Spring Auction). For the second-year students in our study, we also examine data on their attitude towards the Wharton Auction as measured on the preceding spring's stakeholder survey. Our second-year subjects were almost identical to the universe of second-year subjects in reports in the effectiveness of the Wharton Auction.

Table A1: Representativeness of Experimental Subjects

	Subjects	Wharton MBAs	
	132	1660	
Demographics			p-value
First Year Student	51.7%	50.8%	0.83
Female	42.0%	47.0%	0.27
From United States	37.1%	34.3%	0.52
Finance Major	23.5%	25.7%	0.57
Total Registered Credits	17.1	17.0	0.96
Wharton Credits	11.5	11.3	0.56
White	48.5%	37.2%	0.01***
Asian	20.5%	27.0%	0.10*
Black, Non-Hispanic	5.3%	4.0%	0.46
Hispanic	3.0%	3.4%	0.83
Multi-Race	8.3%	7.2%	0.62
No race reported	14.4%	21.1%	0.07*
GPA	Subjects directionally higher		0.14
Auction Behavior			
Points at Start of Spring Auction	6899.6	6966.4	0.79
Points in 4th Round of Spring Auction	4992.3	4960.7	0.92
	Subjects	Wharton MBAs	
	62	731	
Auction Beliefs (Second years only)			
Reported Auction Effectiveness (0 to 7)	4.69	4.68	0.96

Table A1 reports data provided by Wharton. Due to Wharton's policy of grade non-disclosure, GPA levels cannot be reported. The auction beliefs data came from a stakeholder survey completed by rising second year students the preceding spring, so we only have it for the second-year students.

Appendix D

Table A2: Prevalence and magnitude of preference-reporting contradictions for people who use (and do not use) adjustments

Type of Comparison	# Comparisons with this Utility Difference	Accurate	Contradictions			
			All	Weak Preference	Preference	Weak Preference
All	1,662	84.42%	15.58%	5.17%	6.92%	3.49%
Did not use Adjustments	878	85.31%	14.69%	5.01%	6.38%	3.30%
Used Adjustments	784	83.42%	16.58%	5.36%	7.53%	3.70%

Table A2 shows all 1,662 comparisons. *Did Not Use Adjustments* indicates that the subject did not make an adjustment in the CEEI preference reports. *Used Adjustments* indicates that the subject made at least one adjustment in the CEEI preference reports. *Accurate* reports the percentage of these comparisons where the binary comparison choice confirms the CEEI preference report prediction. The *Contradictions* columns report the percentage of binary comparisons that contradicted the CEEI preference reports overall and at each level of preference.

Appendix E. Qualitative Questions

After subjects used each course allocation mechanism they answered qualitative questions about the schedule they received and the mechanism they had just used. These questions addressed the simplicity and clarity of the course allocation mechanism, general liking of the mechanisms, and the strategic effort required to use them. After subjects had used both mechanisms, they answered additional questions. All of these qualitative questions were asked to conduct the “side effects” exercise described in the paper. The full list of questions asked of subjects is listed in Table A3. In addition, subjects were given a page to write free responses at the end of the experiment.

Table A3: Qualitative Questions Subjects Answered

Question Wording	Timing	Scale
The way courses are allocated through this course allocation system is fair.	After using the first mechanism and again after using the second mechanism	“Strongly Disagree” “Disagree” “Somewhat Disagree” “Neither Agree or Disagree” “Somewhat Agree” “Agree” “Strongly Agree”
This course allocation system is easy for me to use.		
I understand how this course allocation system works.		
This course allocation system led to the best outcome I could hope for.		
I am satisfied with my course outcome.		
I enjoyed participating in this course allocation system.		
I like this course allocation system.		
My fellow students will like this course allocation system.		
I felt like I had control over my schedule in this course allocation system.		
This course allocation system is simple.		
I had to think strategically about what other students would do in this course allocation system.		
Someone with perfect knowledge of the historical supply and demand for courses could have had an advantage over me in this system.		
Which course allocation system did you prefer?	After using both mechanisms and completing binary comparisons	“Strongly Prefer 1 st ” “Prefer 1 st ” “Slightly Prefer 1 st ” “Unsure Which I Prefer” “Slightly Prefer 2 nd ” “Prefer 2 nd ” “Strongly Prefer 2 nd ”
Which course allocation system do you think your fellow students would prefer?		
In which course allocation system did you get a better schedule?		

Table A3 shows all the qualitative questions subjects were asked, when the questions were asked, and the responses available to the subjects.