The Multi-unit Assignment Problem: Theory and Evidence from Course Allocation at Harvard[†]

By Eric Budish and Estelle Cantillon*

We use theory and field data to study the draft mechanism used to allocate courses at Harvard Business School. We show that the draft is manipulable in theory, manipulated in practice, and that these manipulations cause significant welfare loss. Nevertheless, we find that welfare is higher than under its widely studied strategyproof alternative. We identify a new link between fairness and welfare that explains why the draft performs well despite the costs of strategic behavior, and then design a new draft that reduces these costs. We draw several broader lessons for market design, regarding Pareto efficiency, fairness, and strategyproofness. (*JEL* D63, D82, I23)

Educational institutions commonly place limits on the number of students in any particular class. These class-size limits create an instance of the *multi-unit assignment problem*: if it is not possible for all students to take their most desired schedule of courses, then how should seats in courses be allocated?[1] Other examples of multi-unit assignment problems include the assignment of tasks or shifts to workers, players to sports teams, shared scientific resources to their users, and airport takeoff-and-landing slots to airlines.

Perhaps the central difference between how practitioners and economists have approached multi-unit assignment and other similar market design problems is the attention paid to incentives. Practitioners often design mechanisms that would lead to economically desirable outcomes if agents told the truth about their preferences, but which fail to provide agents with an incentive to do so. Economists try where possible to design mechanisms where truthful reporting not only leads

to desirable outcomes, but is a dominant strategy. Such *strategyproof* mechanisms eliminate agents "having to engage in costly and risky strategic behavior" (Roth 2008) and are robust in the sense of Wilson (1987) and Bergemann and Morris (2005). Strategyproofness has played an important role in economists' market design recommendations in a variety of settings, including school choice procedures, auctions, two-sided matching markets, and kidney exchanges.[2] For multi-unit assignment specifically, the theoretical literature has focused almost exclusively on strategyproof mechanisms.

In this paper we use theory and field data to (i) show that the non-strategyproofness of a widely used multi-unit assignment mechanism has a large negative impact on welfare; (ii) show that nevertheless the mechanism from practice outperforms the strategyproof alternative emphasized by the extant theory; (iii) identify an important design consideration that the theory missed by limiting attention to strategyproof mechanisms; and (iv) propose a new mechanism that integrates the lessons of the analysis, and that is more attractive in the data than what either theory or practice had developed on its own. Overall, our analysis suggests a nuanced perspective on the role of strategyproofness in design: on the one hand we provide some of the most direct evidence to date on the costs of using a manipulable mechanism, but on the other hand we show that the costs imposed by restricting attention to strategyproof mechanisms can be larger still.

The mechanism we study is the *draft* used since the mid-1990s to allocate courses to students at Harvard Business School (HBS).[3] In the HBS draft, students report their preferences over individual courses to a computer, which then chooses courses for them one at a time over a series of rounds; the choosing order is random in the first round, and then reverses in subsequent rounds. Our data consist of students' actual (potentially strategic) reports of their preferences, as well as their *underlying truthful preferences*, from a survey conducted by the HBS administration.[4] This combination of truthful and stated preferences is powerful, first because it allows us to directly observe students' strategic manipulations and hence quantify their effect on welfare, and second because we can use the truthful preferences to simulate counterfactual mechanisms.

We begin by asking: *Does the non-strategyproofness of the HBS draft really matter in practice*? We approach this question in three steps. In the first step we study the theoretical properties of the HBS draft mechanism, with an eye towards the empirics. We first show that the HBS draft is simple to manipulate: students should overreport how much they like popular courses and underreport how much they like unpopular courses, so they do not waste early-round draft picks on courses

---

[2] Abdulkadiroğlu and Sönmez (2003) propose two strategyproof mechanisms for school choice, the first based on Gale and Shapley's deferred acceptance algorithm (since adopted by Boston and New York), and the second based on Gale's top trading cycles algorithm. Vickrey's (1961) auction is strategyproof and is the foundation for much of modern combinatorial auction design; see Milgrom (2004) and Cramton, Shoham, and Steinberg (2006). Roth and Peranson (1999) describe the role of strategyproofness concerns in the redesign of the medical match, which uses a variant of Gale-Shapley. Roth, Sönmez, and Ünver (2004) propose a strategyproof mechanism for kidney exchange that is based on top trading cycles.

[3] There are many different drafts in practice, both formal and informal, in settings ranging from firms (the allocation of tasks or shifts), to professional sports leagues (new players), to households (chores, divorce settlement), to dispute resolution (issues). See Brams and Straffin (1979) and Brams and Taylor (1996, 1999) for examples and design permutations.

[4] We describe in detail in Section IV our argument that the survey data are indeed truthful.

they can get in later rounds. We then show that equilibrium strategic behavior can affect students' welfare through two distinct channels. First, strategic behavior leads to congestion (popular courses reach capacity faster) which hurts students whose preferred courses are popular amongst other students. Second, strategic behavior can lead to ex post Pareto inefficient allocations, because a student may strategically underreport some class he likes and then turn out not to get it. It is possible that, ex ante, all students are worse off under equilibrium strategic play than under truthful behavior.

In the second step, we ask whether students behave strategically in the data, which cover the allocation of roughly 9,000 course seats amongst 900 MBA students in academic year 2005–2006. At the aggregate level, truthful preferences and actual submitted reports differ significantly and in a way that is consistent with equilibrium strategic behavior. At the individual level, we find that students' submitted preferences rearrange popular and unpopular courses in a manner that is mostly consistent with our characterization of equilibrium. Their optimization is not perfect, however, which suggests that strategic mistakes should also be counted as a potential cost of strategic behavior.

In the last step, we quantify the costs of students' strategic behavior, by comparing actual play of the HBS draft with what would have occurred if students had reported their preferences truthfully. We first document that, as predicted by the theory, strategic behavior causes congestion and ex post inefficient allocations: on average our Pareto-improvement-seeking integer program is able to find beneficial trades involving 64 percent of students and 9 percent of course seats. Ex ante welfare comparisons are more subtle: our data consist of ordinal preferences over individual courses, but welfare depends on cardinal preferences over bundles. We develop a computational method that uses first-order stochastic dominance and allows us to draw welfare comparisons for many students based on this limited ordinal information: nearly half of students are unambiguously harmed by strategic play, whereas only 10 percent unambiguously benefit. To reach comparisons for the other students, or to evaluate welfare from the perspective of a utilitarian social planner, we put increasingly more structure on preferences. Our approach is motivated by the HBS administration's emphasis on the number of students who obtain their single favorite course and on the average rank of the ten courses students receive. If students have what we call lexicographic preferences or what we call average-rank preferences, we can conclude that a large majority of individual students and a utilitarian social planner regard strategic play as harmful to welfare. The magnitudes are meaningful: strategic play reduces the number of students who receive their favorite course from 82 percent to 63 percent, and increases the average rank of the ten courses in a student's schedule from 7.66 to 7.99 (higher is worse, 5.50 is bliss).

Thus, the non-strategyproofness of the HBS draft indeed matters in practice. The mechanism is simple to manipulate in theory, is actually manipulated by students in practice, and these manipulations lead to significant welfare loss. This motivates our next question: *Should HBS switch to a strategyproof mechanism*?

There is a simple way to modify the draft to make it strategyproof: rather than have students choose courses one at a time over a series of rounds, have students take turns choosing their entire bundle of courses, over a single round. If the choosing order is random, as in HBS's draft, this mechanism is called the random serial dictatorship

(RSD). It is straightforward to show that dictatorships are the only strategyproof mechanisms within the class of random priority mechanisms (i.e., mechanisms where agents take turns choosing objects in some random order).[5] In fact, a series of papers have established more generally that dictatorships are the only mechanisms that are strategyproof and ex post Pareto efficient, and on those grounds these papers have recommended that RSD be considered for use in practice.[6]

We compare students' outcomes under the actual strategic play of the HBS draft with their outcomes under equilibrium truthful play of RSD. Ex post, the dictatorship is Pareto efficient whereas the draft is not. But when we compare the two mechanisms ex ante, the conclusion from the ex post analysis reverses: both the large majority of individual students and a utilitarian social planner prefer the HBS draft to RSD. Again, the magnitudes are large: switching to RSD increases the average rank from 7.99 to 8.74 and reduces the number of students who get their favorite course from 63 percent to 49 percent.

This surprising finding raises our next question: *What was the prior theory missing? What explains the ex post efficient RSD's poor ex ante welfare performance relative to the draft?* By limiting attention to strategyproof dictatorships, the prior theory also limited attention to mechanisms that yield a highly unequal distribution of outcomes, in which some students get all of the courses they like most, whereas others get few or none. Ex post, these outcomes are Pareto efficient, because it is impossible to improve the allocation of students late in the dictatorship order without harming the allocation of students early in the dictatorship order. But ex ante, this inequality harms welfare. The basic intuition is that each lucky student with a good random draw makes her second, third, …, last choices independently of whether they would be some unlucky student's *first* choice; that is, the lucky gain less than the unlucky lose. Drafts like HBS's do not have this problem—formally, no student gets to make her $n + 1$th choice before all students get to make their nth choice—hence the draft's outperformance of the dictatorship in the data.

Thus, both the draft from practice and the dictatorship from theory each have an important design weakness. Our last question is, quite simply, *are there better mechanisms for multi-unit assignment*? We integrate the lessons of our analysis into the design of a new mechanism called the "proxy draft." The proxy draft maintains the key strength of the HBS draft relative to the dictatorship, which is its distribution of outcomes. But, it modifies the HBS draft in two simple ways to mitigate its key weakness, the costs of strategic behavior. First, we have the mechanism play strategically on each student's behalf; this mitigates the welfare loss from strategic mistakes. Second, we allow the strategic proxy to utilize the student's *realized* position in the choosing order, whereas in the original HBS draft students decide on their strategic play without knowing the realization of this lottery; this mitigates the welfare loss from strategic risk-taking. In the data, the proxy draft increases ex ante welfare relative to the actual play of the HBS draft. These results demonstrate that the lessons of our analysis are easily implementable.

---

[5] See footnote 26 for a short proof.

[6] Pápai (2001, p. 270) concludes that "[t]he implications are clear (…) if strategic manipulation is an issue, one should seriously consider using a serial dictatorship, however restrictive it may seem." Similar conclusions are drawn by Ehlers and Klaus (2003, p. 266) and Hatfield (2009, p. 514).

Our paper offers insights on three concepts that are of interest beyond the specific problem of multi-unit assignment: ex post Pareto efficiency, ex post fairness, and strategyproofness.

First, our comparison between the draft and the dictatorship is a useful reminder that ex post Pareto efficiency need not be a proxy for ex ante welfare. In allocation problems with transferable utility (e.g., auctions) there typically is no distinction between the two concepts, and in the context of single-unit assignment, empirical results in Pathak (2006) and theoretical results in Che and Kojima (2010) suggest that the shortcut of focusing on ex post Pareto efficiency may be justified. But in settings like ours, with a large number of ex post Pareto efficient allocations with very different utility profiles, conclusions based on ex post Pareto efficiency can be misleading.

Second, our analysis highlights that there need not always be a trade-off between fairness and welfare; the draft is preferable to the dictatorship on both dimensions. Market design practitioners often describe fairness as an explicit design objective alongside welfare;[7] our analysis suggests that this may be sensible policy even if welfare is the ultimate objective.

Last, our paper provides some perspective on strategyproofness, which has long been viewed as an important desideratum in practical market design. Our field data allow us to directly document that students at HBS—real-life participants in a one-shot high-stakes setting—figure out how to strategically manipulate the non-strategyproof HBS draft mechanism. Further, we show that these manipulations have real welfare consequences. Taken on their own, these findings could be interpreted as strong evidence in support of the importance of strategyproofness in design. But our finding that the strategyproof alternative is less attractive still is strong evidence against imposing strategyproofness as an inflexible design requirement. The broader lesson, simple but worth stating explicitly, is that strategyproofness has both benefits and costs.[8] Limiting attention to strategyproof mechanisms implicitly treats the benefits as lexicographically more important than the costs. In our view this is rarely justified.

Organization of the Paper.—The remainder of this paper is organized as follows. Section I describes the environment and the HBS draft mechanism. Section II provides a theory of strategic behavior in the draft. Section III describes the data. Section IV asks whether students behave strategically. Section V quantifies the costs of strategic behavior. Section VI compares the draft to the strategyproof dictatorship. Section VII provides a theoretical explanation for the ex post efficient dictatorship's poor ex ante welfare performance. Section VIII introduces the proxy draft. Section IX concludes.

---

[7] For instance, Wharton (2009) writes that its course allocation system is "designed to achieve an *equitable* and *efficient* allocation of seats in elective courses when demand exceeds supply" (emphasis in original); several other examples are described in Sönmez and Ünver (2010). The HBS administration echoed this joint emphasis on efficiency and fairness in private conversations.

[8] We are not the first to observe that strategyproofness has both costs and benefits, and likely not to be the last. See Abdulkadiroğlu, Pathak, and Roth (2009) for an analysis in the context of school choice, and Section 7.4.2 of Fudenberg and Tirole (1991) for a textbook discussion.

I. The HBS Draft Mechanism for Course Allocation

A. *Environment*

Courses.—There is a finite set of C courses, $\mathcal{C}$.[9] Courses have capacities $\mathbf{q} = (q_1, \ldots, q_C) \in [0,1]^C$.

Preferences.—There is a continuum of students described by the interval $[0,1]$ and endowed with the Lebesgue measure.[10] Each student s is endowed with a von Neumann-Morgenstern (vNM) utility function u_s that indicates her utility from each bundle of courses, including singletons. As a regularity condition, we assume that the set of students who prefer one bundle over another is measurable, for any two pairs of bundles. Students can take 0 or 1 seats in each course, and at most $m > 1$ courses in total. Let P_s denote the ordinal preference relation generated by u_s.

We restrict attention to utility functions that generate strict ordinal preferences over individual courses and responsive ordinal preferences over bundles.[11] The main reason for focusing on responsive preferences in our theoretical analysis is that the HBS draft mechanism takes rank order lists over *individual* courses as its inputs, thereby implicitly assuming that preferences take this form. Responsiveness may seem a strong assumption in the context of course allocation, where students value the bundle of courses they ultimately get. However, we note that responsiveness is compatible with mild complementarities and substitutabilities across courses and that the HBS curriculum is designed to minimize such complementarities and substitutabilities. Moreover, survey evidence is broadly consistent with mild complementarities and substitutabilities.[12] While we will maintain the responsiveness assumption in our theoretical analysis, we will revisit it when we look at our data in Section IV.

Information.—We assume that students' preferences are common knowledge. Since we are working with a continuum and the HBS draft mechanism is anonymous, this is analogous to assuming that students' preferences are private information but that the distribution of preferences is common knowledge (Mas-Colell 1984).

Allocations.—An allocation in this environment is a measurable assignment of courses to students. We denote by a_s student s's allocation of courses. An allocation

---

[9] We use the terms "students" and "courses" because of our application. We could equally use the generic terms "agents" and "objects."

[10] The use of a continuum of students is a technical, rather than substantive, assumption. It simplifies proofs and helps clarify the key forces behind the results. For other treatments of assignment problems with a continuum of agents, see Abdulkadiroğlu, Che, and Yasuda (2008), and Che and Kojima (2010).

[11] Responsiveness requires that a student's ordinal preferences over bundles "respond" to their ordinal preferences over individual courses; essentially, it is an ordinal version of additive preferences. Formally, student s's preferences are responsive if, for any bundle of courses X and any courses $c, c' \notin X$, $c P_s c' \Leftrightarrow (X \cup \{c\}) P_s (X \cup \{c'\})$. Also, $c P_s \emptyset \Leftrightarrow (X \cup \{c\}) P_s X$ (Roth 1985).

[12] Specifically, we asked students in January 2006 whether they viewed any pair or group of courses as substitutes (in the sense that they would not want to take them together) or complements (in the sense that they would want to take them together). One third of respondents did not view any set of courses as substitutes and two thirds of respondents did not view any set of courses as complements. Most sets of substitutable or complementary courses were named by only one or two students. The main source of substitutability across courses is a preference for diversity. The main source of complementarity across courses arises from scheduling constraints for half-courses that meet in the first part of the semester and half-courses meeting in the second part of the semester.

is feasible if $|a_s| \leq m$ for almost all s and $\int 1_{\{s \in [0,1] | c \in a_s\}} ds \leq q_c$ for all c. A random allocation is a probability distribution over feasible allocations.

Pareto Efficiency.—A feasible allocation is ex post Pareto efficient if there is no other feasible allocation that all students weakly prefer and a strictly positive measure of students strictly prefers. A random allocation is ex ante Pareto efficient if there is no other random allocation that all students weakly prefer and a strictly positive measure of students strictly prefers.

B. *The HBS Draft Mechanism*

The allocation of courses at HBS takes place over two phases. Both phases take ordinal information about students' preferences over individual courses as their inputs.

In the first phase (initial allocation), students take turns choosing courses one-at-a-time. Specifically, each student s reports a rank-order list (ROL) $\hat{P}_s$ indicating her ordinal preferences over individual courses. We write $\hat{P}_s : c_1, c_2, c_3, \ldots$ to describe that student s puts course c_1 ahead of c_2, course c_2 ahead of c_3, and so on (with a slight abuse of notation, we will also write $P_s : c_1, c_2, c_3, \ldots$, to describe her true preferences over individual courses). Then, students are randomly ordered. Following Che and Kojima (2010), we model this by having each student draw a random priority number independently from the uniform distribution on $[0, 1]$. Students are allocated courses one-at-a-time over a series of m rounds. In odd rounds, which occur during time intervals $[0, 1], [2, 3], \ldots$, students are allocated courses one-at-a-time in ascending order of their random priority number. In even rounds, which occur during time intervals $[1, 2], [3, 4], \ldots$, students are allocated courses one-at-a-time in descending order of their random priority number. When it is student s's turn, she is allocated her most-preferred course on $\hat{P}_s$ that (i) she has not already received in a previous round; and (ii) has spare capacity. A formal description of the allocation generated by this mechanism is given in the online Appendix. At this stage, it is useful to mention a property that follows directly from our continuum assumption: given a submitted strategy profile by students, the HBS draft generates deterministic course run-out times.

The initial allocation phase is followed by an "add-drop" phase, during which students may add courses with excess capacity to their schedule, and drop courses they obtained in the initial allocation that they no longer want. We model this phase as a random serial dictatorship for the excess capacity from the initial allocation. Specifically, students submit new rank-order lists. Then, there is a single round taking place from time $[0, 1]$. When it is student s's turn, she is allocated her most-preferred schedule out of the courses she got in the initial allocation and whatever of the excess capacity courses still have spare capacity.[13] Note that if students report the same preferences in the add-drop phase as in the initial allocation phase there will be no add-drop activity.

---

[13] The add-drop phase in practice is a multi-pass version of our modeling abstraction. The difference is that in the multi-pass version it is possible that a course reaches capacity in the initial allocation phase but becomes available again in some pass of the add-drop phase. This possibility complicates the theoretical analysis without additional insight. Note, critically, that neither the actual nor the modeled add-drop phase allows for Pareto improving trades of courses at capacity.

Equilibrium.—A pure strategy Nash equilibrium in our environment is a measurable mapping from the set of students to the set of rank order lists, one for each phase. Because the second phase of the HBS draft is equivalent to RSD and RSD is dominant-strategy incentive compatible, we restrict attention to strategies where students report truthfully in that phase. When we write that student s plays strategy $\hat{P}_s$ we mean that s submits the ROL $\hat{P}_s$ in the initial allocation phase and submits P_s in the add-drop phase. Existence of a pure strategy Nash equilibrium in the full game is guaranteed because the action space is finite, students' expected utilities are continuous in the strategies of the other students, and students' expected utilities only depend on the fraction of students who report each preference profile (Schmeidler 1973; Mas-Colell 1984).

II. Strategic Behavior in the HBS Draft Mechanism

We start with a simple example that illustrates the mechanics of the HBS draft and suggests the behavior we can expect in equilibrium.

EXAMPLE 1 (Over-reporting and Congestion): *Let $m = 2$ and suppose there are 4 courses with capacity of $\frac{2}{3}$ seats each. Preferences are as follows:*

Proportion of population	Type	Preferences
$\frac{1}{3}$	P_1	c_1, c_2, c_3, c_4
$\frac{1}{3}$	P_2	c_2, c_1, c_3, c_4
$\frac{1}{3}$	P_3	c_1, c_3, c_4, c_2

Truthful play is not a best response for the P_2 types. If all other students play truthfully, a P_2 type who plays truthfully gets c_2 in the first round and c_3 in the second round. If a P_2 type instead submits preferences $\hat{P}_2 : c_1, c_2, c_3, c_4$, then he gets c_1 in the first round and gets c_2 in the second round. Since P_2 types prefer $\{c_1, c_2\}$ to $\{c_2, c_3\}$ this is a profitable deviation.

In fact, the P_1 and P_3 types reporting truthfully and the P_2 types reporting $\hat{P}_2 : c_1, c_2, c_3, c_4$ is a Nash equilibrium, for any vNM utilities consistent with these ordinal preferences. In this equilibrium, the P_1 and P_2 types get $\{c_1, c_2\}$ with probability $\frac{2}{3}$ and $\{c_2, c_3\}$ with probability $\frac{1}{3}$ and the P_3 types get $\{c_1, c_3\}$ with probability $\frac{2}{3}$ and $\{c_3, c_4\}$ with probability $\frac{1}{3}$. All students rank c_1 first, causing c_1 to reach capacity earlier than under truthful play.

Example 1 suggests that students can profitably manipulate the HBS draft by overreporting how much they like popular courses, and this causes those courses to reach capacity earlier. The intuition is straightforward: a P_2 type should not waste his first-round choice on c_2, since he can get c_2 for sure in the second round, and if he waits until round two to ask for c_1 he is sure not to get c_1. We develop a more detailed understanding of this strategic misreporting in Section IIA.

Example 1 also shows that strategic behavior has welfare consequences. P_2 students benefit from being able to act strategically while the P_1 and P_3 students are worse off from this strategic behavior. We further explore the welfare properties of equilibrium in Section IIB.

A. *Strategic Misreporting in the HBS Draft*

Our first result in this section shows that it is simple to find profitable manipulations of the HBS draft. To describe the manipulation formally we first need to define popularity.

DEFINITION 1 (Popularity): *Course c is $\hat{\mathbf{P}}$-popular if it reaches capacity under play $\hat{\mathbf{P}}$. Otherwise it is $\hat{\mathbf{P}}$-unpopular. For two $\hat{\mathbf{P}}$-popular courses c, c' we say that c is more $\hat{\mathbf{P}}$-popular than c' if c reaches capacity earlier than does c'.*

Because any strategy profile fully pins down deterministic course run-out times, the notion of popularity is well-defined and the "more popular" order is a complete order. By the continuum assumption, $\hat{\mathbf{P}}$-popular courses are also $(\hat{P}_{s'}, \hat{\mathbf{P}}_{-s})$-popular for any $\hat{P}_{s'}$, so we can as well talk about $\hat{\mathbf{P}}_{-s}$-popular courses. When the strategy profile under consideration is clear from the context we will simply say popular and unpopular.

As long as students can predict which courses run out in equilibrium they have the following profitable manipulation:

THEOREM 1 (Simple Manipulations): *Fix $\hat{\mathbf{P}}_{-s}$. Form the strategy $\hat{P}_s^{simple}$ by taking the first m courses in P_s and rearranging them so that $c\hat{P}_s^{simple}c'$ whenever:*

  (i) *$cP_s c'$ and both are popular or both are unpopular; or*

 (ii) *c is popular and c' is unpopular.*

The strategy $\hat{P}_s^{simple}$ generates weakly greater utility than truthful play P_s, for all realized priority orders.

The proof of Theorem 1 (see the online Appendix) is by induction, breaking down the difference between P_s and $\hat{P}_s^{simple}$ into a sequence of smaller deviations, each of which involves changing the relative position of a single unpopular course. We show that, for any realized priority order, each such deviation changes the student's *outcome* by at most a single course, and whenever there is such a change it is one that the student strictly prefers. Thus, even though P_s and $\hat{P}_s^{simple}$ may yield quite different lotteries over bundles, we are able to compare them based only on the student's ordinal preferences over individual courses and the assumption of responsiveness. Formally, the distribution of course bundles s receives under $(\hat{P}_s^{simple}, \hat{\mathbf{P}}_{-s})$ first-order stochastically dominates the distribution he receives under $(P_s, \hat{\mathbf{P}}_{-s})$, for *any* complete order over bundles that is compatible with P_s and responsiveness.[14]

While Theorem 1 shows that it is simple to find a strategic misreport that is preferable to truthful play, it does not claim that $\hat{P}_s^{simple}$ is a best response. Our next result explores the constraints that equilibrium behavior places on the relationship between submitted ROLs and students' ordinal preferences. Specifically, we ask the following

---

[14]For other papers exploiting first order stochastic dominance in the context of random assignment, see e.g., Roth and Rothblum (1999) and Ehlers and Massó (2008). We face the added difficulty in our multi-unit environment that ordinal preferences over individual courses pin down only a partial order over bundles.

question: Suppose we observe a strategy profile $\hat{\mathbf{P}}$ and students' ordinal preferences over individual courses $\mathbf{P}$. Under what conditions is $\hat{\mathbf{P}}$ compatible with equilibrium behavior? Theorem 2 provides necessary conditions for a student's play to be a best response; they have the interpretation of "mistakes" that best responders must avoid.

THEOREM 2 (Best Responses): *Consider any strategy profile $\hat{\mathbf{P}}$ and any student s. Let c be a popular course that s receives with probability less than one under $\hat{\mathbf{P}}$ and that is amongst s's m most truthfully preferred courses. If $\hat{P}_s$ is a best response to $\hat{\mathbf{P}}_{-s}$ then there cannot exist course c' ranked higher than c on $\hat{P}_s$ such that:*

(i) *c' is unpopular and s gets course c with positive probability by moving c ahead of c' on $\hat{P}_s$; or*

(ii) *$c\,P_s\,c'$, s never gets c and gets c' for sure under $\hat{\mathbf{P}}$, and he would get course c for sure by moving c ahead of c' on $\hat{P}_s$; or*

(iii) *$c'P_s\,c$, s never gets c under $\hat{\mathbf{P}}$, but he would get c with positive probability, and receive all popular courses that he truthfully prefers to c with unchanged probability, by moving c ahead of c' on $\hat{P}_s$.*

For each condition, the proof of Theorem 2 (see the online Appendix) identifies an alternative strategy that yields a distribution over bundles that first-order stochastically dominates the distribution student s gets under $\hat{\mathbf{P}}$.

The first condition in Theorem 2 confirms the intuition from Theorem 1 that good strategy involves underreporting unpopular courses relative to popular courses, while conditions (ii)–(iii) provide necessary conditions on how students should further rearrange popular courses, taking their relative popularity into account (which, in the statement of the theorem, translates into probabilities of getting these courses). Condition (ii) says that students *should not* reverse the relative ordering of two courses, if doing so makes them miss the preferred course for sure. By contrast, condition (iii) says that students *should* reverse this relative ordering (placing the more popular course ahead) if this does not come at the cost of missing a more preferred course. Together, these three conditions, which depend only on the kind of information our data contain, form the basis of our test for equilibrium behavior in Section IV.

We conclude this section with two further comments on the nature of equilibrium in the HBS draft. First, there can be multiple equilibria. This is intuitive since course popularities are endogenous: if students think course c is more popular than course c' they will tend to place course c ahead of course c', and vice versa. Second, there exist environments where truthful play is an equilibrium. This will follow as a special case of Lemma 1 in Section VII.

B. *Welfare*

Strategic behavior in the HBS draft has redistributive consequences. On the one hand, it helps students who, by overreporting their preferences for popular courses, increase their chances of getting them. On the other hand, it creates congestion for these courses and hurts students who value them highly. These points can be seen in Example 1.

Strategic behavior also has efficiency consequences, due to the intrinsic trade-off that students face between upgrading a less preferred but very popular course and the risk of potentially missing a seat in a preferred but less popular course. Example 2 illustrates.

EXAMPLE 2 (Ex post Inefficiency of Equilibrium Strategic Play): *Let* $m = 2$. *Courses* c_1, c_2 *have excess demand with respective capacity* 0.6 *and* 0.8. *All other courses do not. Suppose preferences are as follows (where "other" stands for courses other than* c_1 *or* c_2*)*

Proportion of population	Type	Preferences
0.3	P_1	c_1, c_2, other
0.4	P_2	c_2, c_1, other
0.3	P_3	c_2, other

Consider the strategy profile where all students of type P_1 *play* $\hat{P}_1 : c_2, c_1$, *other, and the* P_2 *and* P_3 *types submit truthful ROLs. Under this strategy profile, all students request* c_2 *in round* 1, *which means that only the first* 0.8 *are successful. Those who do not get their first choice get their second choice in round* 1. *At the beginning of round* 2, 0.46 *seats* $(0.6 - (0.7)(0.2))$ *remain in* c_1, *whereas* 0.56 *students will request it. This means the first* 82 *percent of these students are successful. Thus, the* P_1 *and* P_2 *types face the following lottery:*

$$(1) \qquad \left[0.66 : \{c_1, c_2\}; \quad 0.20 : \{c_1, \text{other}\}; \quad 0.14 : \{c_2, \text{other}\} \right].$$

To check that this strategy profile is an equilibrium, we only need to look at the opportunity for a P_1 *student to deviate and submit his truthful preferences. If he does, he gets the deterministic outcome* $\{c_1, \text{other}\}$. *Thus* $\hat{P}_1, P_2, P_3$ *is an equilibrium if all* P_1 *students prefer the lottery in* (1) *over the deterministic outcome* $\{c_1, \text{other}\}$. *This equilibrium is ex post inefficient because it is possible that a* P_1 *student ends up with* $\{c_2, \text{other}\}$ *and that a* P_2 *student ends up with* $\{c_1, \text{other}\}$. *Those students would prefer to trade.*

Because equilibrium play can result in ex post inefficient allocations, the HBS draft is ex ante inefficient. We next show something stronger: it is possible that equilibrium strategic behavior strictly harms all students relative to truthful play.

EXAMPLE 2 (Continued) (Strategic Behavior May Hurt All Students Ex Ante): *Consider again Example* 2. *The following table compares the lotteries that students face under truthful behavior and under the equilibrium identified in Example* 2 *above.*

Type	Lottery under truthful play	Lottery under strategic behavior
P_1	$\left[0.33 : \{c_1, c_2\}; 0.67 : \{c_1, \text{other}\} \right]$	$\left[0.66 : \{c_1, c_2\}; 0.20 : \{c_1, \text{other}\}; 0.14 : \{c_2, \text{other}\} \right]$
P_2	$\left[0.75 : \{c_1, c_2\}; 0.25 : \{c_2, \text{other}\} \right]$	$\left[0.66 : \{c_1, c_2\}; 0.14 : \{c_2, \text{other}\}; 0.20 : \{c_1, \text{other}\} \right]$
P_3	$\left[1 : \{c_2, \text{other}\} \right]$	$\left[0.80 : \{c_2, \text{other}\}; 0.20 : \{\text{other, other}\} \right]$

Clearly, the P_2 and P_3 types are worse off under strategic behavior, independently of cardinal information about their preferences. The P_1 types are worse off if, for example $u_s(\{c_1, c_2\}) = 5$, $u_s(\{c_1, other\}) = 4$ and $u_s(\{c_2, other\}) = 1$.

Example 2 suggests that there is a Prisoners' Dilemma aspect to strategic behavior. Of course, each individual student prefers the option to play strategically rather than be constrained to play truthfully. But when all students have this option, everyone can be made worse off. We summarize the implications of Example 2 as follows.

THEOREM 3 (Efficiency Consequences of Strategic Behavior):

   (*i*) (*ex post*) *The outcome of the HBS draft mechanism under equilibrium play can be ex post inefficient.*

  (*ii*) (*ex ante*) *It is possible that all students strictly prefer the distribution of outcomes they receive from truthful play to that which they receive under equilibrium play.*

Theorem 3 summarizes the efficiency properties of equilibrium behavior in the HBS draft. In practice, students may also make strategic mistakes and it is easy to show that these have ex post and ex ante efficiency consequences as well. We will empirically explore the efficiency consequences of strategic play in Section V.

III. Description of Data

Our dataset covers the allocation of second-year elective courses to second-year MBA students at Harvard Business School during the 2005–2006 academic year. Students choose up to ten elective full courses each, five for each semester. Courses for both semesters are allocated in a single allocation process.

A. *Timing of Actions and Information*

Students are asked to report their ranking over individual courses at three separate times: in early May, in mid-May, and in mid-July. Prior to this, students have information on the past enrollment of each course and they receive both official and unofficial course evaluation information.

In *early May*, students are asked by the administration to voluntarily participate in a survey in which they rank their top five favorite courses. The survey explicitly asks students to report their preferences truthfully, and we are not aware of any compelling reason to disobey this request. The results are used to aggregate information about demand and adjust some course capacities. The students have access to the full results, except for the student identities which are removed.

In *mid May*, students participate in a trial run of the allocation mechanism. Participation is compulsory and students can rank up to 30 courses. The algorithm is run a single time, and then the administration reports the resulting course enrollments. For courses at capacity, students are told by how many times the course was oversubscribed. In addition, students are told which ten courses were most often

ranked first in the submitted ROLs, along with the number of students who ranked each first. Students do not receive any feedback on their individual assignment of courses from the trial run.

Finally, students submit their ROLs for the real run of the mechanism in *mid July*.

B. *Course Characteristics*

Our data contain all course capacities as they were available at the time of the May trial run and the July actual run of the algorithm. Seats in 70 courses and 22 half-courses were offered in May for a total capacity of 11,871 seats. Course capacities ranged from 12 to 404 students. The numbers for July were 67 courses and 22 half-courses for a total of 10,898 seats.[15] The capacities range was 12–401.

C. *Submitted Preferences*

Our data contain students' submitted ordinal preferences over individual courses (with student identifiers) in the May 2005 poll, May 2005 trial run, and the July 2005 actual run of the algorithm. In addition, we conducted an auxiliary survey in January 2006 in which we asked students to rank their 30 favorite courses. The poll was conducted before second-semester courses started. In the poll, we explicitly asked students to rank the courses according to their true preferences, independently of whether they got the course or not. The stated objective of the poll was to collect data on preferences to investigate potential improvements to the HBS allocation mechanism.

Table 1 summarizes the number of students and courses covered by the data on each occasion. Because participation was compulsory, the May trial run data and the July actual run data cover the entire population. The small discrepancy in numbers is due to students leaving for or returning from military duty, maternity leave, or any other leave of absence.

IV. Evidence of Strategic Behavior

In this section, we provide evidence, both at the aggregate level and at the individual level, that students at HBS strategically misreport their preferences in a manner consistent with our theoretical analysis in Section II. The analysis lends support to the interpretation of May poll responses as truthful and July ROLs as equilibrium strategic play, which is natural in the context of our data. However, there are two nuances. First, there is evidence, both at the aggregate level and the individual level, of some preference changes between May and July. We will attempt to correct for such preference changes when we conduct the welfare analysis. Second, there is evidence at the individual level that some students make strategic mistakes. We will *not* attempt to correct for such mistakes in the welfare analysis, since the complexity of calculating a best response is an intrinsic part of the HBS draft mechanism.

---

[15] Four courses were added between the May poll and the trial run. Between the trial run and the July run, one course was added, four courses were cancelled, one full semester course was changed into a half course, and several courses had their capacities increased or decreased slightly.

TABLE 1—DESCRIPTIVE STATISTICS: SUBMITTED PREFERENCES

	May '05 poll	May '05 trial run	July '05 run	Jan '06 poll
Number of students	460	922	916	163
Average number of courses per ROL	5	22.33	22.12	17.46
Std. Dev. of number of courses per ROL	0	5.13	4.92	7.31
Number of courses listed at least once	84	92	89	92

A. *Evidence Based on Aggregate Data*

If May poll preferences are truthful and July ROLs represent equilibrium behavior, then aggregate demand in May and in July should differ in a way that is related to the underlying popularity of courses: roughly speaking, we expect the aggregate distribution of students' preferences for popular courses to be higher in July during the real play of the mechanism than in the May poll, and vice versa for unpopular courses.

To test for the equality of aggregate demand for a course in the May poll and in the July actual play, we use Gehan's (1965) extension of the Wilcoxon rank-based test for discrete and censored data (censoring in our data arises from the fact that students only rank five courses in the May poll and not all students rank 30 courses in July). Consider course c and a sample of students. Course c's distribution of ranks in that sample, $D_c(r)$, is the proportion of students that place course c at rank r or before in that sample. The null hypothesis of the test is that $D_c^{May}(r) = D_c^{July}(r)$ for all $r \leq 5$, where D_c^{May} and D_c^{July} are the distributions of ranks in the May poll and in July.[16]

Table 2 reports the results for the 84 courses that appear in both the May poll and the July run. As a proxy for students' beliefs about a course's popularity, we use the administration's reports after the trial run about predicted enrollment in each course. For courses at capacity, they also report very explicitly by how much the course was oversubscribed. In the table, we categorize courses not at capacity as "unpopular" and courses at capacity according to their level of oversubscription.[17]

Table 2 shows that the null hypothesis of unchanged demand between the May poll and the July run is rejected for 42 out of the 84 courses (50 percent). For courses that are multiple times oversubscribed, the reason for rejection is always because demand is higher in July. For courses that are unpopular or popular but less than once oversubscribed, rejection occurs because demand is lower in July. This pattern is consistent with strategic behavior and equilibrium play by students.

However, systematic preference changes, either through new information correlated to popularity, or through social learning (for example, through the feedback from the May poll or the trial run), could also generate the pattern seen in Table 2. To distinguish between these preference-based explanations and an explanation based on strategic behavior, we compare aggregate demand in the May poll with the aggregate demand in the poll that we carried out in January 2006. If the results of Table 2 are due to preference changes, we should find the same pattern of rejection of the

---

[16] The restriction of the test to $r \leq 5$ is due to the censoring of preferences in the May poll.
[17] Alternative definitions of popularity based on the May poll or the July ROLs yield very similar results.

TABLE 2—COMPARISON BETWEEN MAY POLL DEMAND AND JULY RUN DEMAND
(*5 Percent level*)

	N	July demand lower	No difference	July demand higher
Unpopular	48	27	21	0
Popular, $0 \times$ oversubscribed	16	6	10	0
Popular, $1 \times$ oversubscribed	8	1	7	0
Popular, $2 \times$ oversubscribed	6	0	4	2
Popular, $3 \times$ oversubscribed	3	0	0	3
Popular, $4 \times$ oversubscribed	1	0	0	1
Popular, $5 \times$ oversubscribed	2	0	0	2

null hypothesis when we compare poll responses in May 2005 to poll responses in January 2006. Table 3 shows that this is not the case.

Despite the fact that eight months elapsed between the May poll and the January poll, we are only able to reject the null hypothesis of unchanged demand for 15 percent of courses, as compared with 50 percent above.[18] Additional evidence comes from the pattern of rejections in Table 3. Whereas in July, 27 of 48 unpopular courses had significantly lower demand, in January only three do, and one unpopular course even has significantly higher demand. Amongst courses that are at least one-time oversubscribed, seven had significantly higher demand in July versus one significantly lower, whereas the numbers are four versus three in January. Still, there are nine courses in total for which the null hypothesis is rejected both in July and January in the same direction, which is consistent with a persistent shock to preferences for these courses.

We conclude that there is substantial change in aggregate reported preferences between the May poll and July run, and that most of this change can be explained by short-term strategic considerations rather than long-lasting preference changes due to social learning or new information.

B. *Evidence Based on Individual Data*

If May poll preferences are truthful and July ROLs represent equilibrium behavior, then each individual student's July ROL should satisfy Theorem 2's necessary conditions for a best response given the student's own truthful (May poll) preferences and the realized (July ROL) popularity.

To run this test, we first identify popular and unpopular courses based on Definition 1 and 10,000 runs of the HBS draft using the July run submitted preferences and randomly drawn priority orders. We then focus on the 456 students who participated in both the May poll and the July run. Out of the 2,280 courses that appear as their five most-preferred courses according to the poll, 1,775 are popular.

For each of these courses, we check whether Theorem 2's three necessary conditions are satisfied. Out of the 1,775 popular course entries, 1,438 (81 percent) are compliant with Theorem 2, while 337 (19 percent) are in violation.

Conceptually, there are two possible explanations for each Theorem 2 violation: either the student's reported preferences in May no longer reflect his true preferences

---

[18] During this time period students took Fall 2005 classes, and their experience of these courses may have affected their preferences. The rejection rate for the subsample of Winter 2006 courses is the same (15 percent).

TABLE 3—COMPARISON BETWEEN MAY POLL DEMAND AND JANUARY POLL DEMAND
(*5 Percent level*)

	N	January demand lower	No difference	January demand higher
Unpopular	49	3	45	1
Popular, $0 \times$ oversubscribed	16	2	14	0
Popular, $1 \times$ oversubscribed	8	2	6	0
Popular, $2 \times$ oversubscribed	6	0	5	1
Popular, $3 \times$ oversubscribed	3	1	1	1
Popular, $4 \times$ oversubscribed	1	0	0	1
Popular, $5 \times$ oversubscribed	2	0	1	1

in July, or the student made a mistake. In general, it is impossible to distinguish between these two explanations. Instead what we will do in the rest of this section is try to identify those cases that seem likely to be caused by preference changes, and those that seem likely to be caused by strategic error.

Preference Changes.—The violations that seem most likely to be caused by preference change are those cases where a student ranked a popular course amongst her top five courses in May, and then does not even rank the course in July. 152 of the 337 violations fit this pattern. An alternate explanation is that the student mistakenly expected the course to be unpopular. In 98 of these cases however the course was at least once oversubscribed in May, so this alternate explanation is unlikely for those cases.

A second kind of violation likely to be caused by preference change are those cases where we suspect social learning or new information based on our analysis of Section IVA. There are four courses for which aggregate demand was lower both in July and January relative to the May poll, and 44 of the violations correspond to students downgrading or dropping one of these four courses. Together, these likely preference changes can account for 55 percent of Theorem 2 violations.

Analogous to preference changes are failures of our maintained assumption that preferences are responsive. If a student's preferences are non-responsive, what looks to us like a failure to best respond could in fact be a rational response to substitutabilities or complementarities. However, even when we define substitutabilities and complementarities broadly, we can explain only a small number of the violations.[19]

Strategic Mistakes.—Students may make strategic mistakes if they have incorrect beliefs about courses' equilibrium popularity or fail to optimize correctly.

While we cannot directly observe beliefs, we can identify cases where incorrect beliefs are a plausible explanation of the violation. 31 violations involve courses

---

[19] If a student regards two courses as close substitutes for each other, it may be rational for the student to omit one of them from her July ROL. For each Theorem 2 violation involving downgrading a popular course, we check whether this course is part of a pair or group of courses described as substitutes by *any* student in the January 2006 survey. If this is the case, we then check whether the student prefers another course in the pair or group, and if so, whether the student gets the preferred course with positive probability. Twenty-nine violations can be explained this way, of which 11 involve the student dropping the course entirely. Likewise, if a student regards two courses as complements, it may be rational to omit both. Half courses are the main source of complementarity in our application. Eleven of the violations involve downgrading a popular half course, of which 6 involve the student dropping the course entirely.

that appeared to be unpopular based on the May trial run but then became popular in the July actual run of the mechanism. A student who expected the course to remain unpopular might rationally downgrade the course on their July ROL, which then looks to us like a Theorem 2 mistake. Likewise, 33 violations can be explained by courses that appeared to be popular based on the May trial run but ended up not being popular in July. Some students needlessly placed these courses ahead of other popular courses.

We also cannot directly observe optimization errors, but we can identify students whose play, though flawed, is at least directionally consistent with equilibrium behavior. Out of the 190 violations of Theorem 2 that involve incorrect rearranging of popular courses (i.e., cases (ii)–(iii)), 148 instances correspond to cases where the popular course placed too high in the student's ROL is strictly more popular than the popular course placed too low. Such play is consistent with a basic understanding of the incentive to use early-round choices on courses that are especially popular, and while flawed is at least directionally sensible given course popularities and the student's underlying preferences. Together, these likely strategic mistakes can account for 61 percent of Theorem 2 violations.

Summary.—Our takeaways from the analysis at the individual level are as follows. The joint hypothesis that May poll responses represent truthful preferences and July ROLs represent equilibrium behavior does a reasonable job of organizing the individual level data. The data do suggest however that there are a meaningful number of strategic mistakes and preference changes. Our welfare analysis aims to include the strategic mistakes—the difficulty of optimal play can be considered an intrinsic part of the HBS draft—but correct for the preference changes, which are a flaw of the data and not the market design.

V. Welfare Consequences of Strategic Play

This section quantifies the welfare consequences of students' strategic play under the HBS draft mechanism. Section II identified two channels through which strategic behavior affects welfare: increased congestion and ex post inefficiencies due to strategic risk-taking. This section begins by documenting that strategic behavior indeed has an effect through both channels, and that the magnitudes are large. We then turn to the analysis of students' ex ante welfare. Specifically, we compare students' distributions of outcomes under the actual strategic play of the HBS draft to their distribution of outcomes under a non-equilibrium counterfactual in which students report their preferences truthfully. We find that many more students are harmed by strategic behavior than benefit. Welfare comparisons at the aggregate level confirm this: a utilitarian social planner prefers truthful play over strategic play, as long as students are risk neutral or risk averse.

The main inputs for our analysis in this section are students' May poll responses and July submitted ROLs, which we take as their truthful preferences and actual strategic play, respectively. A limitation of May poll responses is that they are restricted to students' top-five favorite courses. For positions six and below, we append the courses not amongst the student's top five so as to preserve their relative ordering in the student's July ROL. To avoid counting preference changes as strategic behavior, we also

adjust for the cases of likely preference change identified in Section IVB. The online Appendix provides further details on the construction of preferences for our main specification, and describes the robustness checks we carry out. In particular, one of our robustness checks treats all Theorem 2 violations as due to preference change, and another treats all Theorem 2 violations as due to strategic error.

We have both May poll responses and July submitted ROLs for 456 students, so we work with an economy consisting of these students with course capacities adjusted accordingly. May-poll respondents are representative of the entire population both in terms of their strategic behavior in July and their reported preferences in the follow-up survey in January.[20]

A. *Effect of Strategic Behavior on Congestion*

To evaluate the effect of strategic behavior on congestion, we run the HBS draft mechanism for 10,000 random priority orders using both truthful and strategic preferences, and record the time at which each course reaches capacity. We say that a course reaches capacity earlier under strategic (truthful) play if the time it reaches capacity is strictly earlier than it is under truthful (strategic) play for at least 99 percent of the priority orders.

As suggested by our theoretical analysis, congestion is related to popularity. Nearly all of the courses that were at least one time oversubscribed based on the trial run reach capacity earlier under strategic play (18 out of 20). Very few of the other courses experience such congestion (4 of 64). The magnitudes are large, especially amongst the most popular courses. For the 31 courses that reach capacity under both truthful and strategic play, the round by which they do so goes from 6.45 to 5.74 on average. If we focus on the 12 courses that are at least twice oversubscribed, the average round at which they reach capacity goes from 3.79 under truthful play to 2.61 under strategic play.

B. *Ex Post Efficiency Consequences of Strategic Play*

To assess the magnitude of ex post inefficiency in the HBS draft, we run the draft for 100 random priority orders using students' July run preferences. For each priority order, we compute the number of ex post Pareto improving trades we can find based on students' truthful preferences.

Because our data consist of ordinal preferences over individual courses, there are some profitable trades that we will not be able to find. For instance, if a student's truthful ROL is $P_s : c_1, c_2, c_3, c_4$ and his allocation is $\{c_1, c_4\}$ then we know that he is willing to trade c_4 for c_2 or c_3, but we do not know whether he is willing to trade the bundle $\{c_1, c_4\}$ for the bundle $\{c_2, c_3\}$. Subject to this caveat, it is without loss of

---

generality to restrict attention to one-for-one trades: whatever many-for-many trades a student is willing to execute can be described using multiple one-for-one trades.

Let $x_{scc'}$ indicate whether we execute a one-for-one trade in which student s gives c and gets c'; in a multi-way trade (which is allowed) student s gives c to one student and gets c' from another. For each priority order, denoted λ, we maximize the number of Pareto-improving trades subject to the constraint that each student trades each course at most once. The result is the following binary integer program:

$$\max_{\mathbf{x} \in \{0,1\}^{SC^2}} \sum_{s,c,c'} x_{scc'}$$

such that

$$(2) \qquad\qquad \sum_{s}\sum_{c'} x_{scc'} - x_{sc'c} = 0, \forall c$$

$$(3) \qquad\qquad \sum_{c'} x_{scc'} + x_{sc'c} \leq 1, \forall s, c$$

$$(4) \qquad\qquad x_{scc'} = 1 \Rightarrow c \in a_s(\hat{\mathbf{P}}, \lambda), c' \notin a_s(\hat{\mathbf{P}}, \lambda), c' P_s c,$$

where constraint (2) captures the condition that each course must be given as often as it is received, constraint (3) prevents a student from trading the same course twice, both to ensure feasibility and to avoid double-counting, and constraint (4) ensures that the trade is both feasible and desirable. Table 4 suggests that the level of ex post inefficiency in the HBS draft is substantial: on average 9 percent of course seats can be profitably reallocated, involving 64 percent of students.[21] For comparison, if students played truthfully it is easy to see that there should not be any one-for-one Pareto-improving trades.

C. *Consequences of Strategic Behavior on Ex Ante Individual Welfare*

In this section, we compare students' ex ante welfare under the actual play of the HBS draft to their welfare under a non-equilibrium counterfactual in which all students submit their true preferences. The technical challenge we face is that our data consist of students' ordinal preferences over individual courses, whereas ex ante welfare comparisons require information about students' cardinal preferences over bundles of courses.

To begin, we investigate what can be said based solely on the maintained assumption of responsive preferences. This assumption generates, for each student, a partial order over bundles of courses and thus, a fortiori, a partial order over random

---

[21] The magnitudes in Table 4 are consistent with the results of an informal survey conducted by two HBS students in the Spring of 2005. As part of a class project related to the HBS draft mechanism, these students surveyed 160 of their classmates. One of the questions was: "Did you know of a trade with another student that could have made you both better off?" 58.1 percent responded yes. This is suggestive of both the magnitude and students' awareness of ex post Pareto inefficiency.

TABLE 4—EX POST PARETO IMPROVING TRADES

	Mean	Std. Dev.
Number of executed trades per student	0.91	(0.03)
Percent of allocated course seats traded	9.1 %	(0.3 %)
Percent of students executing		
0 trades	35.8 %	(1.6 %)
1 trade	42.3	(2.1)
2 trades	17.5	(1.6)
3+ trades	4.4	(0.8)

allocations. This partial order over random allocations yields the following comparison criterion:

COMPARISON RESULT 1 (Responsive Preferences): *Suppose that student s has responsive preferences. Then student s prefers truthful play* $\mathbf{P}$ *to strategic play* $\hat{\mathbf{P}}$ *if, for every complete order over bundles consistent with the partial order implied by P_s and responsiveness, the lottery over bundles he receives under* $\mathbf{P}$ *first-order stochastically dominates that under* $\hat{\mathbf{P}}$. *He prefers* $\hat{\mathbf{P}}$ *if the reverse relationship holds.*

Comparison result 1 (CR1) is a very strong test because it requires first-order stochastic dominance, which is a demanding order, and does so for *every possible complete order over bundles* compatible with *s*'s preferences over individual courses and responsiveness.

We introduce the following method to implement CR1. First, we run the HBS draft mechanism for both truthful play and strategic play for 100 randomly drawn priority orders, which we denote $\lambda_1, \ldots, \lambda_{100}$. Second, for each student *s*, we form the following bipartite graph:[22] one set of nodes is her outcomes from truthful play, $a_s(\mathbf{P}, \lambda_1), \ldots, a_s(\mathbf{P}, \lambda_{100})$; the other set of nodes is her outcomes from strategic play, $a_s(\hat{\mathbf{P}}, \lambda_1), \ldots, a_s(\hat{\mathbf{P}}, \lambda_{100})$; for each pair λ_j, λ_k an edge is drawn from node $a_s(\mathbf{P}, \lambda_j)$ to node $a_s(\hat{\mathbf{P}}, \lambda_k)$ if we know from P_s and responsiveness that *s* weakly prefers $a_s(\mathbf{P}, \lambda_j)$ to $a_s(\hat{\mathbf{P}}, \lambda_k)$. Third, we check if the resulting bipartite graph has a perfect matching, i.e., a subset of edges such that each node in the truthful-play node set is connected to exactly one node in the strategic-play node set. If there is a perfect match, this means that the set of outcomes under $\mathbf{P}$ first-order stochastically dominates that under $\hat{\mathbf{P}}$ for every complete order consistent with the responsiveness partial order. To check if $\hat{\mathbf{P}}$ dominates $\mathbf{P}$ we reverse the way the edges are drawn.

We find that 41 percent of students are unambiguously harmed by strategic play, versus just 10 percent that unambiguously benefit. 1 percent of students are indifferent. For the remaining 48 percent the comparison is indeterminate. Ignoring the indeterminate cases for now, we note that these results are already suggestive of our intuition from Section II about the asymmetry between the costs and benefits of strategic play: the students who are harmed by strategic play are those whose preferred courses are popular, and by definition there are more of them than students whose preferred courses are

---

[22] A graph is bipartite if its nodes can be divided into two disjoint sets, such that every edge of the graph connects a node in one of these sets to a node in the other set.

not popular and who then gain from the ability to overreport popular courses. The fact that ex post inefficiencies hurt everyone further exacerbates this asymmetry.

To confirm these results, however, we need to pin down the indeterminate cases. Our approach is to place more structure on the relationship between students' ROLs and their vNM utility functions.

DEFINITION 2: *Student s is said to have additive preferences if there exist numbers $v_s(c)$ for all courses in $\mathcal{C}$, and a strictly monotonic function $f_s(\cdot)$, such that s's vNM utility from any allocation a_s can be written as $u_s(a_s) = f_s\left(\sum_{c \in a_s} v_s(c)\right)$. The sum $\sum_{c \in a_s} v_s(c)$ itself is called student s's value of bundle a_s. If $f_s(\cdot)$ is (weakly) concave then s is (weakly) risk averse with respect to bundle values. If $f_s(\cdot)$ is linear then s is risk neutral with respect to bundle values.*

Additive preferences are a special case of responsive preferences. If s has additive preferences and is risk neutral then his expected utility under strategy profile $\hat{\mathbf{P}}$ can be expressed as $\sum_{\lambda \in \Lambda} \left(\frac{1}{|\Lambda|}\right) \sum_{c \in a_s(\hat{\mathbf{P}},\lambda)} v_s(c)$, where λ denotes a priority order and Λ denotes the set of all priority orders. This yields our second comparison result:

COMPARISON RESULT 2 (Additive Preferences, Risk Neutral): *Suppose that student s has additive preferences and is risk neutral. Student s prefers truthful play* **P** *to strategic play* $\hat{\mathbf{P}}$ *if, for all $n \in \mathbb{N}$, the expected number of top-n courses he gets under* **P** *exceeds that under* $\hat{\mathbf{P}}$. *He prefers* $\hat{\mathbf{P}}$ *if the reverse relationship holds.*

A special case of additive preferences that gives us a complete order over bundles is when the difference in v_s's derived from the student's 1st and 2nd highest ranked courses is the same as that between the nth and $n + 1$th ranked courses, for any n. In this case, a sufficient statistic for his realized utility is the *average rank* of the courses in his allocation. Average rank is a measure of mechanism performance emphasized by the HBS administration. Average-rank preferences yield the following comparison results:

COMPARISON RESULT 3 (Average-rank Preferences): *Assume student s has average-rank preferences and let $\bar{r}_s(\mathbf{P}, \lambda)$ $\left(\bar{r}_s(\hat{\mathbf{P}}, \lambda)\right)$ denote the average rank of the courses he gets under strategy profile* $\mathbf{P}(\hat{\mathbf{P}})$ *for the priority order λ:*

  (i) *Independently of his attitude towards risk, student s prefers truthful play* **P** *to strategic play* $\hat{\mathbf{P}}$ *if $-\bar{r}_s(\mathbf{P}, \cdot)$ first-order stochastically dominates $-\bar{r}_s(\hat{\mathbf{P}}, \cdot)$. He prefers* $\hat{\mathbf{P}}$ *to* **P** *if the converse holds.*

  (ii) *If student s is risk averse, he prefers truthful play* **P** *to strategic play* $\hat{\mathbf{P}}$ *if $-\bar{r}_s(\mathbf{P}, \cdot)$ second-order stochastically dominates $-\bar{r}_s(\hat{\mathbf{P}}, \cdot)$. He prefers* $\hat{\mathbf{P}}$ *to* **P** *if the converse holds.*[23]

---

[23] For two cumulative distributions of average ranks, say F and G, with ranks distributed on $\left[\underline{\rho}, \overline{\rho}\right]$, F second-order stochastically dominates G if and only if $\int_x^{\overline{\rho}} (1 - F(x))\, dx \leq \int_x^{\overline{\rho}} (1 - G(x))\, dx$ for all $x \in \left[\underline{\rho}, \overline{\rho}\right]$. The difference versus the usual formula (Gollier 2001; Section 3.2) is due to the fact that lower is better.

(iii) *If student s is risk neutral, he prefers truthful play* **P** *to strategic play* **P̂** *if* $\sum_\lambda \bar{r}_s(\mathbf{P}, \lambda) \leq \sum_\lambda \bar{r}_s(\hat{\mathbf{P}}, \lambda)$. *He prefers* **P̂** *to* **P** *if the converse holds.*

Another special case of additive preferences is lexicographic preferences, which can be seen as the other extreme from average-rank preferences. Formally, we define lexicographic preferences as the limit case of additive risk-neutral preferences where $v_s(c)/v_s(c') \to \infty$ for any cP_sc'. The HBS administration implicitly assumes lexicographic preferences when they evaluate the performance of the mechanism by the number of students who get their favorite course. Lexicographic preferences also generate a complete order over random allocations, and yield the following comparison result.

COMPARISON RESULT 4 (Lexicographic Preferences): *Assume student s has lexicographic preferences. He prefers truthful play* **P** *to strategic play* **P̂** *if he gets his first choice course strictly more often under* **P** *than under* **P̂**, *or if he gets each of his n favorite courses equally as often under both profiles but gets his n + 1th favorite course strictly more often under* **P**, *for some n. He prefers* **P̂** *if the reverse relationship holds.*

To implement Comparison Results 2–4, we run the HBS draft mechanism for 10,000 random priority orders over students using both truthful and strategic preferences, and record each student's distribution over outcomes. Table 5 reports the results.

The results confirm the basic asymmetry between the benefits and costs of strategic play: by each comparison criterion, strategic play harms more students than it benefits. This asymmetry is especially acute when preferences are lexicographic because it is impossible to overreport one's favorite course.

D. *Consequences on Ex Ante Social Welfare*

We now turn to social welfare. In general, we cannot Pareto rank truthful play and strategic play based on the assumption of responsive preferences alone: some students prefer strategic play and others prefer truthful play. So in this section we impose some inter-personal comparability of utilities and assume additive preferences and a utilitarian social planner. The "social" analogues of Comparison Results 2–4 are as follows:

COMPARISON RESULT 5 (Additive Preferences, Risk Neutral): *Assume that students have additive preferences and are risk neutral. Society prefers truthful play* **P** *to strategic play* **P̂**, *if, for all* $n \in \mathbb{N}$, *the expected number of top-n courses allocated to students under* **P** *exceeds that under* **P̂**. *Society prefers* **P̂** *if the converse holds.*

COMPARISON RESULT 6 (Average-rank Preferences): *Assume students have average-rank preferences*:

(i) *Independently of students' attitude towards risk, society prefers truthful play* **P** *to strategic play* **P̂** *if* $-\bar{r}_.(\mathbf{P}, \cdot)$ *first-order stochastically dominates* $-\bar{r}_.(\hat{\mathbf{P}}, \cdot)$. *Society prefers* **P̂** *if the converse holds. (The notation* $\bar{r}_.(\mathbf{P}, \cdot)$ *indicates that the distribution is taken over priority orders and students.)*

Table 5—Individual Preferences over Play of the HBS Draft Mechanism using CR1–4

| | Assumption on preferences | | | | | |
| | Responsive | Additive | Average-rank | | | Lexicographic |
	Any risk attitude (CR1)	Risk neutral (CR2)	Any risk attitude (CR3(i))	Risk averse (CR3(ii))	Risk neutral (CR3(iii))	Risk neutral (CR4)
Outcome (percent)						
Prefers HBS truthful	41	44	49	60	65	86
Prefers HBS strategic	10	10	18	20	34	13
Indifferent	1	1	1	1	1	1
Indeterminate	48	46	32	18	0	0

(ii) *If students are risk averse, society prefers truthful play* $\mathbf{P}$ *to strategic play* $\hat{\mathbf{P}}$ *if* $-\bar{r}_.(\mathbf{P},\cdot)$ *second-order stochastically dominates* $-\bar{r}_.(\hat{\mathbf{P}},\cdot)$. *Society prefers* $\hat{\mathbf{P}}$ *if the converse holds.*

(iii) *If students are risk neutral, society prefers truthful play* $\mathbf{P}$ *to strategic play* $\hat{\mathbf{P}}$ *if* $\sum_\lambda \sum_s \bar{r}_s(\mathbf{P},\lambda) < \sum_\lambda \sum_s \bar{r}_s(\hat{\mathbf{P}},\lambda)$. *Society prefers* $\hat{\mathbf{P}}$ *if the converse holds.*

COMPARISON RESULT 7 (Lexicographic Preferences): *Assume students have lexicographic preferences. Society prefers truthful play* $\mathbf{P}$ *to strategic play* $\hat{\mathbf{P}}$ *if the expected number of students who get their first choice course is strictly higher under* $\mathbf{P}$ *than under* $\hat{\mathbf{P}}$, *or if the expected number of students who get their* 1st, ..., *nth favorite courses is the same under both strategy profiles, for some n, but the expected number of students who get their* $n+$ 1th *favorite course is strictly higher under* $\mathbf{P}$. *Society prefers* $\hat{\mathbf{P}}$ *if the reverse relationship holds.*

Figure 1 shows the average number of courses that students get among their top *n* choices. There is a first-order stochastic dominance relationship between the distribution of outcomes under truthful and strategic play: students get more of their favorite course, more of their top two favorite courses and so on under truthful play than under strategic play.[24] Thus CR5 obtains and by consequence CR6(iii) and CR7 obtain as well since both are special cases of risk-neutral additive preferences. In other words, if students are risk neutral, a utilitarian social planner unambiguously prefers truthful play of the HBS draft. The difference is economically meaningful. 82 percent of students receive their favorite course under truthful play, and they receive 2.46 of their top three courses, versus 63 percent and 1.99 under strategic play.

Figure 2 plots the distribution of the average rank of course allocations in the population. There is a bit more mass at the very best outcomes under strategic play than under truthful play; this is due to the targeted opportunism of students who are fortunate to mainly like unpopular courses. For this reason truthful play does not first-order stochastically dominate strategic play, but second-order stochastic

---

[24] The kink in the HBS Truthful line at rank six is a mechanical effect due to the way we construct truthful preferences. Students report their top-five truthful preferences in the May Poll. Their sixth favorite course is the first course they rank in the strategic rank order list that they did not rank in the May Poll. If this course is rated highly by many other students in the May Poll, then the student will never obtain it under truthful play, but might obtain it under strategic play if he ranks it highly enough.
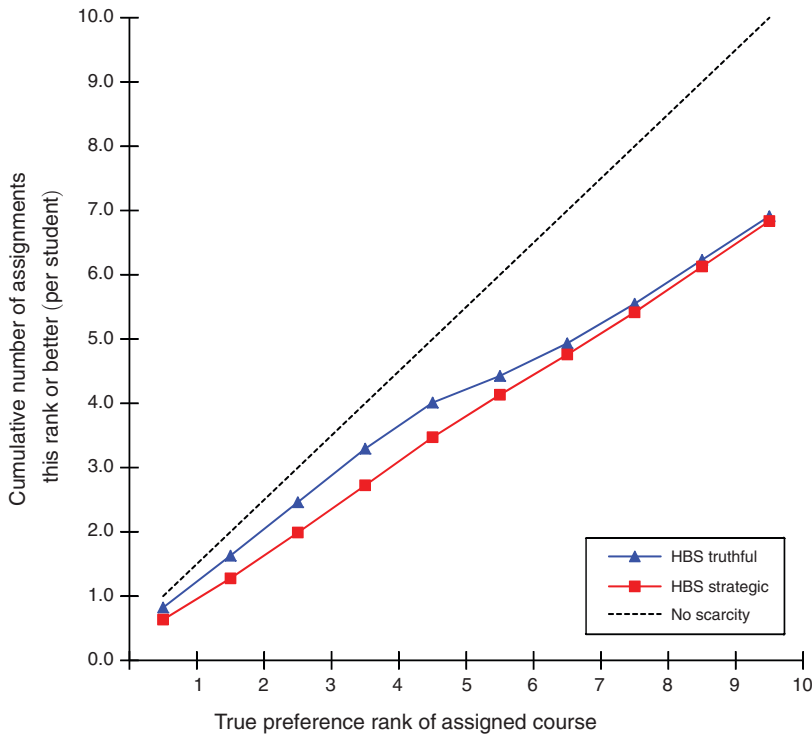
FIGURE 1. CUMULATIVE DISTRIBUTION OF THE TRUE PREFERENCE RANK OF STUDENTS' ASSIGNED COURSES: TRUTHFUL VERSUS STRATEGIC PLAY OF THE HBS DRAFT MECHANISM

Note: The distribution under truthful play first-order stochastically dominates that under strategic play, so CR5, CR6(iii), and CR7 obtain.

dominance does obtain. Thus, strategic behavior hurts ex ante social welfare if students have average-rank preferences and are weakly risk averse.

VI. Comparison of the HBS Draft to the Strategyproof Alternative

In the previous section we showed that strategic play of the HBS draft mechanism harms efficiency, assessed either ex ante or ex post. To get a different perspective on the costs of strategic behavior, we now compare the HBS draft with its strategyproof alternative, the random serial dictatorship (RSD). In the RSD, students are randomly ordered and, in a single round, take turns choosing their *entire* bundle of courses out of those courses still with remaining capacity.

The first thing to note is that RSD is ex post efficient, whereas we found in Section VB that the HBS draft is highly inefficient ex post. The second thing to note is that we will need to impose additional structure on preferences beyond responsiveness in order to assess ex ante welfare. Under RSD, students will often obtain their ideal bundle of courses, but will also often obtain a very poor bundle. The responsiveness assumption does not rule out the possibility that a student only places value on obtaining his ideal bundle, nor does it rule out that the student only cares about maximizing the minimum bundle he obtains.
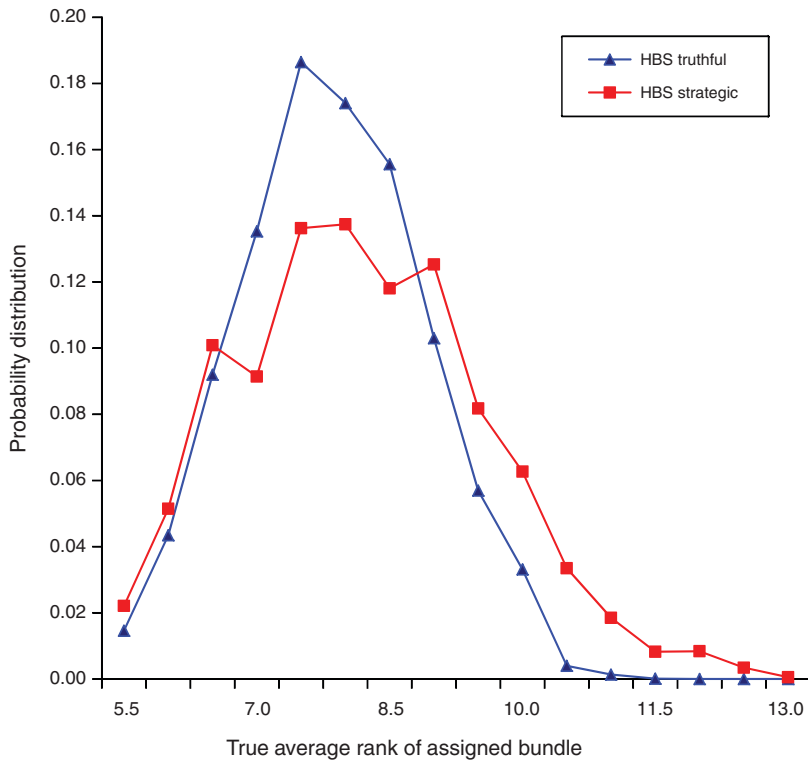
FIGURE 2. PROBABILITY DISTRIBUTION OF THE TRUE-PREFERENCE AVERAGE RANK OF STUDENTS' ASSIGNED BUNDLES: TRUTHFUL VERSUS STRATEGIC PLAY OF THE HBS DRAFT MECHANISM

Note: The distribution under truthful play second-order stochastically dominates that under strategic play, so CR6(ii), and CR6(iii) obtain.

TABLE 6—INDIVIDUAL PREFERENCES BETWEEN HBS AND RSD: CR1–4

	Assumption on preferences					
	Responsive	Additive	Average-rank			Lexicographic
	Any risk attitude (CR1)	Risk neutral (CR2)	Any risk attitude (CR3(i))	Risk averse (CR3(ii))	Risk neutral (CR3(iii))	Risk neutral (CR4)
Outcome (percent)						
Prefers RSD	0	0	0	0	11	20
Prefers HBS strategic	2	31	3	89	89	80
Indifferent	0	0	0	0	0	0
Indeterminate	98	69	97	11	0	0

A. *Ex Ante Individual Welfare*

We repeat the methodology of Section VC. Table 6 compares HBS to RSD under additive, average rank, and lexicographic preferences using Comparison Results 1–4:

Table 6 confirms that without assumptions on students' risk preferences, the comparison is almost entirely indeterminate (columns CR1 and CR3(i)). If we assume that students are risk averse and have average-rank preferences (column CR3(ii)), then most of the indeterminacies are resolved in favor of the HBS draft. No student
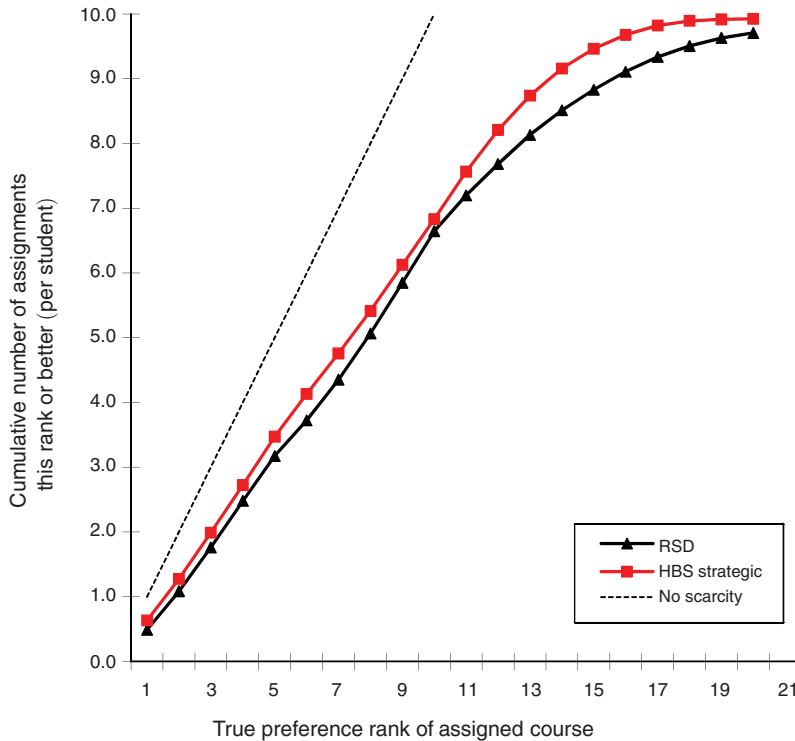
FIGURE 3. CUMULATIVE DISTRIBUTION OF THE TRUE PREFERENCE RANK OF STUDENTS' ASSIGNED COURSES: RSD VERSUS STRATEGIC PLAY OF THE HBS DRAFT MECHANISM

Note: The distribution under HBS first-order stochastically dominates that under RSD, so CR5, CR6(iii), and CR7 obtain.

unambiguously prefers RSD and the vast majority unambiguously prefer HBS. This reflects the well-known criticism of RSD that it exposes students to risk.

There is, however, more in Table 6 than risk. Even if students are risk neutral (columns CR2, CR3(iii), and CR4), the large majority of students prefer the HBS draft to RSD.

B. *Ex Ante Social Welfare*

We repeat the methodology of Section VD. Figure 3 compares the aggregate rank distributions of HBS and RSD. The distribution under strategic play of the HBS draft first-order stochastically dominates that under truthful play of RSD. Thus, a utilitarian social planner prefers HBS to RSD when students are risk neutral and have any additive preferences (CR5, CR6(iii), CR7). This result confirms the picture at the individual level but is more surprising at the social level, given RSD's ex post efficiency. It suggests that RSD's ex post efficiency is not a good proxy for ex ante social welfare.

The magnitudes are of the most economic importance in the tails. Students receive their favorite course with 63 percent probability under HBS, but with only 49 percent probability under RSD. Students receive roughly the same number of
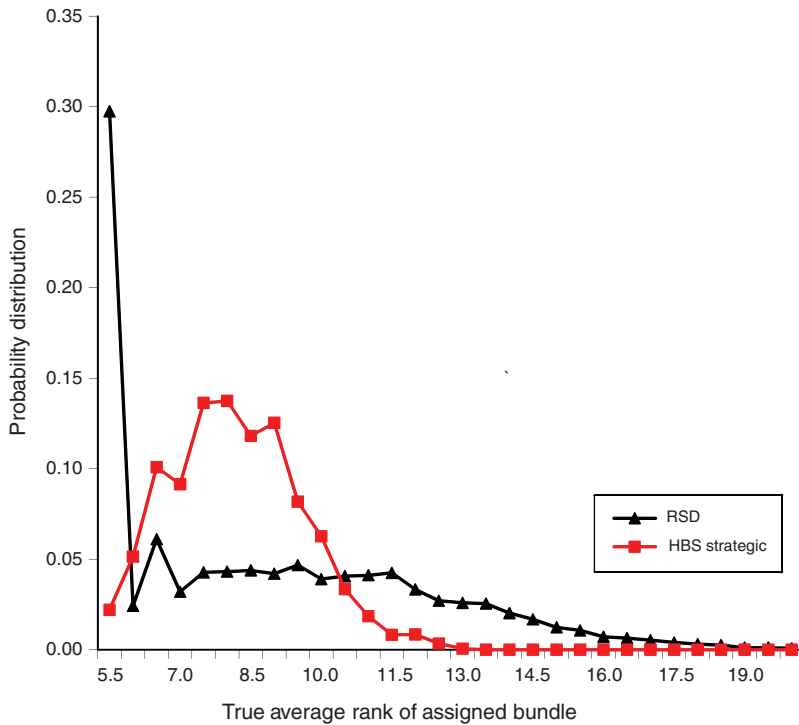
FIGURE 4. DISTRIBUTION OF THE AVERAGE RANK RECEIVED:
RSD VERSUS STRATEGIC PLAY OF THE HBS DRAFT MECHANISM

Note: The distribution under HBS second-order stochastically dominates that under RSD, so CR6(ii) and CR6(iii) obtain.

their second–tenth favorite courses under the two mechanisms, but then receive more courses ranked lower than tenth under RSD than HBS. As a result, the mean average rank under RSD is 8.74, versus 7.99 under HBS. This is an economically meaningful difference, around twice that of the difference between truthful and strategic play of the HBS draft.

Finally, Figure 4 illustrates the risk to which RSD exposes students by examining the distribution of the average ranks of their received bundles. RSD puts much more weight on the tails of the distribution, and indeed is second-order stochastically dominated by HBS (CR6 (ii)). Thus, a utilitarian social planner prefers HBS to RSD if students have average-rank preferences and are weakly risk averse.

VII. Why is There no Efficiency-Fairness Tradeoff?

In the previous section we found that even a utilitarian social planner, unconcerned with fairness per se, prefers the draft to the dictatorship. Moreover, this was not just because RSD exposes students to risk; the draft was preferable even if students are risk neutral.

In this section we provide a simple explanation for this finding: the ex post unfairness of RSD (in a way made precise below) directly hurts its ex ante welfare performance when students are either risk neutral or risk averse. We illustrate the

main idea of this section with an example before articulating more formally how the distribution of outcomes affects ex ante welfare.

EXAMPLE 3 (No Efficiency-Fairness Tradeoff): *There are four courses, with $q_1 = q_2 = q_3 = q_4 = $ ½. Proportion ½ of students are of type $P_1 : c_1, c_2, c_3, c_4$, and proportion ½ are of type $P_2 : c_2, c_1, c_4, c_3$. Students require two courses each.*

In this environment, truthful play is an equilibrium of the HBS draft (this will follow from Lemma 1 below). In equilibrium, all P_1 types ask for and receive c_1 in the first round, while all P_2 types ask for and receive c_2 in the first round. Both c_1 and c_2 reach capacity at time 1. In the second round P_1 types obtain c_3 and P_2 types obtain c_4.

Truthful play is always an equilibrium of RSD. In equilibrium, students who draw priority numbers in $[0, ½]$ obtain $\{c_1, c_2\}$, regardless of whether they are type P_1 or type P_2. The remaining students obtain $\{c_3, c_4\}$, again, regardless of type.

If students have additive preferences (cf. Section VC), then the distribution of bundle values each student receives under HBS second-order stochastically dominates that received under RSD: under HBS all students get their first and third favorite courses, whereas under RSD half get their first and second favorites while the other half get their third and fourth favorites. Thus, if students are risk neutral or risk averse, all individual students and society as a whole strictly prefer HBS to RSD.

Mechanically, what drives RSD's performance in Example 3 is that each student with priority in $[0, ½]$ uses her *second* choice to take the *first* choice course of some student with priority in $[½, 1]$. If students have additive preferences and are risk neutral or risk averse, this is bad for welfare. Further, the magnitudes are large: changing from RSD to HBS increases the proportion of students who get their single favorite course from 50 percent to 100 percent, and improves the mean average rank from 2.5 to 2.0. This echoes our empirical findings in Section VI.

We now generalize Example 3 to show the welfare benefits of imposing a simple ex post fairness condition on the distribution of outcomes. First, we consider the class of anonymous multi-unit random priority mechanisms that take students' ROLs over individual courses as their inputs. Each mechanism is characterized by an atomless joint probability distribution with support on $\{(t_1, \ldots, t_m) \in [0, m]^m$ such that $t_1 \leq t_2, \ldots, \leq t_m\}$, from which students independently and identically draw m choosing times. Then, starting at time $t = 0$ and continuing until time $t = m$, each student with choosing time t is allocated her most-preferred available course (based on her submitted ROL) that she has not yet received and that has remaining capacity. A student may have multiple choices at the same time. A formal description is provided in the online Appendix.

We draw the following distinction amongst random priority mechanisms.

DEFINITION 3 (Callousness): *An anonymous multi-unit random priority mechanism is callous if there exists $n \in \{2, \ldots, m\}$ and a random priority draw, such that a strictly positive measure of students get their nth choosing time before another set of students get their $(n - 1)$th choosing time. Otherwise it is non-callous.*

RSD is an extreme example of a callous mechanism, because for any two students, one takes all her turns before the other takes any. A less-obvious example is

the random priority mechanism in which students draw m choosing times uniform randomly from $[0, m]$. The HBS draft mechanism is non-callous, as are most other draft mechanisms observed in the field.

Callousness can be seen as an ex post fairness criterion for random priority mechanisms because it places constraints on realized choosing times. It is a weak criterion, however. For example, the draft mechanism where students choose courses one at a time, *but in the same order in every round*, is also non-callous, even though students who are early in each round clearly have a better set of choosing times than students who are late in each round.[25]

We next generalize the environment of Example 3 as follows. Suppose that the set of courses can be partitioned into categories, $\mathcal{C}_1, \ldots, \mathcal{C}_L$ with $\mathcal{C} = \mathcal{C}_1 \cup, \ldots, \cup \mathcal{C}_L$ such that (i) each course in $\mathcal{C}_j$ has the same capacity, q_j (by slight abuse of notation), with $q_j | \mathcal{C}_j |$ an integer; (ii) students prefer any course in $\mathcal{C}_j$ to any course in $\mathcal{C}_{j'}$ if $j < j'$; and (iii) students' ordinal preferences over courses within each set $\mathcal{C}_j$ are distributed uniformly. Such preferences are called "block correlated" by Coles, Kushnir, and Niederle (2011). Block correlation allows for some kinds of realistic correlation in students' preferences (e.g., all students prefer high-quality professors to low-quality professors) while enabling the following useful characterization of equilibrium:

LEMMA 1: *If students' preferences are block-correlated, truthful play is an equilibrium of any anonymous multi-unit random priority mechanism.*

Lemma 1 is proved in the online Appendix. The key is that all courses in the same category reach capacity at exactly the same time when students play truthfully. This makes it useless to manipulate the relative ordering of courses within a category. (Note that a direct corollary is that truthful play is an equilibrium in the HBS mechanism when preferences are block-correlated).

Our next result provides a welfare comparison between callous and non-callous mechanisms:

THEOREM 4 (Welfare Costs of Callousness): *Consider any course-allocation environment in which preferences are block correlated. Suppose that students have additive preferences. Under equilibrium truthful play of any non-callous mechanism and any callous mechanism:*

(i) *For each individual student and for society as a whole, the distribution of bundle values under the non-callous mechanism second-order stochastically dominates that under the callous mechanism. Further, for each callous mechanism there exists an environment such that the dominance is strict for every non-callous mechanism.*

(ii) *For each individual student and for society as a whole, there is not a strict first-order stochastic dominance relationship between the distribution of*

---

[25] Drafts with the same order in every round are used to allocate new players to professional sports teams in the United States. For example see www.nfl.com/draft for a description of professional football's draft, and en.wikipedia.org/wiki/Major_League_Baseball_Draft for a description of professional baseball's draft.

bundle values under the non-callous mechanism and that under the callous mechanism.

The environment of Theorem 4 is admittedly special, but part (i) is striking nonetheless: if students are risk neutral or risk averse, any non-callous mechanism is ex ante preferable to any callous mechanism, both for each individual student and for society as a whole.

We end with three remarks. First, while callousness is about distribution, we stress that it is not simply about some mechanisms yielding a more risky distribution of outcomes than others. Theorem 4 shows that there are welfare costs of callousness for risk-neutral students, and Example 3 shows that these can be sizeable. Aversion to risk only exacerbates the unattractiveness of callous mechanisms like RSD.

Second, we clarify the relationship between callousness and Bogomolnaia and Moulin's ((BM) 2001) critique of RSD. In the single-unit demand setting, BM show that the distribution over outcomes under RSD is first-order stochastically dominated. Theorem 4(ii) implies that RSD is not first-order stochastically dominated by any non-callous mechanism. Instead, callousness is about second-order stochastic dominance. If students are sufficiently risk loving they prefer callous mechanisms like RSD to non-callous ones like HBS. Another difference versus BM is the magnitude of the effect. Che and Kojima (2010) show theoretically that the inefficiency BM address goes to zero as the market grows large, and Pathak (2006) finds small magnitudes empirically. This contrasts with the large magnitudes we found in Example 3 and Section VI, respectively.

Finally, it is straightforward to show that dictatorships are the only random priority mechanisms for which truthful reporting is a dominant strategy.[26] Thus, within this class of mechanisms, callousness can be interpreted as a cost of strategyproofness.

VIII. Designing New Multi-Unit Assignment Mechanisms

Our analysis thus far has highlighted that strategic behavior and ex post fairness each matter for welfare in multi-unit assignment. In this section, we explore the performance of a new draft mechanism that integrates these two lessons. Section VIIIA describes the mechanism, and Section VIIIB describes its empirical performance.

A. *The Proxy Draft*

The proxy draft is similar to the HBS draft, but for two modifications. First, strategic play is centralized: students report their preferences to a proxy, which plays strategically on their behalf. Second, the timing of the HBS draft is modified in a simple but important way: the proxy observes the *realization* of the student's random

---

[26] Fix arbitrary $s, P_s : c_1, c_2, \ldots$. In any non-dictatorial mechanism, there exists some $n \in \{2, \ldots, m\}$ such that, with strictly positive probability, s's $(n-1)$th choosing time is strictly earlier than her nth choosing time; denote these times by $t_{n-1} < t_n$. Set $\hat{P}_{-s}$ such that the only course out of s's top-m favorites to reach capacity is c_n, and that the time at which it reaches capacity is strictly between t_{n-1} and t_n. Then $\hat{P}_s : c_n, c_1, c_2, \ldots, c_{n-1}, c_{n+1}, \ldots,$ is a profitable manipulation for student s, and for any other student with the same ordinal preferences over individual courses as student s. This contradicts strategyproofness. The argument and intuition when the number of students is finite is analogous. The only difference is that run-out times will generally be stochastic in discrete environments.

priority before deciding on her play. These modifications are intended to mitigate, respectively, the strategic mistakes we observed in Section IVB, and the incentive for strategic risk-taking we observed in Section IIB.

We describe the proxy draft formally in the context of the continuum economy of Sections I and II. First, each student s reports a rank-order list over individual courses, denoted P_s, to the computer, which will act as her strategic proxy. Second, students are randomly ordered exactly as in the original HBS draft. Specifically, each student draws a random priority number independently from the uniform distribution on $[0, 1]$, and a student who draws priority x gets choosing time x in the first round, $2 - x$ in the second round, $2 + x$ in the third round, etc. Third, the computer looks for a profile of strategies, $\hat{\mathbf{P}}$, that constitutes a Nash equilibrium in ROLs in the complete information draft game induced by students' preferences and the realized priority order. Last, the computer plays this draft game on students' behalf using play $\hat{\mathbf{P}}$. As with the original HBS draft, any strategy profile $\hat{\mathbf{P}}$ generates deterministic run-out times, which we denote for future reference by t^*.

We show as Lemma 2 (in the online Appendix) that s's ordinal preferences over individual courses are sufficient for the proxy to best respond on her behalf. The reason we do not need cardinal preference information to calculate best responses—in contrast to the original HBS draft—is the timing modification. The proxy knows exactly where s is in the priority order, so for any course run-out times t^* the payoff to each strategy is deterministic. The proof that we need only ordinal preferences over individual courses, as opposed to over bundles, is by construction using a greedy algorithm.

The formal sense in which the proxy draft reduces the risk of strategic mistakes is:

THEOREM 5 (Proxy Draft Incentives): *Truthful reporting is a weakly dominant strategy under the proxy draft.*

The intuition is that students in the continuum are "run-out time takers" (i.e., they cannot affect t^*), and for any realized t^* each student can do no better than to report her preferences truthfully and let the proxy optimize on her behalf. By contrast, in the HBS draft, optimal strategic play depends on correct beliefs about course run-out times, a detailed understanding of the rules of the game, and potentially complex calculations involving lotteries.

The formal sense in which the proxy draft mitigates strategic risk taking is:

THEOREM 6 (Proxy Draft Ex Post Efficiency): *For any profile of students' ordinal preferences over individual courses, there exists a profile of students' preferences over bundles that is compatible with responsiveness and for which the equilibrium proxy draft allocation is ex post Pareto efficient.*

In other words, the proxy draft does not leave Pareto-improving one-for-one trades on the table. By contrast, the HBS draft leaves such trades on the table both in equilibrium (cf. Example 2) and in the data (cf. Section VB).

Note that while the proxy draft mitigates both strategic mistakes and strategic risk taking, it does not eliminate all costs of strategic behavior. In particular, we still expect congestion due to overreporting of popular courses by the proxy.

B. *Welfare Performance of the Proxy Draft*

We now explore the performance of the proxy draft in our data. We make two simplifying assumptions for the empirical analysis. First, for computational tractability, we look for a fixed point in course run-out times rather than in student ROLs.[27] Second, we assume that students report their preferences truthfully to the proxy. While truthful reporting is a dominant strategy in the continuum economy, it is only approximately optimal in large finite economies.[28]

Following the methodology of Section V, we draw 100 random priority orders and for each order run the proxy draft and both truthful and strategic play of the HBS draft. At both the individual level and the social level, on almost all of our measures, the proxy draft "lands in between" truthful and strategic play of the HBS draft: the proxy draft is ex ante preferable to the actual play of the HBS draft, but non-equilibrium truthful play of the HBS draft would be better still.

We highlight two specific results. First, the societal distribution of assigned course-ranks (Comparison Result 5) under the proxy draft first-order stochastically dominates that under strategic play of the HBS draft, but is itself dominated by truthful play of the HBS draft. Second, the societal distribution of assigned bundle average-ranks (Comparison Result 6) under the proxy draft second-order stochastically dominates the distribution under strategic play of the HBS draft, but is itself dominated by truthful play of the HBS draft. Full results at both the individual and societal level are reported in the online Appendix.

These "lands in between" results make sense because the proxy draft eliminates some but not all costs of strategic behavior. To give a sense of magnitudes, the average rank of students' assigned bundles under the proxy draft is equal to 7.84, as compared to 7.66 under truthful play of the HBS draft, 7.99 under the actual strategic play of the HBS draft, and 8.74 under RSD.

While we make no claim here as to the optimality of the proxy draft mechanism, these results suggest that the lessons from our analysis are easily implementable.[29] They also provide a perspective on the relative importance of the distribution of outcomes and the incentive constraint on welfare. The proxy draft mechanism circumvents two avoidable costs of strategic behavior, namely strategic mistakes and the ex post inefficiency due to strategic risk-taking. It does not eliminate all strategic behavior since the computer acts strategically on students' behalf. The fact

---

[27] That is, we look for a set of course run-out times t^* such that, when the proxy optimizes on each student's behalf given t^*, the times t^* are in fact realized. The computational procedure we use for finding the fixed point is described in the online Appendix. In the continuum economy, there is no distinction between finding a fixed point in run-out times and finding a Nash equilibrium in ROLs, but the two only approximate each other in large finite economies, because run-out times are not necessarily exogenous from the perspective of each student. The probability that an individual student can affect run-out times decreases as the number of students increases and as course capacities increase.

[28] Again, the issue is that course run-out times are not necessarily exogenous from the perspective of each individual student.

[29] Several other papers have proposed new multi-unit assignment mechanisms in response to earlier versions of this paper, which highlighted strategic behavior and ex post fairness as important design issues but did not propose a specific mechanism. See Budish (2011) and Budish, Che, Kojima, and Milgrom (forthcoming) for approaches based on competitive equilibrium from equal incomes. See Kominers, Ruberry, and Ullman (2010) for an approach that, like ours, combines the HBS draft with a kind of proxy agent; whereas our proxy agent best responds to realized run-out times (which involves finding a fixed point), the Kominers, Ruberry, and Ullman proxy agent follows a heuristic analogous to straightforward bidding in a simultaneous ascending auction.

that the average rank of the assigned bundles under the proxy draft mechanism is close to the average rank of the assigned bundles under (non-equilibrium) truthful behavior suggests that the cost of the incentive constraint is small relative to the welfare benefit of the draft's more equal distribution of outcomes.

IX. Conclusion

Economists have increasingly played an active role in designing solutions to real-world resource allocation problems like course allocation. The usual starting point for this work is a mechanism that is attractive in a stylized theoretical environment, which is then adapted to some of the complexities that arise in practice. One well-known example that fits this pattern is the redesign of the market for medical residencies in the United States, on which Roth (2002) remarks that "the simple theory of Gale and Shapley (1962) offered a surprisingly good guide to the design of the Roth-Peranson (1999) algorithm." Other prominent examples that fit this pattern include the design of school-choice procedures, advertising markets on internet search engines, and combinatorial auctions.[30]

Here, the relationship between theory and practice is different, and invited a different approach. The theory literature on multi-unit assignment has yielded mostly negative results; there is no widely admired "classic" solution akin to Gale and Shapley (1962) for matching, or Vickrey (1961) for auctions. In light of all the negative findings, several theory papers concluded that dictatorships are the best possible solution, because they are strategyproof and ex post Pareto efficient. Practitioners meanwhile designed their own mechanisms, on an ad hoc basis without guidance from economic theory. It seemed to us worthwhile to study the design choices practitioners made, to see what theory and practice could teach each other.

Studying the HBS draft mechanism in isolation highlighted what theory could teach practice: the importance of incentives in practical market design. We showed that the HBS draft is simple to manipulate in theory, is indeed manipulated by students in practice, and that these manipulations cause significant welfare loss. These findings constitute some of the most direct evidence to date on the costs of using a manipulable mechanism.

Comparing the draft from practice to the dictatorship from theory uncovered what the prior theory had missed: the link between fairness and welfare in multi-unit assignment. A natural prior based on reading the extant theory is that the dictatorship is less fair than the draft, but nevertheless the right design choice because the manipulability of the draft leads to inefficiency. Instead, we found that the draft's ex post fairness relative to the dictatorship—its "non-callousness"—causes it to be better for ex ante welfare than the dictatorship. These findings serve as a useful reminder that ex post Pareto efficiency can be a very poor proxy for ex ante welfare, and, quite simply, that strategyproofness has costs as well as benefits.

---

[30] New school-choice procedures in Boston and New York are based on an adaptation of Gale and Shapley's deferred acceptance algorithm (Abdulkadiroğlu and Sönmez 2003). Google cites the influence of Vickrey's (1961) "Nobel Prize-winning economic theory" in the design of its auction for advertising slots (Edelman, Ostrovsky, and Schwarz 2007). On the relationship between theory and practice for combinatorial auction design, see Cramton et al. (2006) and Milgrom's (2004) aptly named text "Putting Auction Theory to Work."

Thus, the mechanisms from theory and practice each have flaws, and both theory and practice each have something to teach the other. A particularly satisfying result of this two-way dialogue is that it ultimately yielded new solutions to the multi-unit assignment problem, improving welfare relative both to what theory previously had to offer and to what practitioners had designed on their own.

REFERENCES

Abdulkadiroğlu, Atila, Yeon-Koo Che, and Yosuke Yasuda. 2008. "Expanding 'Choice' in School Choice." http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1308730.

Abdulkadiroğlu, Atila, Parag A. Pathak, and Alvin E. Roth. 2009. "Strategy-Proofness versus Efficiency in Matching with Indifferences: Redesigning the NYC High School Match." *American Economic Review* 99 (5): 1954–78.

Abdulkadiroğlu, Atila, and Tayfun Sönmez. 2003. "School Choice: A Mechanism Design Approach." *American Economic Review* 93 (3): 729–47.

Bartlett, Thomas. 2008. "Class Warfare: When Getting in is the Hardest Part." *Chronicle of Higher Education,* February 15.

Bergemann, Dirk, and Stephen Morris. 2005. "Robust Mechanism Design." *Econometrica* 73 (6): 1771–1813.

Bogomolnaia, Anna, and Hervé Moulin. 2001. "A New Solution to the Random Assignment Problem." *Journal of Economic Theory* 100 (2): 295–328.

Brams, Steven J., and Philip D. Straffin, Jr. 1979. "Prisoners' Dilemma and Professional Sports Drafts." *The American Mathematical Monthly* 86 (2): 80–88.

Brams, Steven J., and Alan D. Taylor. 1996. *Fair Division: From Cake-cutting to Dispute Resolution*. New York: Cambridge University Press.

Brams, Steven J., and Alan D. Taylor. 1999. *The Win-Win Solution: Guaranteeing Fair Shares to Everybody*. New York: W.W. Norton and Co.

Budish, Eric. 2011. "The Combinatorial Assignment Problem: Approximate Competitive Equilibrium from Equal Incomes." *Journal of Political Economy* 119 (6): 1060–1103.

Budish, Eric, and Estelle Cantillon. 2012. "The Multi-unit Assignment Problem: Theory and Evidence from Course Allocation at Harvard: Dataset." *American Economic Review*. http://dx.doi.org/10.1257/aer.102.5.2237.

Budish, Eric, Yeon-Koo Che, Fuhito Kojima, and Paul Milgrom. Forthcoming**.** "Designing Random Allocation Mechanisms: Theory and Applications." *American Economic Review*.

Che, Yeon-Koo, and Fuhito Kojima. 2010. "Asymptotic Equivalence of Probabilistic Serial and Random Priority Mechanisms." *Econometrica* 78 (5): 1625–72.

Coles, Peter, Alexey Kushnir, and Muriel Niederle. 2011. "Preference Signaling in Matching Markets." http://www.stanford.edu/~niederle/SignalingPaper.pdf.

Cramton, Peter, Yoav Shoham, and Richard Steinberg, eds. 2006. *Combinatorial Auctions*. Cambridge, MA: MIT Press.

Edelman, Benjamin, Michael Ostrovsky, and Michael Schwarz. 2007. "Internet Advertising and the Generalized Second-Price Auction: Selling Billions of Dollars Worth of Keywords." *American Economic Review* 97 (1): 242–59.

Ehlers, Lars, and Bettina Klaus. 2003. "Coalitional Strategy-Proof and Resource-Monotonic Solutions for Multiple Assignment Problems." *Social Choice and Welfare* 21 (2): 265–80.

Ehlers, Lars, and Jordi Massó. 2008. "Matching Markets under (In)complete Information." Unpublished.

Fudenberg, Drew, and Jean Tirole. 1991. *Game Theory*. Cambridge, MA: MIT Press.

Gale, David, and Lloyd Shapley. 1962. "College Admissions and the Stability of Marriage." *American Mathematical Monthly* 69 (1): 9–15.

Gehan, Edmund A. 1965. "A Generalized Wilcoxon Test for Comparing Arbitrarily Singly-Censored Samples." *Biometrika* 52 (1): 203–23.

Guernsey, Lisa. 1999. "Business School Puts Courses in Hands of an On-Line Market." *New York Times,* September 9.

Gollier, Christian. 2001. *The Economics of Risk and Time*. Cambridge, MA: MIT Press.

Hatfield, John William. 2009. "Strategy-Proof, Efficient, and Nonbossy Quota Allocations." *Social Choice and Welfare* 33 (3): 505–15.

Kominers, Scott, Mike Ruberry, and Jonathan Ullman. 2010. "Course Allocation by Proxy Auction." In *Internet and Network Economics: 6th International Workshop, WINE 2010, Stanford, CA, USA December 13–17, 2010, Proceedings*, edited by Amin Saberi, 551–58. New York: Springer.

Lehrer, Jim. 2008. "College Students Squeezed by Rising Costs, Less Aid: Straining the System." *NewsHour*, PBS, December 9, 2008. http://www.pbs.org/newshour/bb/education/july-dec08/collegecosts_12-09.html.

Mas-Colell, Andreu. 1984. "On a Theorem of Schmeidler." *Journal of Mathematical Economics* 13 (3): 201–06.

Milgrom, Paul. 2004. *Putting Auction Theory to Work*. Churchill Lectures in Economics. New York: Cambridge University Press.

Pápai, Szilvia. 2001. "Strategyproof and Nonbossy Multiple Assignments." *Journal of Public Economic Theory* 3 (3): 257–71.

Pathak, Parag. 2006. "Lotteries in Student Assignment." Unpublished.

Roth, Alvin E. 1985. "The College Admissions Problem Is Not Equivalent to the Marriage Problem." *Journal of Economic Theory* 36 (2): 277–88.

Roth, Alvin E. 2002. "The Economist as Engineer: Game Theory, Experimentation, and Computation as Tools for Design Economics." *Econometrica* 70 (4): 1341–78.

Roth, Alvin E. 2008. "What Have We Learned from Market Design?" *Economic Journal* 118 (527): 285–310.

Roth, Alvin E., and Elliott Peranson. 1999. "The Redesign of the Matching Market for American Physicians: Some Engineering Aspects of Economic Design." *American Economic Review* 89 (4): 748–80.

Roth, Alvin E., and Uriel G. Rothblum. 1999. "Truncation Strategies in Matching Markets—In Search of Advice for Participants." *Econometrica* 67 (1): 21–43.

Roth, Alvin E., Tayfun Sönmez, and M. Utku Ünver. 2004. "Kidney Exchange." *Quarterly Journal of Economics* 119 (2): 457–88.

Schmeidler, David. 1973. "Equilibrium Points of Nonatomic Games." *Journal of Statistical Physics* 7 (4): 295–300.

Sönmez, Tayfun, and M. Utku Ünver. 2010. "Course Bidding at Business Schools." *International Economic Review* 51 (1): 99–123.

Vickrey, W. 1961. "Counterspeculation, Auctions, and Competitive Sealed Tenders." *Journal of Finance* 16 (1): 8–37.

Wharton. 2009. "The Course Registration Auction for MBA Electives." Version 16.0. University of Pennsylvania, The Wharton School, Graduate Division, MBA Program Office.

Wilson, Robert B. 1987. "Game-Theoretic Analyses of Trading Processes." In *Advances in Economic Theory: Fifth World Congress*, edited by T. F. Bewley, 33–70. Econometric Society Monographs Series, 12. New York: Cambridge University Press.