Web Appendix for "The Multi-unit Assignment Problem: Theory and Evidence from Course Allocation at Harvard"

Eric Budish and Estelle Cantillon

Forthcoming, *American Economic Review*

June 30, 2011

A Proofs

A.1 Proof of Theorem 1 (Simple Manipulations)

Consider any arbitrary strategy $\widehat{P}_s$ and relabel courses such that $\widehat{P}_s : c_1, c_2, c_3, ..., c_C$. Let $\widehat{P}_s^{c_k \downarrow l}$ denote the ROL that corresponds to $\widehat{P}_s$ except that course c_k is moved down to position l ($l > k$). Denote by λ a realization of the priority order.

Claim 1: Consider $k \leq m$ and suppose c_k is unpopular. Then, for all λ, $\widehat{P}_s^{c_k \downarrow l}$ gets exactly the same courses as $\widehat{P}_s$ or exactly one more course in $\{c_{k+1}, ..., c_l\}$ than $\widehat{P}_s$, at the cost of a course in $\{c_{l+1}, ..., c_C\}$.

Proof of Claim 1: Fix an arbitrary λ. Because $\widehat{P}_s$ and $\widehat{P}_s^{c_k \downarrow l}$ only differ from position k onwards, the game proceeds identically until the time at which $\widehat{P}_s$ requests c_k (and $\widehat{P}_s^{c_k \downarrow l}$ requests course c_{k+1}). Let r_k be the round at which this happens. By construction, c_k is available when $\widehat{P}_s$ requests it in round r_k. Because student s has zero mass, the fact that his outcome in round r_k is different across the two strategies does not affect course seat availabilities and thus, a fortiori, the allocation and requests of other students in any given round.

From round $r_k + 1$ onwards, student s requests courses one round earlier under strategy $\widehat{P}_s^{c_k \downarrow l}$ than under strategy $\widehat{P}_s$, until we either reach a course, say $c_{k'}$, in $\{c_{k+1}, ..., c_l\}$ that student s gets under $\widehat{P}_s^{c_k \downarrow l}$ but not under $\widehat{P}_s$, or reach position l in student s's ROL. We consider each case in turn:

1. There exists $c_{k'}$ in $\{c_{k+1}, ..., c_l\}$ that student s gets under $\widehat{P}_s^{c_k \downarrow l}$ but not under $\widehat{P}_s$.

   Let $r_{k'}$ be the round at which student s requests but does not get this course under $\widehat{P}_s$. From round $r_{k'}$ onwards, student s's requests are in synch under both strategies (i.e. he asks for the same courses in the same round under both strategy profiles) and thus he gets the same outcome until the algorithm reaches position l in his ROL.

When the algorithm reaches the request in position l, student s requests (and gets) course c_k under $\widehat{P}_s^{c_k \downarrow l}$. From then on, student s requests courses one round earlier under $\widehat{P}_s$. This has two possible consequences: either there exists a course that he gets under $\widehat{P}_s$ but not under $\widehat{P}_s^{c_k \downarrow l}$ (after which his requests are in synch and there is no more discrepancy between the two outcomes), or the algorithm reaches round m (and thus the course that the student requests in round m under $\widehat{P}_s$ is never requested by $\widehat{P}_s^{c_k \downarrow l}$). In both cases, there is a single course in $\{c_{l+1}, ..., c_C\}$ that student s gets under $\widehat{P}_s$ instead of $c_{k'}$ that he does not get under $\widehat{P}_s^{c_k \downarrow l}$.

2. The algorithm reaches position l in student s's ROL without any difference in allocations between the two strategies

   At that round, $\widehat{P}_s^{c_k \downarrow l}$ requests c_k and student s's requests become in synch again. There is thus no more difference in outcomes.

Claim 2: Let c_k be the lowest-ranked unpopular course among the top m courses in P_s. Let $\widehat{P}_s^1 = P_s^{c_k \downarrow m}$. Student s is weakly better off using $\widehat{P}_s^1$ than P_s for all λ.

Proof of Claim 2: By claim 1, $\widehat{P}_s^1$ gets exactly the same courses or exactly one additional course in $\{c_{k+1}, ..., c_m\}$ than P_s, at the cost of a course in $\{c_{m+1}, ..., c_C\}$. Because all courses in $\{c_{k+1}, ..., c_m\}$ are strictly preferred to courses in $\{c_{m+1}, ..., c_C\}$, student s is either indifferent or strictly better off using $\widehat{P}_s^1$ (here we are using the fact that preferences are responsive and that students have vNM preferences over lotteries).

Claim 3: Let c_j be the n^{th} lowest unpopular course among the top m courses in P_s. Let $\widehat{P}_s^n = \widehat{P}_s^{n-1 \;\; c_j \downarrow m-n+1}$ (student s downgrades course c_j just above all the other less preferred unpopular courses that he has already downgraded). Student s is weakly better off using $\widehat{P}_s^n$ than $\widehat{P}_s^{n-1}$ for all λ.

Proof of Claim 3: By claim 1, $\widehat{P}_s^n$ gets either exactly the same courses or exactly one additional course among the popular courses that were between c_j and position $m-n+1$ in $\widehat{P}_s^{n-1}$. This comes at the expense of a course in $\{c_{m+1}, ..., c_C\}$. Given that preferences are responsive and take the vNM form, student s is weakly better off using $\widehat{P}_s^n$ over $\widehat{P}_s^{n-1}$.

We continue until there is no further unpopular course to downgrade. At each deviation, student s is weakly better off for all λ. The claim then follows by transitivity.

A.2 Proof of Theorem 2 (Best-Response Characterization)

(i) Consider the alternative strategy $\widetilde{P}_s$ that is identical to $\widehat{P}_s$, except that c is placed in the original position of c' and c' is dropped from the ROL. Partition the set of priority orders according to whether s gets c under $(\widetilde{P}_s, \widehat{\mathbf{P}}_{-s})$. We first argue that, in both cases, students s is at least as well off under $(\widetilde{P}_s, \widehat{\mathbf{P}}_{-s})$ as under $\widehat{\mathbf{P}}$, for all realizations of the priority order. Indeed, if student s gets

c under $(\widetilde{P}_s, \widehat{\mathbf{P}}_{-s})$, the two bundles after the initial allocation differ by exactly one course (student s gets c' under $\widehat{\mathbf{P}}$ and c or some other course under $(\widetilde{P}_s, \widehat{\mathbf{P}}_{-s})$) and student s can always replicate the bundle he gets under $\widehat{\mathbf{P}}$ by asking for c' in the add-drop phase. If, instead, student s does not get c under $(\widetilde{P}_s, \widehat{\mathbf{P}}_{-s})$, he requests all the courses following the failed request for c one round earlier and (using arguments similar to those in the proof of Theorem 1) also gets a bundle at the end of the initial allocation that differs by exactly one course from the bundle he gets under $\widehat{\mathbf{P}}$. Because c' is available in the add-drop phase, student s can ensure that his final bundle is at least as good as the one he gets under $\widehat{\mathbf{P}}$. We next argue that student s is strictly better off using $\widetilde{P}_s$. This follows from the fact that c is a top m course and that student s gets it with strictly higher probability under $(\widetilde{P}_s, \widehat{\mathbf{P}}_{-s})$.

(ii) Consider the alternative strategy $\widetilde{P}_s$ that is identical to $\widehat{P}_s$, except that c is placed in the original position of c' and c' is dropped from the ROL. For all realizations of the priority order, this strategy yields the same course bundle except that c' is replaced by the preferred course c.

(iii) Consider the alternative strategy $\widetilde{P}_s$ that is identical to $\widehat{P}_s$, except that c is placed ahead of c'. Partition the set of priority orders according to whether s gets c under $(\widetilde{P}_s, \widehat{\mathbf{P}}_{-s})$. For all priority orders for which s does not get c, the outcomes under $(\widetilde{P}_s, \widehat{\mathbf{P}}_{-s})$ and under $\widehat{\mathbf{P}}$ are identical. Instead, when student s gets c under $(\widetilde{P}_s, \widehat{\mathbf{P}}_{-s})$, he requests the next courses one round later than under $\widehat{\mathbf{P}}$, which may cause him to miss a course he used to get under the original strategy. As in the proof of Theorem 1, it is easy to estabish that the final bundle will differ by at most one course. Condition (iii) guarantees that student s does not get c at the cost of a more preferred popular course, ensuring (modulo the use of the add-drop phase) that he weakly prefers this bundle, for all realizations of the priority order. Student s is strictly better off using $\widetilde{P}_s$ than $\widehat{P}_s$ because he gets c, a top m course, with higher probability.

A.3 Proof of Lemma 1 (Truthful Play if Preferences are Block-Correlated)

We first consider non-callous mechanisms. As argued formally in Appendix B course run-out times are deterministic. Moreover, if students play truthfully, all courses in $\mathcal{C}_1$ run out at time $q_1|\mathcal{C}_1|$. Likewise, courses in $\mathcal{C}_2$ run-out at time $q_1|\mathcal{C}_1| + q_2|\mathcal{C}_2|$, and so on. Let t_1^*, t_2^*, ... describe the run-out times ($t_1^* < t_2^* < ...$) for the courses in $\mathcal{C}_1$, $\mathcal{C}_2$, ... These run-out times are exogenous from the point of view of a single student.

Because the allocation in non-callous mechanisms takes place over rounds (with all students getting their n^{th} pick before any student gets his $n+1^{th}$ pick), truthful play guarantees that student s gets his top $q_1|\mathcal{C}_1|$ most preferred courses. Any strategy that does not place student s's top $q_1|\mathcal{C}_1|$ courses on top of his ROL is dominated by one that does. Among those that do, the strategy that places all courses in $\mathcal{C}_1$ in truthful order does as well as the others. Truthful play also guarantees that student s gets his best $q_2|\mathcal{C}_2|$ courses out of $\mathcal{C}_2$. The only way to get more than $q_2|\mathcal{C}_2|$ courses out of

$\mathcal{C}_2$ would be to get that many fewer courses from $\mathcal{C}_1$, which is not advantageous by responsiveness. We can repeat this argument all the way until we reach $\mathcal{C}_L$.

We now consider callous mechanisms. Like for non-callous mechanisms, course run-out times are deterministic and equal for all courses within the same course category. Let t_1^*, t_2^*, ... describe the run-out times for the courses in $\mathcal{C}_1, \mathcal{C}_2$, ... We have $t_1^* \leq t_2^* \leq ...$ under truthful play. These run-out times are exogenous from the point of view of a single student. A similar argument as above applies to argue that truthful play is an equilibrium if $t_1^* < t_2^* < ...$. If two run-out times are identical, the argument goes through by merging the two course categories that have the same run-out times and repeating the argument for this new partition of courses with strictly different run-out times.

A.4 Proof of Theorem 4 (Welfare Costs of Callousness)

(i) As argued in the proof of Lemma 1, each student gets for sure his top $q_j|\mathcal{C}_j|$ courses from $\mathcal{C}j$, $j = 1, ..., L$, under truthful play of any non-callous mechanism. In contrast, under any callous mechanism, each student gets $q_j|\mathcal{C}_j|$ courses *in expectation* from $\mathcal{C}_j$, for $j = 1, ..., L$: it is possible that a student sometimes gets more than $q_1|\mathcal{C}_1|$ courses from $\mathcal{C}_1$, in which case the extra courses are at rank $q_1|\mathcal{C}_1| + 1, q_1|\mathcal{C}_1| + 2, ...$; and sometimes that student gets fewer than $q_1|\mathcal{C}_1|$ courses from $\mathcal{C}_1$, in which case the courses she loses are those ranked $q_1|\mathcal{C}_1|, q_1|\mathcal{C}_1| - 1, ...$ Since preferences are strict and additive and students are risk neutral or risk averse, such deviations from getting exactly $q_j|\mathcal{C}_j|$ courses from each course category $\mathcal{C}_j$ are strictly welfare reducing. This yields the weak version of (ii).

For the strict version, we identify an environment for which such strictly welfare reducing events occur with positive probability. Consider an arbitrary callous mechanism. By definition, there exists n such that with strictly positive probability some students get their n^{th} choice before others get their $n-1^{th}$. Set $q_1|\mathcal{C}_1| = n-1$. Now, each student with strictly positive probability gets strictly more than $q_1|\mathcal{C}_1|$ courses from $\mathcal{C}_1$, and since she still gets $q_1|\mathcal{C}_1|$ courses from $\mathcal{C}_1$ in expectation, her welfare is strictly lower than under any non-callous mechanism. Since each individual student is strictly worse off in expectation, social welfare is lower as well.

(ii) The arguments above are sufficient to establish that the distribution of bundle values from a callous mechanism never strictly first-order stochastically dominates the distribution under any non-callous mechanism. For the converse, note that whenever a callous mechanism yields a different distribution over outcomes than a non-callous mechanism, it involves each student sometimes getting strictly more of their most-preferred courses from $\mathcal{C}_1 \cup \mathcal{C}_2 \cup ... \cup \mathcal{C}_j$ for some $j \geq 1$, than under the non-callous mechanism.

A.5 Statement and Proof of Lemma 2 (Proxy Draft Best Response Lemma)

Lemma 2: Best responses in step three of the proxy draft mechanism can be fully derived based on students' ordinal preferences over individual courses.

Proof. Fix $\widehat{\mathbf{P}}_{-s}$ and λ. Renumber courses such that $P_s : c_1, c_2, ...$ and let t_1^*, t_2^*, ..., t_C^* describe the run-out times for courses c_1, c_2, ..., c_C. In the continuum econony these run-out times can be seen as exogenous when deriving individual students' best responses. We now describe an algorithm that constructs a best-response ROL for student s.[1]

<u>Initialization:</u> Let $\{\tau_1, \tau_2, ..., \tau_m\}$ denote the m choosing times of student s given λ, ordered so that $\tau_1 < \tau_2 < ... < \tau_m$. Call this the set of remaining choosing times for student s.

<u>Step 1</u>. If $\tau_1 > t_1^*$, s cannot obtain c_1. Proceed to Step 2. If $\tau_1 \leq t_1^*$, find the latest of s's choosing times that is earlier than t_1^*, denoted $\tau_{f(c_1)}$. Set the $f(c_1)^{th}$ position on s's ROL to c_1, and remove $\tau_{f(c_1)}$ from the set of remaining choosing times.

<u>Step 2</u>. Find the latest of s's remaining choosing times that is earlier than t_2^* and denote it $\tau_{f(c_2)}$ (if there is no such choosing time, proceed to Step 3). Set the $f(c_2)^{th}$ position on s's ROL to c_2, and remove $\tau_{f(c_2)}$ from the set of remaining choosing times.

<u>Step j</u>. Find the latest of s's remaining choosing times that is earlier than t_j^*, denoted $\tau_{f(c_j)}$ (if there is no such choosing time, instead proceed to Step $j+1$). Set the $f(c_j)^{th}$ position on s's ROL to c_j, and remove $\tau_{f(c_j)}$ from the set of remaining choosing times.

Stop when either the set of remaining choosing times is empty, or when all courses on P_s have been exhausted.

We prove by induction that this constructed ROL $\widehat{P}_s$ obtains the best possible bundle for s. Let c' be the first course on P_s that s does not obtain under $\widehat{P}_s$. There is no way to obtain c' without losing some course that is earlier on P_s than c'. Therefore, since preferences are responsive, any achievable bundle of courses that includes c' is dominated by one that does not. Let c'' be the second course on P_s that s does not obtain under $\widehat{P}_s$. By the same argument, any achievable bundle of courses that includes c'' is dominated by one that does not. Continuing in this fashion we conclude that $\widehat{P}_s$ is a best response. We have only used s's ordinal preferences over individual courses to compute $\widehat{P}_s$, as required.

A.6 Proof of Theorem 5 (Proxy Draft Incentives)

Given that run-out times in the draft stage are exogenous to an individual student's (computer-generated) play in the draft stage and thus report in the preference submission stage, the best a student can do is to report truthfully and let the computer optimize on his behalf for each realization of λ.

[1] Best-responses may not be unique but each yields the same bundle.

A.7 Proof of Theorem 6 (Proxy Draft Ex-Post Efficiency)

It is sufficient to consider trades in which each participant in the trade gives and gets a single course. Suppose there exists a Pareto improving trade amongst n types of students. Label the types of students and courses involved in the trade such that students of type s_1 give c_1 to s_2, students of type s_2 give c_2 to s_3,..., students of type s_n give c_n to s_1. Since the trade is Pareto improving, each student prefers the course he receives to the course he gives up. Since in equilibrium students report their preferences truthfully (or misreport in a way that yields an equivalent allocation), this means that each student was unable to select the course he receives in the trade at the time he selected the course he gives up in the trade. But this is a contradiction, because it implies that $t_1^* > t_2^* > ... > t_n^* > t_1^*$, where t_k^* denotes the run-out time of course c_k.

B Formal Description of the HBS Mechanism and other Random Priority Mechanisms

B.1 The HBS Mechanism

The allocation generated by the HBS draft mechanism during the initial allocation phase can be described formally as follows. Let $\mathbf{P}$ denote the ROLs submitted by the students. Let $\mathcal{X}_i \subset 2^{\mathcal{C}}$ describe the set of the subsets of $\mathcal{C}$ of cardinality i and denote by X_i a typical element of this set ($|X_i| = i$). The function Ch maps a student profile (his submitted preferences, P, and the set of courses that he already has, X_i) to his most preferred course among a set of available courses $\mathcal{C}' \subset \mathcal{C}$:

$$Ch_{(P,X_i)}(\mathcal{C}') = \{c \in \mathcal{C}' \backslash X_i | cPc' \text{ for all } c' \in \mathcal{C}' \backslash X_i\}$$

Let $\mu(P, X_{r-1})$ denote the measure of students who submitted P and get the courses in X_{r-1} by the end of round $r-1$ under play $\mathbf{P}$ of the HBS draft (we will argue below that this object is well-defined). Then, leveraging Che and Kojima (2010)'s trick of modeling students' random priorities as iid draws from the uniform distribution, and a weak law of large numbers due to Uhlig (1996), the proportion of students requesting course c within any interval of round r when all courses in $\mathcal{C}' \subset \mathcal{C}$ are available is given by:[2]

$$m_c(\mathcal{C}', r) = \sum_{\{(P,X_{r-1})|c=Ch_{(P,X_{r-1})}(\mathcal{C}')\}} \mu(P, X_{r-1}) \qquad (1)$$

Next, let $\mathcal{C}(t)$ describe the set of courses available at time t, $x_c(t)$ the used capacity of course c at time t, t_c^* the run-out time of course c, and let round(t) be equal to the smallest integer strictly greater than t. Initialize $\mathcal{C}(0) = \mathcal{C}$, $x_c(0) = 0$ and $t_c^* = m$. Let $\mu(P, X_0)$ be equal to the measure of students using strategy P in play $\mathbf{P}$. Finally, we introduce two functions, $\tau_c(.)$ and $\tau(.)$ whose

[2] Recall that both the set of ROLs and the set of courses are finite. There is thus a finite number of pairs (P, X_{r-1}).

value at the different steps of the algorithm will be defined endogenously. We initialize them at $\tau_c(0) = 0$ and $\tau(0) = 0$. Let $v = 1$. Course run-out times, course requests and allocations in the HBS draft are jointly determined iteratively through the following process.

1. For each course c, define the time it would run out if the proportion of students requesting it in any interval remains equal to the current rate, $m_c(\mathcal{C}(\tau(v-1)),\text{round}(\tau(v-1)))$:

   $$\tau_c(v) = \sup\{t \in [\tau(v-1), m] | x_c(v-1) + m_c\left(\mathcal{C}(\tau(v-1)), \text{round}(\tau(v-1))\right)(t - \tau(v-1)) < q_c\}$$

2. Let $\tau(v) = \min\{\min_{c\in\mathcal{C}(\tau(v-1))}\{\tau_c(v)\}, \text{round}(\tau(v-1))\}$ which corresponds to the end of the current round or the earliest run-out time, whichever happens earlier. If $\tau(v) = \tau_c(v)$ for some c, set $t^*_c = \tau_c(v)$.

3. Adjustment of the set of available courses: If $\tau(v) = \tau_c(v)$ for some c, then $\mathcal{C}(\tau(v)) = \mathcal{C}(\tau(v-1))\backslash\{c \in \mathcal{C}(\tau(v-1)) | \tau_c(v) = \tau(v)\}$. Otherwise, $\mathcal{C}(\tau(v)) = \mathcal{C}(\tau(v-1))$.

4. Adjustment of capacities: For all c, $x_c(v) = x_c(v-1) + m_c(\mathcal{C}(\tau(v-1)), \text{round}(\tau(v-1))(\tau(v) - \tau(v-1))$

5. If $\tau(v) \neq \text{round}(\tau(v-1))$, go to Step 6. If $\tau(v) = \text{round}(\tau(v-1))$ ($\tau(v)$ corresponds to the end of a round), then $\mu(P, X_r)$ is defined for all pairs (P, X_r) where $r = \tau(v)$ as follows:[3]

   (a) For all (P, X_{r-1}) such that $\mu(P, X_{r-1}) > 0$, consider $c_1 = Ch_{(P,X_{r-1})}(\mathcal{C}(r-1))$.

      i. If $t^*_{c_1} \geq r$, then $\mu(P, X_r) = \begin{cases} \mu(P, X_{r-1}) & \text{if } X_r = X_{r-1} \cup c_1 \\ 0 & \text{otherwise} \end{cases}$

      ii. If $t^*_{c_1} < r$ (which means that course c_1 runs out during round r so that students with profile (P, X_{r-1}) cannot be guaranteed to get c_1), denote $c_2 = Ch_{(P,X_{r-1})}(\mathcal{C}(t^*_{c_1}))$. If $t^*_{c_2} < r$, reiterate as many times as needed so that we reach a course c_k with $t^*_{c_k} \geq r$.

      $$\mu(P, X_r) = \begin{cases} (t^*_{c_1} - r + 1)\mu(P, X_{r-1}) & \text{if } X_r = X_{r-1} \cup c_1 \\ (t^*_{c_2} - t^*_{c_1})\mu(P, X_{r-1}) & \text{if } X_r = X_{r-1} \cup c_2 \\ \dots & \\ (r - t^*_{c_{k-1}})\mu(P, X_{r-1}) & \text{if } X_r = X_{r-1} \cup c_k \\ 0 & \text{otherwise} \end{cases}$$

   (b) For all other (P, X_r), $\mu(P, X_r) = 0$

6. Increment v by 1 and repeat until $\tau(v) = m$

Because both the number of courses and the number of rounds is finite this process ends in a finite number of stages. The random allocation of students who submitted profile P is given by $\frac{\mu(P,X_m)}{\mu(P,X_0)}$ for X_m. Note that every strategy profile generates deterministic run-out times for courses.

[3]Because $\mathcal{P}$ and $\mathcal{C}$ are both finite, there is a finite number of pairs to consider.

B.2 Formal Description of Random Priority Mechanisms

We describe formally the allocation generated by any multi-unit random priority mechanism (defined in Section 7). As with the formal description of the HBS draft, we fix $\mathbf{P}$ as the ROLs submitted by the students and introduce the notation X_i to denote a subset of $\mathcal{C}$ of cardinality i and the notation $Ch_{(P,X_i)}(\mathcal{C}')$ to describe the course chosen in $\mathcal{C}'$ by a student who submitted preferences P and already has courses in X_i. Finally, let f describe the atomless density function of choosing times that characterizes the multi-unit random priority mechanism at hand. f has support on $\{(t_1, ..., t_m) \in [0, m]^m$ such that $t_1 \leq t_2 \leq ... \leq t_m\}$.[4]

The main difference with the HBS draft mechanism is that we need to track the measure of students who submitted P and already have courses in X_i *at all times*, rather than just at the end of each round. We denote this measure $\mu(P, X_i, t)$. Likewise, the proportion of students requesting course c when courses in $\mathcal{C}'$ are available must be tracked at all times. We denote it by $m_c(\mathcal{C}', t)$.[5] Formally, it is defined as

$$m_c(\mathcal{C}', t) = \sum\nolimits_{\{(P,X_i)|i=1,...,m \text{ and } c=Ch_{(P,X_i)}(\mathcal{C}')\}} \mu(P, X_i, t) \tag{2}$$

Because there is a finite number of (P, X_i), $m_c(\mathcal{C}', t)$ is well defined for all t as long as $\mu(P, X_i, t)$ is well defined. This will be done recursively as part of the allocation process.

Let $\mathcal{C}(t)$ describe the set of courses available at time t, $x_c(t)$ the used capacity of course c at time t and t_c^* the run-out time of course c. Initialize $\mathcal{C}(0) = \mathcal{C}$, $x_c(0) = 0$, and $t_c^* = m$. Let $\mu(P, X_0, 0)$ be equal to the measure of students using strategy P in play $\mathbf{P}$ (and let all other $\mu(P, X, 0) = 0$ for $|X| \neq 0$). We introduce two functions, $\tau_c(.)$ and $\tau(.)$ whose value at the different steps of the algorithm will be defined endogenously. We initialize them at $\tau_c(0) = 0$ and $\tau(0) = 0$. Finally let $v = 1$. Course run-out times, course requests and allocations in any multi-unit random priority mechanisms are jointly determined iteratively through the following process.

1. For each course c, we define the time it would run-out if the proportion of students requesting it is equal to $m_c(\mathcal{C}(\tau(v-1)), s)$:

$$\tau_c(v) = \sup\{t \in [\tau(v-1), m] | x_c(v-1) + \int_{\tau(v-1)}^{t} m_c(\mathcal{C}(\tau(v-1)), s) ds < q_c\}$$

   where $m_c(\mathcal{C}(\tau(v-1)), s)$ is based on (2) with $\mu(P, X_i, t)$ defined as follows. Consider any X_i and P, and label the least preferred course in X_i by c_i, the second least preferred by $c_{i-1}, ...$. Let k denote the number of consecutive courses in X_i, starting from the least preferred c_i

[4] For example, the HBS draft mechanism is characterized by the joint probability distribution with uniform support on $(x, 2-x, 2+x, ...)$ for $x \in [0, 1]$, whereas RSD is characterized by the joint probability distribution with uniform support on $(x, ..., x)$ for $x \in [0, m]$.

[5] As in the description of the HBS draft mechanism, we can leverage Che and Kojima (2010)'s trick of modeling priorities as iid draws to get a weak law of large numbers.

and counting in order of increasing preference, that are in $\mathcal{C}(\tau(v-1))$ (i.e., still available at $\tau(v-1)$). Let $\Pr(i-r, r, \tau(v-1), s)$ denote the probability that a student with preferences P chooses $i-r$ courses before $\tau(v-1)$ and exactly r between $\tau(v-1)$ and s (this depends only on f). For all $s > \tau(v-1)$

$$\begin{aligned} \mu(P, X_i, s) &= \mu(P, X_i, \tau(v-1)) \Pr(i, 0, \tau(v-1), s) + \\ & \sum_{r=1}^{k} \mu(P, X_i \backslash \{c_i, .., c_{i+1-r}\}, \tau(v-1)) \Pr(1-r, r, \tau(v-1), s) \end{aligned} \quad (3)$$

In words, the measure of students who submitted P and already have courses in X_i at time t depends on the measure of students who submitted P and had any possible "antecedent" of X_i at time $\tau(v-1)$, weighted by the probability that they fill their schedule to reach X_i at time t. The set of possible antecedents depends both on P and the courses available at time $\tau(v-1)$ (which are assumed to be available between $\tau(v-1)$ and s). For this reason, $\mu(P, X_i, s)$ needs to be reevaluated at every pass of the algorithm.

2. Let $\tau(v) = \min_{c \in \mathcal{C}(\tau(v-1))} \{\tau_c(v)\}$ which corresponds to the earliest run-out time. Set $t_c^* = \tau_c(v)$ for c such that $\tau(v) = \tau_c(v)$

3. Adjustment of the set of available courses: $\mathcal{C}(\tau(v)) = \mathcal{C}(\tau(v-1)) \backslash \{c \in \mathcal{C}(\tau(v-1)) | \tau_c(v) = \tau(v)\}$.

4. Adjustment of capacities: For all c, $x_c(v) = x_c(v-1) + \int_{\tau(v-1)}^{\tau(v)} m_c(\mathcal{C}(\tau(v-1)), s) ds$

5. Increment v by 1 and repeat until $\tau(v) = m$

Notes:

1. Because there is a finite number of courses, the algorithm reaches an end in a finite number of steps.

2. A direct output of the formal description is that course run-out times are deterministic in multi-unit random priority mechanisms.

3. It is easy to check that, because (a.s.) all students get their n^{th} course before any other student gets to choose his $n+1^{th}$ course, non-callous mechanisms take place over m distinct rounds. As a result, the formalism in Appendix B equally applies to describe the allocation generated by non-callous mechanisms.

C Preference Data and Robustness Checks

The key inputs to our welfare analysis in Sections 5 and 6 are students' truthful and strategically reported preferences. In this appendix we describe in more detail how we construct truthful preferences from our data, and the various robustness checks we conduct.

C.1 Truthful Preferences: Main Specification

Based on the evidence presented in Section 4, we start by assuming that students' truthful top-five courses correspond to their top-five courses in the May poll. We then adjust these preferences to account for the cases of likely preference change as identified in Section 4.2. Specifically, for the cases where a student violated Theorem 2 either by dropping a course (141 cases) or by downgrading a popular course for which there was a negative aggregate preference shock (42 cases) we assume that the student's preference for that course changed and act as if the student did not rank it in May. Courses that were offered in the May poll but no longer available in the July run are also dropped.

We then append all of the other courses that the student ranks in their July submitted ROL in order of their relative ranking in July. This convention will cause us to underestimate the extent of strategic behavior in the HBS draft, because we assume that the relative ranking of courses not in the top five is truthful. To illustrate, suppose a student submitted the ROL c_1, c_2, c_3, c_4, c_5 in the May poll and submitted $c_4, c_3, c_6, c_1, c_2, c_5, c_7, c_8$ in the July run. We construct his truthful preferences as $c_1, c_2, c_3, c_4, c_5, c_6, c_7, c_8$. If in July he did not rank say c_4, in violation of Theorem 2, his constructed truthful preferences would instead be $c_1, c_2, c_3, c_5, c_6, c_7, c_8$.

C.2 Robustness Checks

The results of our welfare analysis in Sections 5 and 6 are robust to a variety of alternate specifications for truthful and strategic preferences.

Our first robustness check uses students' May trial-run preferences instead of their July actual-run preferences, both for students' strategic reports and for the construction of their truthful preferences beyond the top five from the May poll data. The advantage of using the May trial-run preferences is that just 10 days elapsed between the May poll and the May trial-run, so there is less reason to worry about social learning, new information and idiosyncratic preference changes. The disadvantage is that the July run is what actually mattered for welfare, and students may have used the May trial-run to learn about equilibrium best responses.

Our next two robustness checks consider alternative ways of handling violations of Theorem 2 when constructing truthful preferences.[6] In the first, we attribute all Theorem 2 violations to

[6] As described in the text of Section 4.2, ideally welfare analysis should correct for Theorem 2 violations that are

preference change. In the second, we attribute all Theorem 2 violations to strategic error. The first specification is likely to make the HBS draft look a bit better than it actually is because it excludes the possibility of human error; the latter is likely to make the HBS draft look a bit worse than it actually is because it exaggerates the frequency of human error.

Reassuringly, the results of Sections 5 and 6 do not move that much under any of these alternate specifications. Furthermore, they move in the direction we expect given the issues of concern. For instance, treating all Theorem 2 violations as due to preference change narrows the difference in mean average ranks between truthful and strategic play of the HBS draft, whereas treating all violations as due to strategic error widens the gap.

Our last robustness check investigates the mechanical effect of overreporting of popular courses in a simulated economy that resembles the HBS economy. Specifically, there are $1,000$ students each of whom requires 10 courses, and 100 courses with 120 seats each, for 20% excess capacity as in the data. Students have additive risk-neutral preferences. Student s's value for course c is given by

$$v_{sc} = v_c + \epsilon_{sc}$$

where $v_c \sim U[0,1]$ is a common-value quality component for course c, and $\epsilon_{sc} \sim U[0,1]$ represents s's idiosyncratic taste for c. We then assume that students report their preferences under the HBS draft as if their preferences are

$$\widehat{v}_{sc} = 2v_c + \epsilon_{sc}$$

that is, students over-weight the common-value quality component. Our aim is not to model equilibrium behavior but rather to understand, in a simple and transparent way, the mechanical effects of strategic overreporting of popular courses on the welfare measures we care about.

The main patterns we find in Sections 5 and 6 emerge in this simple simulation as well. Thus, a reader who is skeptical of survey data,[7] but who is persuaded that students are likely to strategically overreport popular classes, should be somewhat willing to believe our basic results.

Table C1 reports the most salient moments of our analysis under each of the above-described specifications. Table C2 reports the social comparison results (CR5-7) for each of the specifications. We then re-run the individual-level comparison results (CR1-4), i.e., Tables 5 and 6 from the main text, under each of the specifications.

due to preference change between May and July, but should not correct for Theorem 2 violations that are due to strategic error.

[7]Some of the advantages and disadvantages of using survey data for economic analysis are described in Bertrand and Mullainathan (2001). Fortunately, our context avoids some of the most important disadvantages (as compared e.g. to surveys of political attitudes).

Table C1: Robustness Checks: Summary Statistics

	E(Average Rank)	Pr(Get 1st Favorite)	Pr(Get Top Ten)
Main Specification			
HBS - Truthful Play	7.66	82.2%	1.5%
HBS - Strategic Play	7.99	63.3	2.2
RSD - Truthful Play	8.74	49.0	29.7
May Trial Run instead of July Actual Run			
HBS - Truthful Play	7.55	83.7%	1.0%
HBS - Strategic Play	8.00	60.8	1.9
RSD - Truthful Play	8.59	48.2	31.5
Theorem 2 Violations all Preference Change			
HBS - Truthful Play	7.67	81.8%	1.4%
HBS - Strategic Play	7.86	65.0	2.4
RSD - Truthful Play	8.69	49.3	29.7
Theorem 2 Violations all Strategic Mistakes			
HBS - Truthful Play	7.76	82.7%	0.9%
HBS - Strategic Play	8.25	59.7	1.9
RSD - Truthful Play	9.95	47.3	31.0
Simulation Economy			
HBS - Truthful Play	15.34	99.1%	0.0%
HBS - Strategic Play	16.38	86.5	0.0
RSD - Truthful Play	19.98	29.8	12.0

Table C2: Robustness Checks: Social Comparison Results

	CR5	CR6(i)	CR6(ii)	CR6(iii)	CR7
HBS Truthful vs. HBS Strategic					
Main Specification	HBS-T	Indet.	HBS-T	HBS-T	HBS-T
May Trial Run instead of July Actual Run	HBS-T	Indet.	HBS-T	HBS-T	HBS-T
Thm 2 Violations all Preference Change	HBS-T	Indet.	HBS-T	HBS-T	HBS-T
Thm 2 Violations all Strategic Mistakes	HBS-T	Indet.	HBS-T	HBS-T	HBS-T
Simulation Economy	HBS-T	HBS-T	HBS-T	HBS-T	HBS-T
HBS Strategic vs. RSD					
Main Specification	HBS-S	Indet.	HBS-S	HBS-S	HBS-S
May Trial Run instead of July Actual Run	HBS-S	Indet.	HBS-S	HBS-S	HBS-S
Thm 2 Violations all Preference Change	HBS-S	Indet.	HBS-S	HBS-S	HBS-S
Thm 2 Violations all Strategic Mistakes	HBS-S	Indet.	HBS-S	HBS-S	HBS-S
Simulation Economy	HBS-S	Indet.	HBS-S	HBS-S	HBS-S

Table 5 Robustness Check - May Trial Run instead of July Actual Run

	Assumption on Preferences					
	Responsive	Additive	Average-Rank			Lexicographic
	Any Risk Attitude	Risk Neutral	Any Risk Attitude	Risk Averse	Risk Neutral	Risk Neutral
Outcome	CR1	CR2	CR3(i)	CR3(ii)	CR3(iii)	CR4
Prefers HBS Truthful	41%	44%	49%	60%	66%	87%
Prefers HBS Strategic	8%	8%	14%	20%	32%	12%
Indifferent	1%	1%	1%	1%	1%	1%
Indeterminate	49%	46%	36%	18%	0%	0%

Table 6 Robustness Check - May Trial Run instead of July Actual Run

	Assumption on Preferences					
	Responsive	Additive	Average-Rank			Lexicographic
	Any Risk Attitude	Risk Neutral	Any Risk Attitude	Risk Averse	Risk Neutral	Risk Neutral
Outcome	CR1	CR2	CR3(i)	CR3(ii)	CR3(iii)	CR4
Prefers RSD	0%	0%	0%	0%	18%	22%
Prefers HBS Strategic	2%	30%	2%	81%	81%	77%
Indifferent	0%	0%	0%	0%	0%	0%
Indeterminate	98%	70%	97%	18%	0%	0%

Table 5 Robustness Check - Theorem 2 Violations all Preference Change

	Assumption on Preferences					
	Responsive	Additive	Average-Rank			Lexicographic
	Any Risk Attitude	Risk Neutral	Any Risk Attitude	Risk Averse	Risk Neutral	Risk Neutral
Outcome	CR1	CR2	CR3(i)	CR3(ii)	CR3(iii)	CR4
Prefers HBS Truthful	38%	40%	45%	57%	61%	80%
Prefers HBS Strategic	13%	13%	19%	23%	37%	18%
Indifferent	2%	2%	2%	2%	2%	2%
Indeterminate	47%	45%	34%	19%	0%	0%

Table 6 Robustness Check - Theorem 2 Violations all Preference Change

	Assumption on Preferences					
	Responsive	Additive	Average-Rank			Lexicographic
	Any Risk Attitude	Risk Neutral	Any Risk Attitude	Risk Averse	Risk Neutral	Risk Neutral
Outcome	CR1	CR2	CR3(i)	CR3(ii)	CR3(iii)	CR4
Prefers RSD	0%	0%	0%	0%	7%	17%
Prefers HBS Strategic	2%	36%	3%	92%	93%	83%
Indifferent	0%	0%	0%	0%	0%	0%
Indeterminate	98%	64%	97%	8%	0%	0%

Table 5 Robustness Check - Theorem 2 Violations all Strategic Mistakes

	Assumption on Preferences					
	Responsive	Additive	Average-Rank			Lexicographic
	Any Risk Attitude	Risk Neutral	Any Risk Attitude	Risk Averse	Risk Neutral	Risk Neutral
Outcome	CR1	CR2	CR3(i)	CR3(ii)	CR3(iii)	CR4
Prefers HBS Truthful	41%	43%	51%	64%	68%	89%
Prefers HBS Strategic	7%	6%	14%	18%	31%	10%
Indifferent	1%	1%	1%	1%	1%	1%
Indeterminate	51%	50%	34%	17%	0%	0%

Table 6 Robustness Check - Theorem 2 Violations all Strategic Mistakes

	Assumption on Preferences					
	Responsive	Additive	Average-Rank			Lexicographic
	Any Risk Attitude	Risk Neutral	Any Risk Attitude	Risk Averse	Risk Neutral	Risk Neutral
Outcome	CR1	CR2	CR3(i)	CR3(ii)	CR3(iii)	CR4
Prefers RSD	0%	0%	0%	0%	16%	24%
Prefers HBS Strategic	2%	29%	2%	84%	84%	76%
Indifferent	0%	0%	0%	0%	0%	0%
Indeterminate	98%	71%	98%	16%	0%	0%

Table 5 Robustness Check - Simulation Economy

	Assumption on Preferences					
	Responsive	Additive	Average-Rank			Lexicographic
	Any Risk Attitude	Risk Neutral	Any Risk Attitude	Risk Averse	Risk Neutral	Risk Neutral
Outcome	CR1	CR2	CR3(i)	CR3(ii)	CR3(iii)	CR4
Prefers HBS Truthful	26%	30%	68%	85%	95%	96%
Prefers HBS Strategic	0%	0%	1%	2%	6%	4%
Indifferent	0%	0%	0%	0%	0%	0%
Indeterminate	74%	70%	31%	13%	0%	0%

Table 6 Robustness Check - Simulation Economy

	Assumption on Preferences					
	Responsive	Additive	Average-Rank			Lexicographic
	Any Risk Attitude	Risk Neutral	Any Risk Attitude	Risk Averse	Risk Neutral	Risk Neutral
Outcome	CR1	CR2	CR3(i)	CR3(ii)	CR3(iii)	CR4
Prefers RSD	0%	0%	0%	0%	1%	2%
Prefers HBS Strategic	0%	42%	0%	10%	99%	98%
Indifferent	0%	0%	0%	0%	0%	0%
Indeterminate	100%	58%	100%	90%	0%	0%

D Empirical Analysis of the Proxy Draft

In this Appendix, we describe our computational implementation of the Proxy Draft and report the details of the empirical analysis omitted from Section 8.

D.1 Computational Implementation of the Proxy Draft

We implement the proxy draft computationally by finding a fixed point in course run-out times. Specifically, given preferences $\mathbf{P}$ and a realized priority order λ, we seek to identify a set of run-out times for courses $t^* = \{t^*_c\}_{c \in \mathcal{C}}$ such that, when the proxy optimizes with respect to these run-out times on each student s's behalf (which is possible per Lemma 2) these run-out times are in fact correct.

We use the following best-response dynamic to find a fixed point. We begin by running the standard HBS draft mechanism for priority order λ and truthful preferences $\mathbf{P}$. This yields a set of run-out times, which we denote t^0. We then have the proxy optimize on each student's behalf with respect to t^0;[8] call the set of best responses $\mathbf{P}^1$. Now, run the HBS draft mechanism again under preferences $\mathbf{P}^1$ to obtain a new set of run-out times, t^1, and compute a new set of best responses to t^1; call them $\mathbf{P}^2$. We continue in this fashion in the hopes of finding some $t^k = t^{k-1}$. If this is the case, the strategies $\mathbf{P}^k$ are optimal with respect to t^{k-1} and generate run-out times t^k, so since $t^k = t^{k-1}$ we have found our fixed point. The outcome under the proxy draft is whatever occurs under the original HBS draft when the priority order is λ and students play $\mathbf{P}^k$.

This procedure sometimes fails to converge. If we go too many iterations without convergence we perturb the observed run-out times by a random amount and then continue. In our data we always find a fixed point, typically on the first few tries.

D.2 Numerical Simulations of the Performance of the Proxy Draft

We randomly choose 100 random priority orders, and for each order run the proxy draft mechanism, the HBS draft mechanism under truthful play, and the HBS draft mechanism under strategic play. We then use the methodology of Sections 5.3-5.4 to compare truthful play of the proxy draft to both truthful and strategic play of the HBS draft.

Tables D1 and D2 report ex-ante welfare comparisons at the individual level.

[8] Recall from the proof of Lemma 2 that there can be multiple best responses that are outcome equivalent. We choose the best response that lexicographically minimizes the amount by which most-preferred courses are downgraded. For instance if s's true preferences are $P_s : c_1, c_2, c_3$ and two equally good responses are $P'_s : c_1, c_3, c_2$ and $P''_s : c_3, c_2, c_1$, we choose P'_s because it downgrades the most-preferred course c_1 by less. This convention seems to facilitate convergence.

Table D1. Individual preferences between Proxy Draft and HBS-Truthful

	Assumption on Preferences					
	Responsive	Additive	Average-Rank			Lexicographic
	Any Risk Attitude	Risk Neutral	Any Risk Attitude	Risk Averse	Risk Neutral	Risk Neutral
Outcome	(1)	(2)	(3)	(4)	(5)	(6)
Prefers HBS Truthful	29%	32%	41%	52%	56%	62%
Prefers Proxy Draft	23%	23%	30%	35%	43%	38%
Indifferent	0%	0%	0%	0%	1%	0%
Indeterminate	48%	45%	28%	13%	0%	0%

Table D2. Individual preferences between Proxy Draft and HBS-Strategic

	Assumption on Preferences					
	Responsive	Additive	Average-Rank			Lexicographic
	Any Risk Attitude	Risk Neutral	Any Risk Attitude	Risk Averse	Risk Neutral	Risk Neutral
Outcome	(1)	(2)	(3)	(4)	(5)	(6)
Prefers HBS Strategic	4%	5%	13%	27%	39%	37%
Prefers Proxy Draft	21%	21%	33%	52%	60%	63%
Indifferent	1%	1%	1%	1%	1%	1%
Indeterminate	74%	73%	53%	20%	0%	0%

For each comparison result, a plurality of students prefers truthful play of the HBS draft mechanism to equilibrium play of the proxy draft, and similarly a plurality prefers equilibrium play of the proxy draft to the actual play of the HBS draft mechanism. This is the "lands in between" result at the individual level described in the main text.

Figures D1 and D2 report ex-ante welfare comparisons at the societal level.

A utilitarian social planner prefers truthful play of the HBS draft to equilibrium play of the proxy draft, and prefers equilibrium play of the proxy draft to the actual play of the HBS draft mechanism, for each of CR5, CR6(ii), CR6(iii), and CR7. This is the "lands in between" result at the societal level described in the main text.

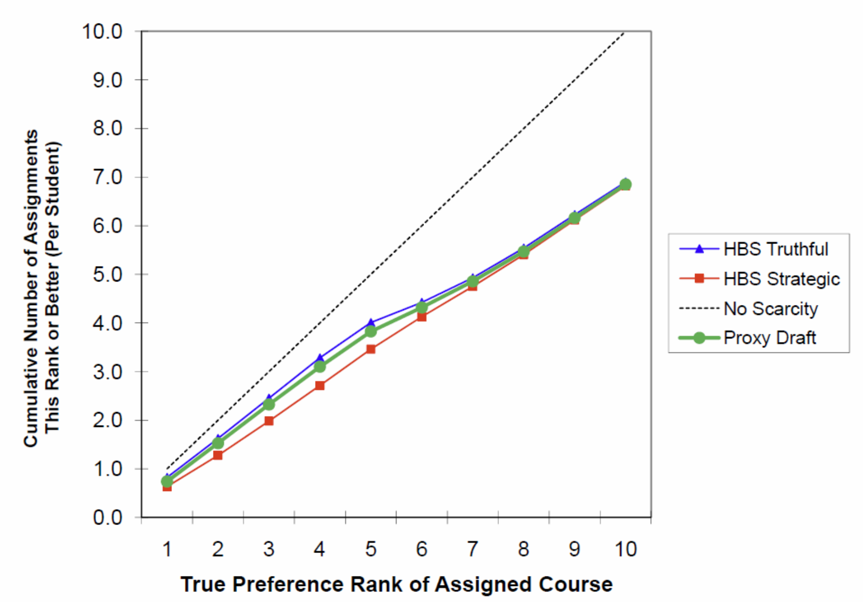

Figure 1: Cumulative distribution of the true preference rank of students' assigned courses: proxy draft versus truthful and strategic play of the HBS draft. The distribution under truthful play of the HBS draft first-order stochastically dominates the distribution under the proxy draft, which itself first-order stochastically dominates the distribution under strategic play of the HBS draft. Thus in each case, CR5, CR6(iii) and CR7 obtain.

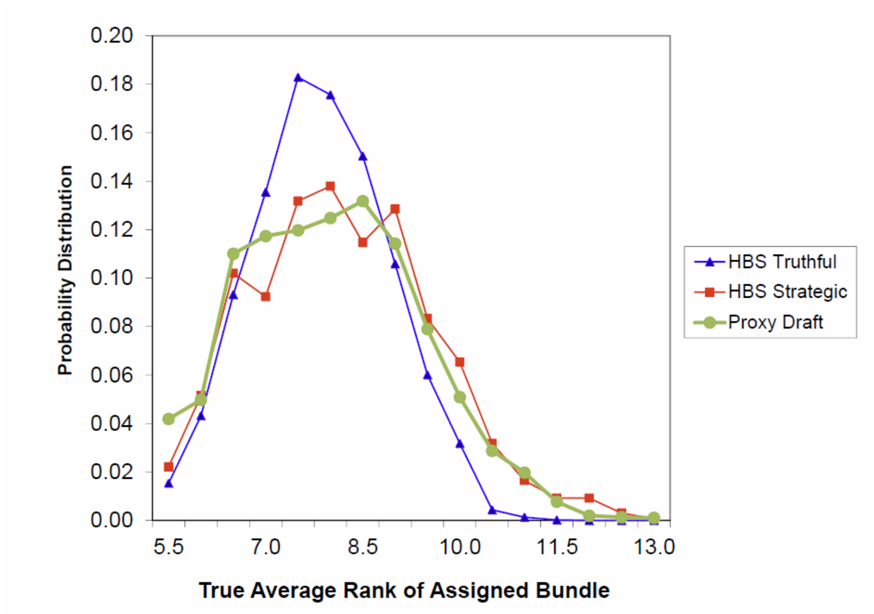

Figure 2: Probability distribution of the true-preference average rank of students' assigned bundles: proxy draft versus truthful and strategic play of the HBS draft. The distribution under truthful play of the HBS draft second-order stochastically dominates that under the proxy draft, which itself first-order stochastically dominates that under strategic play of the HBS draft. So CR6(ii)-(iii) and CR6(i)-(iii) obtain, respectively.

References

[1] Bertrand, Marianne and Sendhil Mullainathan (2001). "Do People Mean What They Say? Implications for Subjective Survey Data." American Economic Review, 116(3), 901-932.

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

[3] Uhlig, Harald (1996). "A Law of Large Numbers for Large Economies." Economic Theory, 8, 41-50.