Supplementary Appendix to "Strategy-proofness in the Large" (for Online Publication)

Eduardo M. Azevedo* and Eric Budish†

July 16, 2018

B Clarifying Example: Multi-Unit Auctions

This section illustrates several of the key definitions in the paper using the example of multi-unit auctions for identical objects, such as government bond auctions. We will see that the uniform-price auction is SP-L whereas the pay-as-bid auction is manipulable in the large. The example also illustrates the large-market limit, the role of the full-support requirement, and the contrast between SP-L and the traditional notion of approximate strategy-proofness based on ex-post realizations of others' play.

Example B.1. (**Multi-Unit Auctions**). There are kn units of a homogeneous good. To simplify notation, we assume that agents assign a constant per-unit value to the good, up to a capacity limit. Specifically, each agent i's type $t_i = (v_i, q_i)$ consists of a per-unit value v_i and a maximum capacity q_i. The set of possible values is $V = \{1, \ldots, \bar{v}\}$, the set of possible capacity limits is $Q = \{0, 1, \ldots, \bar{q}\}$ with $1 < k < \bar{q}$, and $T = V \times Q$. The set of outcomes is $X_0 = (\{1, 2, \cdots, \bar{v}\} \times \{1, 2, \cdots, \bar{q}\}) \cup \{0\}$, with an outcome consisting either of a per-unit payment and an allotted quantity, or 0 to denote that the agent receives no units and makes no payment.

We first describe the uniform-price auction. Bids consist of a per-unit value and a maximum capacity, so the action set $A = T$. Given a vector of n bidders' reports t, denote the demand for the object at a price of p as $D(p; t) = \sum_{i=1}^{n} q_i \cdot 1\{v_i \geq p\}$, where $1\{\cdot\}$ is the indicator function. The market-clearing price $p^*(t)$ is the highest price at which demand

*Wharton, eazevedo@wharton.upenn.edu.
†University of Chicago Booth School of Business, eric.budish@chicagobooth.edu.

exceeds supply. That is,

$$p^*(t) = \max\left\{p \in V : \frac{D(p;t)}{n} \geq k\right\} \tag{B.1}$$

if $D(1,t) \geq k$ and $p^*(t) = 0$ otherwise. The uniform-price auction allocates each bidder i her demanded quantity at $p^*(t)$, with the exception that bids with $v_i = p^*(t)$ are rationed with equal probability. Formally, $\Phi_i^n(t)$ allots each bidder the following number of units of the good,

Reported Value	Expected Number of Units
$v_i < p^*(t)$	0
$v_i = p^*(t)$	$\bar{r} \cdot q_i$
$v_i > p^*(t)$	q_i

at a price per unit of $p^*(t)$, and the rationing probability $\bar{r}$ set so that the market clears.[1]

We now analyze the large-market limit of the uniform-price auction. Let $\rho^*(m)$ denote the price that clears supply and *average demand* given bid distribution m. That is,

$$\rho^*(m) = \max\{p \in V : \mathbb{E}[D(p;t_i)|t_i \sim m] \geq k\} \tag{B.2}$$

if $\mathbb{E}[D(1;t_i)|t_i \sim m] \geq k$ and 0 otherwise.

Generically, expected demand at price $\rho^*(m)$ strictly exceeds supply, that is,

$$\mathbb{E}[D(\rho^*(m);t_i)|t_i \sim m] > k.$$

In this generic case, as the market grows large, the realized price as defined in (B.1) will be equal to $\rho^*(m)$ with probability converging to one. Therefore, the limit mechanism allocates each bidder their demand at $\rho^*(m)$, again with the exception that bidders with value exactly equal to $\rho^*(m)$ are rationed, and with all winning bidders paying $\rho^*(m)$ per unit. Formally,

[1]Since preferences are linear up to the capacity limit, the exact form of the rationing is immaterial in the analysis below. The rationing constant is

$$\bar{r} = \frac{kn - D(p^*(t)+1;t)}{D(p^*(t);t) - D(p^*(t)+1;t)}.$$

$\phi^\infty(t_i, m)$ gives player i

Reported Value	Expected Number of Units
$v_i < \rho^*(m)$	0
$v_i = \rho^*(m)$	$\bar{r} \cdot q_i$
$v_i > \rho^*(m)$	q_i

at a per unit price of $\rho^*(m)$, and the rationing probability $\bar{r}$ is set so that the market clears on average.[2] Note that, in this generic case, the price in the limit is deterministic and is exogenous from the perspective of each individual bidder.

In addition to the generic case, there is a knife-edge case in which expected demand at $\rho^*(m)$ is exactly equal to supply. That is, $\mathbb{E}[D(\rho^*(m); t_i)|t_i \sim m] = k$ and $\rho^*(m) > 0$. In this case, focusing for now on m with full support, the price is stochastic even in the large-market limit. Given large n, the realized per-capita demand at price $\rho^*(m)$ will be weakly greater than per-capita supply k with probability of about $\frac{1}{2}$, and will be strictly smaller than per-capita supply k with probability of about $\frac{1}{2}$.[3] Therefore, the price in the limit will be $\rho^*(m)$ with probability of $\frac{1}{2}$, and $\rho^*(m) - 1$ with probability of $\frac{1}{2}$. $\phi^\infty(t_i, m)$ assigns to player i the following expected number of units,

Reported Value	Expected Number of Units
$v_i < \rho^*(m)$	0
$v_i \geq \rho^*(m)$	q_i

and prices are $\rho^*(m)$ or $\rho^*(m) - 1$ with equal probability. Note that bids of $\rho^*(m)$ are not rationed in the limit. This is so because, in this knife-edge case, average demand is exactly equal to average supply. Moreover, in both cases the price in the limit is exogenous from the perspective of each individual bidder. Even though the price is sometimes $\rho^*(m)$ and sometimes $\rho^*(m) - 1$, the probability that bidder i is pivotal in determining which of the two prices occurs converges to zero.

The argument that the uniform-price auction is SP-L is now straightforward. Choose

[2] That is, $\bar{r}$ satisfies

$$\bar{r} = \frac{k - E[D(\rho^*(m)+1; t_i')|t_i' \sim m]}{E[D(\rho^*(m); t_i')|t_i' \sim m] - E[D(\rho^*(m)+1; t_i')|t_i' \sim m]}.$$

[3] The intuition is that if a fair coin is tossed $n \to \infty$ times, the probability that at least $n/2$ of the tosses are heads converges to $1/2$, just as the probability that less than $n/2$ of the tosses are heads converges to $1/2$, with both probabilities independent of the outcome of the i^{th} toss.

any type t_i and any full support distribution $m \in \bar{\Delta}T$. The description of ϕ^∞ above implies that truthful reporting is a dominant strategy in the limit, hence Definition 4 is satisfied.

Last, we turn to the pay-as-bid auction. The pay-as-bid auction allocates units of the good in exactly the same way as the uniform-price auction. The difference is that winning bidders pay their bid instead of the market-clearing price $p^*(t)$. Clearly, bidders will gain from misreporting their value, even in the large-market limit. If the distribution of opponent bids is m and the limit price is $\rho^*(m)$, then a bidder of type $t_i = (v_i, q_i)$ with $v_i > \rho^*(m) + 1$ strictly prefers to misreport as $t_i' = (\rho^*(m) + 1, q_i)$: he receives the same allocation in the limit but pays a strictly lower price per unit. Hence, the pay-as-bid auction is not SP-L.□

Discussion: SP-L vs. Traditional Approximate SP Observe that the argument that the uniform price auction is SP-L would not go through using a stronger notion of asymptotic strategy-proofness based on realizations of opponents' reports rather than probability distributions. To see this, consider the case where there are $k = 2$ objects per bidder, and bidder i knows that all other bidders will report a demand of 2 objects for \$100. That is, that all other bidders report a type of $(2, \$100)$ for sure. Then bidder i knows that she is marginal, and can reduce the market-clearing price to 0 by asking for 1 object instead of 2. This example illustrates the importance of the interim perspective in the definition of SP-L, and why SP-L classifies mechanisms in a substantially different way than the traditional ex-post notion of approximate strategy-proofness.

Discussion: Full Support Requirement The uniform-price auction example also illustrates the importance of the full-support requirement in the SP-L definition. If a bidder believes that opponent reports equal $(2, \$100)$ for sure, then she could lower the market-clearing price from \$100 to \$0 by demanding a single unit. While this example uses a degenerate distribution with support on a single type, there are non-degenerate distributions where a bidder can manipulate the uniform-price auction.[4] For example, if bidder i believes that opponents report $(2, \$100)$ or $(2, \$200)$ with equal probability, then she can still, with high probability, drive the prices down from \$100 to \$0 by asking for one unit instead of two. Even though she is uncertain about other players' types, this uncertainty is at a part of the demand curve that is not relevant for the determination of the market-clearing price. That said, such manipulations do not seem very realistic because they require extremely detailed information about opponent play. The full support requirement in the SP-L definition is a

[4] See also Swinkels (2001; Section 5) for an elegant example, with limited support, in which bidders remain pivotal with probability one even in very large markets.

simple way to capture the idea that agents are not likely to have that level of information.

The uniform-price auction example also suggests that, on a case by case basis, it may not always be necessary to assume full support. As long as there is uncertainty about opponents' play in a region that is relevant for price determination, bidding truthfully will be optimal in a large enough market. For example, assume that bidder i believes that her opponents report $(4, \$100)$ or $(4, \$25)$ with 50% chance each. Then, in the limit, the market-clearing price is \$100 or \$25 with 50% chance each, and bidder i cannot meaningfully affect the price. Thus, even in this adversarial case where supply intersects the expected demand curve at a discontinuity, and bidder i thinks that the distribution of opponents' reports has only two elements in its support, reporting truthfully is approximately optimal from the interim perspective.

C Semi-Anonymity

Our main analysis considers anonymous mechanisms, where agents' outcomes depend on their own report and the distribution of all reports. The analysis generalizes straightforwardly, though at some notational burden, to the case of semi-anonymous mechanisms, as defined by Kalai (2004). In this setting, agents are divided into a number of groups, and agents within each group can be treated differently by the mechanism.

In this section, agents belong to **groups** g in a finite set G. The set of types is partitioned into subsets

$$T = T_{g_1} \cup T_{g_2} \cup \cdots \cup T_{g_G}.$$

A **semi-anonymous mechanism** is defined as $\{(\Phi^n)_{n\in\mathbb{N}}, (A_g)_{g\in G}\}$, where the A_g are the sets of actions available to each group g, and

$$A = A_{g_1} \cup \cdots \cup A_{g_G}$$

is the set of actions. As in the anonymous case, the Φ^n are functions

$$\Phi^n : A^n \to \Delta({X_0}^n).$$

The difference with respect to anonymous mechanisms is that agents in group g are restricted to play strategies in A_g. That is, if $t_i \in T_g$ then the support of any strategy $\sigma(t_i)$ is contained in A_g. In a matching setting, for example, the groups may specify whether an agent is a man or a woman, and the agent's traits. Agents are then permitted to misreport their preferences over other match partners, but they cannot misrepresent their gender or their traits. **Limit mechanisms** are defined as in Section 3.1. In particular, we define limit mechanisms with respect to a single distribution $\mu \in \Delta T$, and not distributions of types within groups. Alternatively, one could assume that the number of agents in each group grows in a specific way, and that types are drawn i.i.d. within each group. We now formally define a two-sided matching mechanism, to clarify the definition.

Example C.1. (Two-Sided Matching) This example shows that semi-anonymous mechanisms include matching mechanisms in two-sided markets (Gale and Shapley, 1962). Agents are men and women, who differ on a set of traits. Groups g index both sex and the traits, so that the set of groups is

$$G = \{m_1, m_2, \cdots, m_M\} \cup \{w_1, w_2, \cdots, w_W\}.$$

That is, there are M groups of men and W groups of women. Men and women within each group have the same traits, and are equally good marriage partners. However, within each group, agents may differ in their preferences over the other groups. The way in which the semi-anonymous framework differs from the anonymous setting is that men and women may misreport their preferences, but cannot misreport their sex nor traits.

Formally, agent i's type is

$$t_i = (g_{t_i}, u_{t_i}),$$

where $g_{t_i} \in G$ is the agent's group, and u_{t_i} is a strictly positive utility function over the groups of the opposite sex. The set of outcomes $X_0 = G \cup \emptyset$. That is, each agent only cares about which type of man (woman) she (he) is matched to, or whether she (he) is unmatched. Utilities of each type t_i are given by $u_{t_i}(g)$ if she is matched to someone of the opposite sex. We extend u_{t_i} so that it is 0 if the agent is unmatched or matched to a group of the same sex.

Consider now a stable matching mechanism, using a tie-breaking lottery, as in school choice mechanisms (Abdulkadiroğlu et al., 2009). The mechanism is direct, so that $A_g = T_g$ for each $g \in G$. Men and women report a vector of types t, and therefore traits. This implies a weak preference ordering of each man over each woman and vice versa. The mechanism assigns a lottery number l_i to each agent, uniformly and independently distributed between 0 and 1. Lottery numbers are used to break ties between preferences. That is, preferences are refined to strict preferences, by using the lottery numbers to break ties. Conditional on a vector of lotteries l and a vector of reported types t, the mechanism implements a stable matching $x^n(t,l)$. The function $x^n(t,l)$ is taken to be symmetric, to conform to the semi-anonymity assumption. The mechanism is then defined as

$$\Phi^n(t) = \int_{l\in[0,1]^n} x^n(t,l)\, dl.$$

□

Define a semi-anonymous mechanism as SP-L if no agent wants to misreport as a different type within the same group.

Definition C.1. *The direct, semi-anonymous, mechanism* $\{(\Phi^n)_{\mathbb{N}}, T\}$ *is* ***strategy-proof in the large (SP-L)*** *if, for any* $m \in \bar{\Delta}T$ *and* $\epsilon > 0$ *there exists* n_0 *such that, for all* $n \geq n_0$, $g \in G$, *and all* $t_i, t_i' \in T_g$,

$$u_{t_i}[\phi^n(t_i, m)] \geq u_{t_i}[\phi^n(t_i', m)] - \epsilon.$$

If the mechanism has a limit, this is equivalent to, for any $m \in \bar{\Delta}T$, $g \in G$, *and all* $t_i, t_i' \in T_g$,

$$u_{t_i}[\phi^\infty(t_i, m)] \geq u_{t_i}[\phi^\infty(t_i', m)].$$

Otherwise, the mechanism is ***manipulable in the large.***

The sufficient conditions for a mechanism to be SP-L also have straightforward extensions. The extension of the EF-TB condition is that no agent envies another agent in the same group, and with lower lottery number.

Definition C.2. *A direct semi-anonymous mechanism* $\{(\Phi^n)_{\mathbb{N}}, (T_g)_{g\in G}\}$ *is* ***envy-free but for tie-breaking*** *(EF-TB) if for each* n *there exists a function* $x^n : (T \times [0,1])^N \to \Delta(X_0^n)$, *symmetric over its coordinates, such that*

$$\Phi^n(t) = \int_{l\in[0,1]^n} x^n(t,l) dl$$

and, for all i, j, n, t, *and* l, *if* $l_i \geq l_j$, *and if* t_i *and* t_j *belong to the same group, then*

$$u_{t_i}[x_i^n(t,l)] \geq u_{t_i}[x_j^n(t,l)].$$

With this definition, an extension of Theorem 1 to semi-anonymous mechanisms follows from essentially the same proof.[5] This implies that the stable matching procedure in example C.1 is SP-L, because an agent envying another agent with a lower lottery number would violate the stability condition. In the working paper version of this article we extend a version of Theorem 2 to the semi-anonymous case.

D Details for Table 1

This Section provides supporting details for the classification of non-SP mechanisms presented as Table 1. For each mechanism we provide a formal definition of the mechanism in our setting, a formal proof of the classification, and relevant references.

D.1 Anonymous Mechanisms.

D.1.1 Multi-Unit Auctions

See Appendix B.

D.1.2 Single-Unit Assignment

In single-unit assignment problems, each agent is to be assigned at most one indivisible object, and there are no transfers. We refer the reader to Kojima and Manea (2010) and references therein for a detailed description of the environment and applications.

Formally, we define single-unit assignment as follows. Denote the set of object types by X_0. In a market of size n there are $\{q_{x_0} \cdot n\}$ units of object type x_0 available.[6] An agent of type $t_i \in T$ has a strict utility function u_{t_i} over X_0. It is assumed that X_0 includes a null object $\emptyset$, in supply $n - \sum_{x_0' \neq \emptyset} \{q_{x_0'} \cdot n\} \geq 0$, so that the total quantity of objects equals n. The utility of the null object is normalized to 0. Therefore, we assume that all agents strictly prefer any other object (termed a proper object) to the null object.

Boston Mechanism and Adaptive Boston Mechanism

The Boston mechanism is a mechanism used in many cities to allocate seats in public schools. Abdulkadiroğlu and Sönmez (2003) show that the Boston mechanism is not SP, and

[5] Lemma A.1 holds as is, since it is a statement about the empirical distribution of randomly drawn vectors of types, and therefore does not rely on the definition of a mechanism. Lemma A.2 holds for any two types t_i and t_i' in the same group, using the same proof, as for any such pairs of types the EF-TB condition in the semi-anonymous case implies the same properties as in the anonymous case. Given the two lemmas, the argument in the proof of Theorem 1 in Appendix A.1 holds as is, as long as we take t_i' to be in the same group as t_i, which is all that is needed for the definition of SP-L for semi-anonymous mechanisms.

[6] A bracketed expression denotes the nearest integer to the real number within brackets.

Abdulkadiroğlu et al. (2006) document that it was extensively manipulated in practice. We now formally define the Boston mechanism and show that it is not SP-L. This complements an example given by Kojima and Pathak (2009), in a formally different environment, where the Boston mechanism can be manipulated in a large market. Here we consider the standard version of the Boston mechanism, as opposed to the simplified version used in our application in Section E.

We now define the Boston mechanism. Fix a vector of reports t. To be consistent with the literature we will use the terminology of schools (the objects) and students (the agents). The mechanism first assigns to each student a lottery number l_i, uniformly and independently distributed in $[0, 1]$. The mechanism then proceeds in rounds, following the algorithm below.

1. The mechanism begins in $\texttt{round} = 1$. All students are initially unassigned.
2. Students that are still present in the mechanism take turns, in the order of their lottery number, with higher lottery numbers going first. In her turn, student i is permanently assigned to her $\texttt{round}^{\text{th}}$ choice, as given by u_{t_i}, if there are still seats in that school, or remains unassigned otherwise.
3. If all students have been assigned, finish, otherwise increase $\texttt{round}$ by 1 and go to Step 2.

Note that the algorithm must finish, as eventually all students are assigned either to a proper school or to the null school $x_0 = \emptyset$. Therefore, conditional on a vector of types t and lottery numbers l the mechanism produces a well-defined outcome $x^n(t, l)$. Before lottery draws, the mechanism is defined as

$$\Phi^n(t) = \int_{l \in [0,1]^n} x^n(t, l) dl.$$

We now show that the Boston mechanism is not SP-L. Consider an economy with two proper schools, $x_0 = A, B$, and the null school $x_0 = \emptyset$, corresponding to being unmatched. That is, $X_0 = \{A, B, \emptyset\}$. Let $q_A = q_B = 1/6$. Consider a distribution $m \in \bar{\Delta}T$ such that $2/3$ of the agents prefer school A, while only $1/3$ prefer school B. Then, in a large market, the proper schools are filled in the first round with probability close to 1. Therefore, an agent has a negligible chance of getting her second choice. The chance of getting her first choice is $(1/6)/(2/3) = 1/4$ for school A and $(1/6)/(1/3) = 1/2$ for school B. That is, the limit mechanism is

$$\begin{aligned}\phi^{\infty}(t_i, m) &= \frac{1}{4}\cdot A + \frac{3}{4}\cdot \emptyset \text{ if } u_{t_i}[A] > u_{t_i}[B] \\ & \frac{1}{2}\cdot B + \frac{1}{2}\cdot \emptyset \text{ otherwise.}\end{aligned} \tag{D.1}$$

Note in particular that an agent who prefers school A faces a tradeoff when reporting her preferences. If she announces that she prefers school A, she will be assigned to it with 1/2 the chance she has of receiving school B. Therefore, it is not optimal for an agent with $u_{t_i}[A] > u_{t_i}[B] > u_{t_i}[A]/2$ to report truthfully.

Harless (2014) and Dur (2015) propose a variant of the Boston mechanism called the adaptive Boston mechanism. In the adaptive Boston mechanism, if a student points to a school where there are no seats left, then the student gets to point to the next school in her preference list (see Mennle and Seuken (2015) for a formal definition). The adaptive Boston mechanism is not SP-L. This is clear from our example above, because, in the example, schools A and B both run out of capacity in the first round.

Probabilistic Serial Mechanism

The probabilistic serial mechanism has been proposed as a solution to the assignment problem by Bogomolnaia and Moulin (2001). The mechanism works as follows. With time running continuously, agents "eat" probability shares of their favorite object, out of all objects still available. After probability shares of all objects are assigned, the objects are randomly assigned to agents according to these probabilities. We refer the reader to Kojima and Manea (2010) page 110 for a formal definition of the mechanism, as their analysis includes ours as a particular case.

Bogomolnaia and Moulin (2001) show that the mechanism is EF. Consequently, Theorem 1 guarantees that it is SP-L. Note that the fact that this mechanism is SP-L is a particular case of Kojima and Manea's Theorem 1.

Hylland and Zeckhauser Pseudo-Market Mechanism

Hylland and Zeckhauser (1979) proposed a pseudo-market mechanism for single-unit assignment, in which agents are endowed with equal budgets of an imaginary currency which they use to purchase probability shares of the objects. The mechanism works as follows. First, agents report their types, t. Second, the mechanism allocates each agent an equal budget $B > 0$ of an artificial currency. Third, the mechanism computes a competitive equilibrium price vector $p^* \in \mathbb{R}_+^{|X_0|}$ and a probabilistic allocation of goods to each agent. Each consumer's probabilistic allocation of goods is optimal given prices and the budget.

We refer the reader to the original paper for full details.

Hylland and Zeckhauser (1979) prove existence of competitive equilibria in a setting that is strictly more general than ours (in particular, they allow for indifferences). For each market size n and each possible reported vector of types $t \in T^n$, choose one such equilibrium, and use this equilibrium to define the resulting allocation $\Phi^n(t)$. To make sure that anonymity is satisfied, choose an equilibrium where all players of each type receive the same probabilistic allocation, which is always possible. As Hylland and Zeckhauser (1979) observe on page 307, since each agent has the same budget and faces the same prices, such a mechanism is EF. Consequently, Theorem 1 guarantees that it is SP-L.

D.1.3 Multi-Unit Assignment

In multi-unit assignment problems, each agent is to be assigned a finite number of indivisible objects. Transfers of a numeraire are not allowed. A prototypical application is the allocation of courses to students at business schools. For further details we refer the reader to Budish (2011).

Denote the finite set of object types by J. Each object j is available in supply $\{q_j \cdot n\}$. A bundle $x_0 \in X_0 = \mathcal{P}(J)$ specifies a subset of the object types.[7] A type t_i specifies a utility function u_{t_i} over bundles. We will adopt the terminology of course allocation, denoting object types by courses, and agents by students.

HBS Draft Mechanism

The mechanism used by Harvard Business School to allocate MBA courses was studied empirically by Budish and Cantillon (2012). Using survey data, they showed that students often misreport their preferences. Here we formally define the mechanism and show that it is not SP-L.

The HBS draft mechanism does not allow students to express preferences over bundles of courses. Instead, students submit a preference ordering over single courses. To examine the possibility of truthful reporting, we restrict our attention to preferences over bundles that are responsive to preferences over individual courses, with preferences over individual courses strict. We will say that a student of type t_i prefers course j_A to course j_B if she prefers a bundle consisting only of course j_A to a bundle consisting only of course j_B, that is, $u_{t_i}(\{j_A\}) > u_{t_i}(\{j_B\})$.

The HBS draft mechanism works as follows. First, each student is assigned a lottery number l_i, uniformly distributed in $[0, 1]$. In the first round, students take turns ordered by

[7] $\mathcal{P}(J)$ denotes the power set of J.

their lottery number, with higher lottery numbers going first. At her turn, student i chooses her favorite course out of the ones that are still available. In round two, the same procedure is repeated, but with students with lower lottery numbers going first. The procedure is repeated in the following rounds, with higher lottery numbers going first in the odd rounds and last in the even rounds. The mechanism ends after k rounds, where k is the number of courses required per student.

To see that this mechanism is not SP-L, consider the following example based closely on Example 1 of Budish and Cantillon (2012). There are 4 proper courses, $J = \{j_A, j_B, j_C, j_D\}$, of which students require $k = 2$ courses each. Each course has capacity for $\frac{2}{3}$ of the population, that is $q_j = \frac{2}{3}$ for each $j \in J$. Consider a probability distribution over students' reports where $\frac{1}{3}$ of the population lists courses in the order j_A, j_B, j_C, j_D, $\frac{1}{3}$ lists courses in the order j_B, j_A, j_C, j_D, and $\frac{1}{3}$ lists courses in the order j_A, j_C, j_D, j_B. Given this distribution of reports, the probability that course j_A reaches capacity either in the end of the first round, or early in the second round converges to 1, as the market grows large. Therefore, a student that ranks course j_A as her first choice has probability close to 1 of receiving it, while a student who ranks j_A second has probability close to 0 of receiving it. In contrast, course j_B is very likely to reach capacity either late in the second round, or early in the third round, in a large market. Consequently, a student who ranks course j_B either first or second is very likely to receive it. For this reason, a student whose true preference order is j_B, j_A, j_C, j_D profits by misreporting as j_A, j_B, j_C, j_D. By doing so, the student receives both j_A and j_B, her two favorite courses, rather than courses j_B and j_C if she reports truthfully.[8]

The Bidding Points Auction Mechanism

The bidding points auction mechanism is used by several business schools to allocate MBA courses. It has been described by Sönmez and Ünver (2010) and Krishna and Ünver (2008), who demonstrated that the mechanism is flawed in several important ways, despite its widespread use. We now define the bidding points auction mechanism and show that it is not SP-L.

The mechanism works as follows. Students report vectors of bids, with one bid per course. Students can only spend up to a budget of B points, so that the set of actions is the set of all vectors of bids that sum to at most B. We restrict the bids to be integers, so that

$$A = \{a_i \in \{0, 1, \cdots, B\}_+^J : \sum_j a_{i,j} \leq B\}.$$

[8]This particular profitable misreport is valid for any cardinal preferences consistent with the ordinal preferences j_B, j_A, j_C, j_D. In other examples the profitability of a particular misreport might depend on cardinal preference information.

Given a vector of bids, the mechanism starts with the highest bid and allocates the course to the student, as long as the course still has capacity. Ties are broken randomly.

To examine the possibility of truthful reporting, we assume that students' preferences are additive, meaning that their utility for a bundle of courses is the sum of their utilities from the component courses in that bundle. This allows us to interpret a student's bid vector as an expression of their individual course preferences, and allows us to interpret the bidding points auction as a direct mechanism with $T = A$.

Consider the case where there are three courses, $J = \{j_A, j_B, j_C\}$. Consider an agent who likes the three courses j_A, j_B, j_C equally, and derives no utility of being unmatched. That is,

$$\begin{aligned} u_{t_i}(j_A) &= u_{t_i}(j_B) = u_{t_i}(j_C) = B/3, \\ u_{t_i}(\emptyset) &= 0. \end{aligned} \tag{D.2}$$

Consider a distribution of play m, such that, in the large-market limit, the last accepted bid for the courses j_A, j_B, j_C is $2B/3$ with very high probability. In that case, the agent should not report her true preferences, with bids equal to her utility. If bids are given by equation (D.2), then with very high probability the agent does not receive any course. If instead she bids B for one of the courses she likes, and 0 for the others, she receives at least one of the courses. Therefore, the mechanism is not SP-L.

Approximate Competitive Equilibrium from Equal Incomes (A-CEEI)

Budish (2011) proposed a pseudo-market mechanism for multi-unit assignment problems. Budish's setting is a strict generalization of ours. For that reason, we do not repeat all formal definitions, and refer the reader to the original paper for further details. In our setting, the A-CEEI mechanism can be defined as follows. First, assign each student a lottery number l_i uniformly and identically distributed in $[0, 1]$. Then give each student a budget in an imaginary currency of $1 + l_i \cdot \beta_{(n)}$, where $\beta_{(n)}$ is a strictly positive constant that is weakly decreasing in n, as defined in Budish (2011) page 1081. Budish's Theorem 1 guarantees that given these budgets there exists an approximate competitive equilibrium of the economy where agents purchase courses using the imaginary currency. The A-CEEI mechanism selects one such equilibrium, anonymously, and gives each agent his equilibrium allocation. This defines a function $x^n(\cdot, \cdot)$ giving an assignment of bundles $x^n(t, l) \in X_0^n$, for each vector of types t and lottery draws l. The A-CEEI mechanism is defined as

$$\Phi^n(t) = \int_{l \in [0,1]^n} x^n(t, l) dl.$$

To show that this mechanism is SP-L, we use Theorem 1. By the definition of approximate competitive equilibrium (Budish's Definition 1), after lotteries are drawn, no agent envies another agent with a lower lottery number. Therefore, the CEEI mechanism is EF-TB, and therefore SP-L.

The Generalized Hylland and Zeckhauser Pseudo-Market

Budish et al. (2013) have proposed an extension of the Hylland and Zeckhauser pseudo-market mechanism that can be used for multi-unit assignment problems. In the simplest setting they consider, students have additive preferences over courses. We therefore assume that T only includes additive preferences. With this assumption, their setting is a strict generalization of ours. Budish et al. (2013) then formally define the mechanism. It works similarly to the Hylland and Zeckhauser mechanism, with students purchasing probability shares of courses using a fake currency. The mechanism then calculates a competitive equilibrium allocation of probability shares. Finally, the mechanism implements a lottery over allocations that gives each agent her equilibrium probability share. Budish et al.'s Theorem 6 and Corollary 3 guarantee that the mechanism is well-defined, as both an equilibrium exists and can be implemented by a lottery over feasible assignments. Budish et al.'s Theorem 8 shows that the mechanism is envy-free. Along with our Theorem 1, this implies that the mechanism is SP-L.

D.1.4 Exchange Economies

Walrasian Mechanism

A Walrasian mechanism implements competitive equilibrium allocations in an exchange economy. Several contributions in the literature have considered approximate incentive compatibility of Walrasian mechanisms in large markets, including the classic paper by Roberts and Postlewaite (1976). We refer the reader to Jackson and Manelli (1997) for an overview and references. We note that this example has an infinite set of bundles X_0, which does not fit the framework in the body of the paper. However, the mechanism fits the more general framework considered in Appendix A.1.2, which allows us to use Theorem 1 to classify it as SP-L.

We consider an exchange economy with J goods. A type $t_i = (e_{t_i}, v_{t_i})$ specifies

- An endowment vector $e_{t_i} \in \mathbb{R}^J_+$.
- A continuous utility function v_{t_i} over bundles of goods in $\mathbb{R}^J_+$, taking values in $[0, 1]$.

Assume that the finite set of types T is such that, for any finite n and type vector $t \in T^n$, there always exists at least one competitive equilibrium where all agents of the same type receive the same bundle. This is guaranteed under standard assumptions on the set of utility functions and endowment vectors.

Given a type t_i, we define the utility function u_{t_i} over net trades $x_0 \in \mathbb{R}^J$ as

$$\begin{aligned} u_{t_i} &= v_{t_i}(e_{t_i} + x_0) \text{ if } e_{t_i} + x_0 \in \mathbb{R}^J_+ \\ & \quad -\infty \text{ if } e_{t_i} + x_0 \notin \mathbb{R}^J_+. \end{aligned}$$

We let X_0 be $\mathbb{R}^J$, the set of all possible vectors of net trades.

Having defined X_0 and T, we now define the mechanism. For all n, t, $\Phi^n(t)$ anonymously selects a competitive equilibrium allocation of an economy with the n agents of types in the vector t, such that agents of the same type receive the same bundle, and assigns each agent i her vector of net trades in that equilibrium.

Note that the Walrasian mechanism is EF, as each agent receives her preferred vector of net trades given prices. Furthermore, while X_0 is not finite, it does satisfy the more general assumptions in Remark 1. Namely, X_0 is a measurable subset of Euclidean space, utility is measurable and bounded above by 1, and the utility of telling reporting truthfully is at least 0. Therefore, by Theorem 1, the Walrasian mechanism is SP-L.

D.2 Semi-Anonymous Mechanisms

Semi-anonymity generalizes anonymity to allow a mechanism to treat agents differently if they belong to identifiably distinct groups. Examples include treating men and women differently in a matching mechanism, and treating buyers and sellers differently in a double auction. While the body of the paper deals with the notationally simpler case of anonymous mechanisms, semi-anonymous mechanisms are analyzed in Appendix C. This subsection classifies some of these mechanisms.

D.2.1 Double Auctions

Double auctions have been extensively studied as a simplified model of price formation. We consider auctions where buyers and sellers submit bids, and prices are given as the average of marginal winning and losing bids. See for example Rustichini et al. (1994) for further details and references.

Types t_i specify whether an agent is a potential buyer or seller, and a value. That is, types specify the agent's group, which is $g_{t_i} = b(\text{uyer})$ or $s(\text{eller})$, and her value for the object, which is v_{t_i}. Sellers are endowed with a unit of the object, while buyers are not. The set of types is $T = G \times V$, with $G = \{b, s\}$ and $V = \{1, \cdots, \bar{v}\}$. A bundle x_0 specifies whether the agent trades or not, with a dummy $d_{x_0} = 0$ or 1, and the price of the transaction

$$p_{x_0} \in P = \{(p' + p'')/2 : p', p'' \in V\}.$$

We have $X_0 = \{0, 1\} \times P$. Buyers and sellers have quasilinear utility. The utility of a bundle is 0 if the agent does not trade. If the bundle prescribes a trade, utility is $v_{t_i} - p_{x_0}$ for a buyer, and $p_{x_0} - v_{t_i}$for a seller.

The mechanism works as follows. Given t, let $n_s(t)$ be the number of sellers, and therefore the number of objects. The market clearing price is the average of the $n_s(t)^{\text{st}}$ and $n_s(t)+1^{\text{st}}$ highest valuations. The mechanism assigns bundles x_0 with this price to all agents. The objects are assigned to the agents with the $n_s(t)$ highest valuations, with uniform tie-breaking for agents tied with the lowest winning valuation. Formally, the mechanism $\Phi^n(t)$ assigns bundles x_0 specifying trade to all buyers with valuations higher than the price, all sellers with valuations lower than the price, and randomly rations agents with valuations equal to the price.

Note that the mechanism is envy-free. This is so because all agents pay the same price, and therefore do not envy the price paid by other agents. Moreover, at this price, agents who trade with probability 1 would rather trade than not trade, and likewise agents that trade with probability 0 would rather not trade. Agents that are rationed are indifferent between trading or not trading, and therefore the mechanism is envy-free.[9] Therefore, double auctions are SP-L.

D.2.2 Matching

This setting is defined formally in Section C, Example C.1. That section also defines stable matching mechanisms, which are shown to be SP-L using a semi-anonymous version of the EF-TB condition.

Priority Match

Priority match mechanisms are described by Roth (1991), who proved that these mechanisms can produce unstable outcomes. Roth also documented that labor market clearinghouses using priority matching mechanisms were very likely to fail, and hypothesized that the reason why they failed is that they produce unstable outcomes.

[9] Note that agents are only rationed in the case of a tie between the marginal winning and losing bids, and therefore both of these bids equal the price.

The priority match works as follows. Given a man i (woman) and a woman (man) j define the rank of i on j's preferences as 1 plus the number of men (women) who are strictly preferred to i. Assign to the pair i, j the priority $p_{i,j}$ equal to the rank of the man in the woman's preferences, times the rank of the woman in the man's preferences. The mechanism then proceeds by matching pairs with the lowest priorities first, breaking ties randomly.

To see that the priority match mechanism is not SP-L, consider the case where there is a single trait for men. Then women are indifferent over all men. In this case, the priority match mechanism coincides with the Boston mechanism, which is not SP-L.

It is interesting to note that Roth (1991) conjectured that the reason why stable matching mechanisms seem to succeed in practice, while priority matching mechanisms lead to unravelling and market failures, is stability. Our analysis, however, shows that stable matching mechanisms are SP-L, while priority matching mechanisms are not. Therefore, Roth's empirical finding can be phrased equivalently as saying that SP-L mechanisms succeed while non SP-L mechanisms fail.

E Application: The Boston Mechanism

The school choice literature has debated the desirability of the commonly used Boston mechanism for student assignment (cf. Section 5.4.3). While the earliest papers on the Boston mechanism criticized it for being manipulable and argued in favor of the strategy-proof Gale-Shapley mechanism (Abdulkadiroğlu and Sönmez, 2003; Abdulkadiroğlu et al., 2006), subsequent papers showed that the Boston mechanism has Bayes-Nash equilibria that are more attractive for students, from an ex-ante welfare perspective, than the dominant strategy equilibria of Gale Shapley (Abdulkadiroğlu et al., 2011; Miralles, 2009; Featherstone and Niederle, 2011). This section applies Theorem 2 to show that there exists a mechanism that produces approximately the same outcomes as the desirable Bayes-Nash equilibria of the Boston mechanism, but that is SP-L.

E.1 Definition of the Boston Mechanism

The set of bundles is a set of schools $X_0 = S \cup \{\emptyset\}$. In a market of size n, there are $\lfloor q_s \cdot n \rfloor$ seats available in school s in S, where $q_s \in (0, 1)$ denotes the proportion of the market that s can accommodate and $\lfloor \cdot \rfloor$ is the floor function. It is assumed that X_0 includes a null school $\emptyset$ that is in excess supply. An agent of payoff type $t_i \in T$ has a utility function u_{t_i} over X_0, with no indifferences. The utility of the null school is normalized to 0. In particular, all agents strictly prefer any of the proper schools to the null school.

We consider a simplified version of the Boston mechanism with a single round.[10] The action space is the set of proper schools $A = S$, so that each student points to a school. If the number of students pointing to school s is lower than the number of seats, then all of those students are allocated to school s. If there are more students who point to s than its capacity, then students are randomly rationed, and those who do not obtain a seat in s are allocated to the null school. Formally, given a vector of reports a, the allocation $\Phi_i^n(a)$ assigns i to school a_i with probability

$$\min\left\{\frac{\lfloor q_{a_i} \cdot n \rfloor}{\text{emp}_{a_i}[a] \cdot n}, 1\right\},$$

and to the null school with the remaining probability. Consequently, the limit mechanism is

$$\phi^\infty(s, m) = \min\left\{\frac{q_s}{m_s}, 1\right\} \cdot s,$$

which denotes receiving school s with the probability $\min\{q_s/m_s, 1\}$, which we term the probability of acceptance to school s, and school $\emptyset$ with the remaining probability.

E.2 Results

The next section shows that the Boston mechanism has equilibria σ^* where $\sigma^*(t_i, \mu)$ depends continuously on beliefs μ for $\mu \in \bar{\Delta}T$. Theorem 2 then yields the following corollary:

Corollary E.1 (SP-L implementation of the Boston mechanism)**.** *The Boston mechanism has limit Bayes-Nash equilibria that depend continuously on beliefs. For any such equilibrium σ^*, the direct mechanism constructed according to equation (5.2) is SP-L, and, in the large market limit, for any prior, truthful play of the direct mechanism produces the same outcomes as equilibrium play of σ^*.*

Interestingly, the SP-L mechanism constructed by (5.2) closely resembles the Hylland and Zeckhauser (1979) pseudo-market mechanism for single-unit assignment.[11] In the constructed mechanism, agents report their types, the mechanism computes the equilibrium

[10] This simplified version of the Boston mechanism streamlines the exposition. However, this simplification means that the result in this section is stylized. An interesting question for future research is to extend the result to the standard version of the Boston mechanism, and to variations such as the adaptive Boston mechanism (Harless, 2014; Dur, 2015; Mennle and Seuken, 2015).

[11] See also Miralles (2009), which contains a very nice description of the connection between the Boston mechanism's Bayes-Nash equilibria and Hylland and Zeckhauser (1979).

market-clearing probabilities associated with the distribution of reports, and each student points to their most-preferred school given their reported types and the computed probabilities. In Hylland and Zeckhauser (1979)'s mechanism, agents report their types, the mechanism computes equilibrium market-clearing prices given the distribution of reports, and each student purchases the lottery they like best given their reported types and the computed prices.

E.3 Proof of Corollary E.1

In this section, we denote the Boston mechanism by $((\Phi)_{n\in\mathbb{N}}, S)$. The corollary uses some facts about limit equilibria of the Boston mechanism given a common identically independently distributed prior over payoff types. Let $\Sigma^*(\mu)$ denote the set of limit equilibria of the Boston mechanism given a prior μ in ΔT . Formally, denote by Σ^{**} be the set of limit Bayes-Nash equilibria of the Boston mechanism in the type space Ω^*. Then we define

$$\Sigma^*(\mu) = \{\rho \in \mathbb{R}_+^{T\times S} : \exists \sigma^* \in \Sigma^{**} \text{ such that } \rho(s, t_i) = \sigma^*(s, (t_i, \mu)) \text{ for all } s \in S, t_i \in T\}.$$

That is, each element ρ of $\Sigma^*(\mu)$ specifies the probability $\rho(s, t_i)$ with which type t_i agents play action s in a limit equilibrium of the game with a common identically independently distributed prior μ over payoff types. In other words, ρ is an equilibrium strategy profile of the Boston mechanism with set of types T and a common iid prior μ. Let $P^*(\mu)$ be the set of vectors of probability of acceptance to each school in equilibrium. We then have the following result:

Proposition E.1. *The correspondence $\Sigma^*(\mu)$ is non- empty, convex-valued and continuous in $\bar{\Delta}T$. The correspondence $P^*(\mu)$ is non-empty, single-valued, and continuous in $\bar{\Delta}T$.*

The Proposition shows that, given a prior μ, the Boston mechanism may have multiple equilibria. Nevertheless, the probability of acceptance to each school is the same in any equilibrium. The intuition is that lowering the probability of acceptance to a school weakly reduces the set of students who want to point to it, and weakly increases the set of students who want to point to other schools. Therefore, an argument similar to uniqueness arguments in competitive markets with gross substitutes shows that equilibrium probabilities of acceptance are unique. Moreover, equilibrium delivers well-behaved outcomes because probabilities of acceptance vary continuously.

Before proving the Proposition, we use it to establish Corollary E.1.

Proposition E.1 implies that Σ^* is non-empty, lower hemi-continuous, and convex-valued. The Michael Selection Theorem implies that Σ^* has a continuous selection. Thus, there exists a limit Bayes Nash equilibrium σ^* of the Boston mechanism defined over the type space Ω^*, and moreover this equilibrium $\sigma^*(t_i, \mu)$ varies continuously with μ in $\bar{\Delta}T$. Because outcomes of the Boston mechanism vary continuously with the empirical distribution of types, the social choice function $(F^n)_{n\in\mathbb{N}}$ defined by

$$F^n(\omega) = \Phi^n(\sigma^*(\omega))$$

is continuous and limit Bayes-Nash implementable. Corollary E.1 then follows from Theorem 2.

E.3.1 Proof of Proposition E.1

The Boston mechanism has a limit

$$\phi^\infty(s, m) = \min\{\frac{q_s}{m_s}, 1\}.$$

Therefore, a strategy profile ρ^* is in $\Sigma^*(\mu)$ if and only if, for all t_i and t_i' in T,

$$u_{t_i}[\phi^\infty(\rho^*(t_i'), \rho^*(\mu))] \leq u_{t_i}[\phi^\infty(\rho^*(t_i), \rho^*(\mu))].$$

In that case, we say that ρ^* is a limit Bayes-Nash equilibrium of the Boston mechanism given μ. Given a prior μ and strategy profile ρ, denote by $\rho(\mu)$ the induced distribution over actions.

We establish the Proposition in a series of claims.

Claim 1. The correspondence Σ^* is non-empty and upper hemi-continuous.

Proof. Payoffs

$$u_{t_i}[\phi^\infty(\rho(t_i), \rho(\mu))]$$

vary continuously with σ and μ. Therefore, Σ^* is non-empty and upper hemi-continuous (see Fudenberg and Tirole (1991) p. 30). □

Claim 2. For a fixed $\mu \in \Delta T$, the probabilities of acceptance to each school are the same in any limit Bayes Nash equilibrium.

Proof. Consider an equilibrium ρ. Let the mass of students pointing to school s in this equilibrium be

$$m_s = \sum_{t_i} \rho(t_i)(s) \cdot \mu(t_i)$$

and let the probability of acceptance at school s be p_s. Let the vectors $p = (p_s)_{s \in S}$ and $m = (m_s)_{s \in S}$. To establish the result, consider another equilibrium ρ', with associated vectors of the mass of students pointing to each school m' and probabilities of acceptance p'. Define the set of schools for which $p_s > p'_s$ as S^+ and the set of schools for which $p_s < p'_s$ as S^-.

Consider now the types who, in the equilibrium ρ, choose a school in S^+ with positive probability. All agents with types in

$$T^+ = \{t_i \in T : \max_{s \in S^+} u_{t_i} \cdot p_s > \max_{s \notin S^+} u_{t_i} \cdot p_s\}$$

must choose a school in S^+. That is, all agents who strictly prefer some school in S^+ to any school not in S^+ must point to one of the S^+schools in equilibrium. Therefore,

$$\sum_{t_i \in T^+} \mu_{t_i} \leq \sum_{s \in S^+} m_s.$$

Consider the types who choose a school in S^+ in the equilibrium ρ'. Note that the probability of obtaining entry to any school in S^+ is strictly lower at ρ' than at ρ from how we constructed S^+. Similarly, the probability of obtaining entry to any school not in S^+ is weakly higher. Therefore, in the equilibrium ρ', only agents in T^+ possibly choose a school in S^+ with positive probability. That is,

$$\sum_{s \in S^+} m'_s \leq \sum_{t_i \in T^+} \mu_{t_i}.$$

These two inequalities then imply that

$$\sum_{s \in S^+} m'_s \leq \sum_{s \in S^+} m_s.$$

However, for any $s \in S^+$ we have

$$m_s < m'_s,$$

because $p_s > p'_s$, and because probabilities of acceptance are determined by the mass of students pointing to each school. Taken together, these equations imply that $S^+ = \emptyset$.

Analogously, we can prove that $S^- = \emptyset$, so $p = p'$ as desired. □

Claim 3. P^* is non-empty, single-valued, and continuous.

Proof. The previous claims show that P^* is non-empty and single-valued. Moreover, P^* is upper hemi-continuous, because Σ^* is upper hemi-continuous and probabilities of acceptance depend continuously on equilibrium strategies and the distribution of types. Finally, P^* is continuous because continuity is equivalent to upper hemi-continuity for single-valued and non-empty correspondences. □

Claim 4. Σ^* is convex-valued.

Proof. Fix μ, and consider two equilibria ρ and ρ', and let $\bar{\rho}$ be a convex combination of ρ and ρ'. We must show that the strategy profile $\bar{\rho}$ is an equilibrium. By Claim 2, the probability of acceptance to each school is the same under ρ and ρ'. Therefore, the probability of acceptance is the same under $\bar{\rho}$. Because the support of $\bar{\rho}$ is contained in the union of the supports of ρ and ρ', all types play optimally under $\bar{\rho}$. □

Claim 5. Consider a prior $\mu_0 \in \bar{\Delta}T$, and associated equilibrium ρ_0 such that, for some t_i and s_0, we have $\rho_0(t_i)(s_0) > 0$. Then there exists a neighborhood of μ_0 such that, for all μ in this neighborhood, school s_0 is optimal for t_i given $P^*(\mu)$. That is, for any $s \in S$,

$$P^*_{s_0}(\mu) \cdot u_{t_i}(s_0) \geq P^*_s(\mu) \cdot u_{t_i}(s).$$

Proof. To reach a contradiction, assume that this is not the case for some type t_i' and school s_0. Then there exists a school s_1 and sequence of priors $(\mu_k)_{k\in\mathbb{N}}$ converging to μ_0 such that, for all k,

$$P^*_{s_0}(\mu_k) \cdot u_{t_i'}(s_0) < P^*_{s_1}(\mu_k) \cdot u_{t_i'}(s_1). \tag{E.1}$$

Denote the mass of t_i' types originally pointing to school s_0 as the strictly positive constant

$$C = \rho_0(t_i')(s_0) \cdot \mu_0(t_i').$$

Denote the relative increase in probability of acceptance at school s from prior μ_0 to prior μ_k by $r_s(\mu_k) = P^*_s(\mu_k)/P^*_s(\mu_0)$. We can assume, passing to a subsequence if necessary, that the ordering of schools according to $r_s(\mu_k)$ is the same for all k. Denote the schools where the probability of acceptance increases relatively more than at school s_0 as

$$S^+ = \{s : r_s(\mu_k) > r_{s_0}(\mu_k)\}.$$

Let ρ_k be an equilibrium associated with μ_k. The mass of students pointing to schools in S^+ under ρ_k minus the mass of students pointing to schools in S^+ under ρ_0 equals

$$\sum_{s\in S^+, t_i\in T} \rho_k(t_i)(s)\cdot\mu_k(t_i) - \sum_{s\in S^+, t_i\in T} \rho_0(t_i)(s)\cdot\mu_0(t_i).$$

This sum can be decomposed as

$$\sum_{s\in S^+, t_i\in T} (\rho_k(t_i)(s) - \rho_0(t_i)(s))\cdot\mu_0(t_i) \\ + \sum_{s\in S^+, t_i\in T} \rho_k(t_i)(s)\cdot(\mu_k(t_i) - \mu_0(t_i)). \tag{E.2}$$

Students who point to schools in S^+ under ρ_0 continue to do so under ρ_k. And, because equation (E.1) holds, the mass of students who point to schools in $S\backslash S^+$ under ρ_0 but who point to schools in S^+ under ρ_k is at least C. Hence, the first term in expression (E.2) is bounded below by C. Moreover, the second term converges to 0, because μ_k converges to μ_0. Therefore, for large enough k, the mass of students pointing to schools in S^+ under ρ_k is strictly larger than the mass of students pointing to schools in S^+ under ρ_0.

This implies that there exists a school $s^+ \in S^+$ such that the mass of students pointing to s^+ is strictly greater under ρ_k than under ρ_0. And there exists a school $s^- \in S\backslash S^+$ such that the mass of students pointing to s^- is strictly smaller under ρ_k than under ρ_0. However, from the way we constructed S^+ we have that $r_{s^+}(\mu_k) > r_{s^-}(\mu_k)$, which is a contradiction. □

Claim 6. Consider a prior μ_0, and associated equilibrium ρ_0 such that, for some t_i and school s_0, the mass of students pointing to s_0 is strictly lower than its capacity:

$$\sum_{t_i\in T} \rho_0(t_i)(s_0)\cdot\mu_0(t_i) < q_{s_0}.$$

Then there exists a neighborhood of μ_0 such that, for all μ in this neighborhood, $P^*_{s_0}(\mu) = 1$.

Proof. Denote the excess supply of school s_0 as the strictly positive constant

$$C = q_{s_0} - \sum_{t_i\in T} \rho_0(t_i)(s_0)\cdot\mu_0(t_i).$$

To reach a contradiction, assume that the claim's conclusion does not hold. Then there exists a sequence of priors $(\mu_k)_{k\in\mathbb{N}}$ converging to μ_0 such that, for all k, $P^*_{s_0}(\mu_k) < 1$. Let ρ_k

be an equilibrium given μ_k. The fact that the probability of acceptance at s_0 is lower than 1 under ρ_k implies that the difference between the mass of students pointing to s_0 under ρ_k and ρ_0 is bounded below by C. That is,

$$\sum_{t_i \in T} \rho_k(t_i)(s_0) \cdot \mu_k(t_i) - \sum_{t_i \in T} \rho_0(t_i)(s_0) \cdot \mu_0(t_i) > C.$$

Because μ_k converges to μ_0, this implies that, for large enough k,

$$\sum_{t_i \in T} (\rho_k(t_i)(s_0) - \rho_0(t_i)(s_0)) \cdot \mu_0(t_i) > C/2. \tag{E.3}$$

As in the previous claim's proof, denote the relative increase in the probability of acceptance at school s from prior μ_0 to prior μ_k by $r_s(\mu_k) = P_s^*(\mu_k)/P_s^*(\mu_0)$. We can assume, passing to a subsequence if necessary, that the ordering of schools according to $r_s(\mu_k)$ is the same for all k. Denote the set of schools where the relative probability of acceptance does not increase more than in s_0 by

$$S^- = \{s : r_s(\mu_k) \leq r_{s_0}(\mu_0)\} \backslash \{s_0\}.$$

All students who point to a school in $S^- \cup \{s_0\}$ under ρ_k point to schools in $S^- \cup \{s_0\}$ under ρ_0. Thus,

$$\sum_{s \in S^- \cup \{s_0\}, t_i \in T} (\rho_k(t_i)(s) - \rho_0(t_i)(s)) \cdot \mu_0(t_i) \leq 0.$$

Substituting inequality (E.3) we have that, for large enough k,

$$\sum_{s \in S^-, t_i \in T} (\rho_k(t_i)(s) - \rho_0(t_i)(s)) \cdot \mu_0(t_i) < -C/2. \tag{E.4}$$

The mass of students pointing to schools in S^- under ρ_k minus the mass of students pointing to schools in S^- under ρ_0 equals

$$\sum_{s \in S^-, t_i \in T} \rho_k(t_i)(s) \cdot \mu_k(t_i) - \sum_{s \in S^-, t_i \in T} \rho_0(t_i)(s) \cdot \mu_0(t_i).$$

This sum can be decomposed into

$$\sum_{s\in S^-,t_i\in T} (\rho_k(t_i)(s) - \rho_0(t_i)(s)) \cdot \mu_0(t_i) \\ + \sum_{s\in S^-,t_i\in T} \rho_k(t_i)(s) \cdot (\mu_k(t_i) - \mu_0(t_i)).$$

By inequality (E.4), for large enough k, the first term in the expression above is smaller than $-C/2$. Because the second term converges to 0, we have that, for sufficiently large k, the mass of students pointing to schools in S^- under ρ_k is strictly lower than the mass of students pointing to schools in S^- under ρ_0. Hence, for at least one school s^- in S^-, we have $r_{s^-}(\mu_k) \geq 1$. But this contradicts $r_{s^-}(\mu_k) \leq r_{s_0}(\mu_k) < 1$. □

Claim 7. The correspondence Σ^* is lower hemi-continuous in $\bar{\Delta}T$.

Proof. To prove lower hemi-continuity, fix μ_0, an associated limit equilibrium ρ_0, and consider a sequence $(\mu_k)_{k\geq 1}$ converging to μ_0. Fix $\epsilon > 0$. We will show that there exists a sequence of equilibria $(\rho_k)_{k\geq 1}$, associated with the μ_k, which converges to a strategy profile with distance lower than ϵ to ρ_0.

Part 1: Define the candidate sequence of equilibria.

Let ρ'_k be an equilibrium associated with μ_k. Passing to a subsequence, we can assume that $(\rho'_k)_{k\geq 1}$ converges to an equilibrium ρ'_0 associated with μ_0. Define

$$\rho_k(t_i) = \rho'_k(t_i) + (1-\epsilon) \cdot [\rho_0(t_i) - \rho'_0(t_i)] \cdot \frac{\mu_0(t_i)}{\mu_k(t_i)}.$$

Note that this sequence converges to $\epsilon \cdot \rho'_0 + (1-\epsilon) \cdot \rho_0$. Hence, it converges to a point within ϵ distance from ρ_0.

Part 2: For large enough k, ρ_k is a strategy profile.

Because the sum $\sum_s \rho_k(t_i)(s) = 1$, we only have to demonstrate that every $\rho_k(t_i)(s)$ is nonnegative. To see this, note that ρ_k converges to $\epsilon \cdot \rho'_0 + (1-\epsilon) \cdot \rho_0$. Hence, if either $\rho_0(t_i)(s) > 0$ or $\rho'_0(t_i)(s) > 0$, then $\rho_k(t_i)(s) > 0$ for sufficiently large k. The remaining case is when $\rho_0(t_i)(s) = \rho'_0(t_i)(s) = 0$. In this case we have that $\rho_k(t_i)(s) = \rho'_k(t_i)(s) \geq 0$.

Part 3: For sufficiently large k, the ρ_k are equilibria.

We will begin by proving that, for sufficiently large k, the probabilities of acceptance under ρ_k equal those under ρ'_k. That is, the probabilities of acceptance under ρ_k equal

$P^*(\mu_k)$. To see this, note that the mass of agents pointing to school s under ρ_k equals

$$\sum_{t_i} \rho_k(t_i)(s) \cdot \mu_k(t_i) = \sum_{t_i} \rho'_k(t_i)(s) \cdot \mu_k(t_i) + (1-\epsilon) \cdot \sum_{t_i} [\rho_0(t_i)(s) - \rho'_0(t_i)(s)] \cdot \mu_0(t_i). \quad \text{(E.5)}$$

There are two cases. The first case is when the mass of students pointing to s is strictly lower than q_s under either ρ_0 or ρ'_0. In this case, we have $P^*_s(\mu_0) = 1$, so that, in the mass of students pointing to s is at most equal to q_s under both ρ'_0 and ρ_0. The mass of students pointing to school s under ρ_k converges to

$$\epsilon \cdot (\sum_{t_i \in T} \rho'_0(t_i)(s)) + (1-\epsilon) \cdot (\sum_{t_i \in T} \rho_0(t_i)(s)).$$

That is, to an average of the mass of students pointing to s under ρ'_0 and ρ_0. Because both quantities are weakly smaller than q_s, and at least one of them is strictly lower than q_s, this average is strictly lower than q_s. Thus, for large enough k, the probability of acceptance to s under ρ_k is 1. This is equal to the probability of acceptance under ρ'_k, by Claim 6.

The second case is when the mass of students pointing to school s is at least equal to q_s both under ρ_0 and under ρ'_0. If this is the case, then the mass of students pointing to school s is the same under ρ_0 and under ρ'_0, because probabilities of acceptance are the same in any equilibrium under μ_0. Therefore, the sum

$$\sum_{t_i} [\rho_0(t_i)(s) - \rho'_0(t_i)(s)] \cdot \mu_0(t_i) = 0.$$

Substituting this in Equation (E.5), we have that the probabilities of acceptance under ρ_k and ρ'_k are equal, as desired.

To complete the proof we show that, for large enough k, the strategies ρ_k are optimal given $P^*(\mu_k)$. Consider a school s with $\rho_k(t_i)(s) > 0$. Therefore, either $\rho'_k(t_i)(s) > 0$ or $\rho_0(t_i)(s) > 0$. If $\rho'_k(t_i)(s) > 0$, then it is optimal for type t_i to point to s under $P^*(\mu_k)$, because ρ'_k is an equilibrium. Likewise, if $\rho_0(t_i)(s) > 0$, then Claim 5 implies that, for large enough k, it is optimal for type t_i to report s under $P^*(\mu_k)$. □

The proposition then follows from Claims 1, 3, and 7.

References

Abdulkadiroğlu, Atila and Tayfun Sönmez, "School Choice: A Mechanism Design Approach," *The American Economic Review*, 2003, *93* (3), 729–747.

_ , **Parag A. Pathak, Alvin E. Roth, and Tayfun Sönmez**, "Changing the Boston School Choice Mechanism," *NBER Working Paper No 11965*, 2006.

_ , _ , **and** _ , "Strategy-Proofness versus Efficiency in Matching with Indifferences: Redesigning the NYC High School Match," *The American Economic Review*, 2009, *99* (5), 1954–1978.

_ , **Yeon-Koo Che, and Yosuke Yasuda**, "Resolving Conflicting Preferences in School Choice: The 'Boston Mechanism' Reconsidered," *The American Economic Review*, 2011, *101* (1), 399–410.

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

Budish, Eric, "The Combinatorial Assignment Problem: Approximate Competitive Equilibrium from Equal Incomes," *Journal of Political Economy*, 2011, *119* (6), 1061–1103.

_ **and Estelle Cantillon**, "The Multi-Unit Assignment Problem: Theory and Evidence from Course Allocation at Harvard," *The American Economic Review*, 2012, *102* (5), 2237–2271.

_ , **Yeon-Koo Che, Fuhito Kojima, and Paul Milgrom**, "Designing Random Allocation Mechanisms: Theory and Applications," *The American Economic Review*, 2013, *103* (2), 585–623.

Dur, Umut, "The modified Boston mechanism," Technical Report, Mimeo, North Carolina State University 2015.

Featherstone, Clayton and Muriel Niederle, "School Choice Mechanisms Under Incomplete Information: An Experimental Investigation," *Mimeo, Harvard Business School*, 2011.

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

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

Harless, Patrick, "A School Choice Compromise: Between Immediate and Deferred Acceptance," *Mimeo, University of Rochester*, 2014.

Hylland, Aanund and Richard Zeckhauser, "The Efficient Allocation of Individuals to Positions," *The Journal of Political Economy*, 1979, *87* (2), 293–314.

Jackson, Matthew O. and Alejandro M. Manelli, "Approximately Competitive Equilibria in Large Finite Economies," *Journal of Economic Theory*, 1997, *77* (2), 354–376.

Kalai, Ehud, "Large Robust Games," *Econometrica*, 2004, *72* (6), 1631–1665.

Kojima, Fuhito and Mihai Manea, "Incentives in the Probabilistic Serial Mechanism," *Journal of Economic Theory*, 2010, *145* (1), 106–123.

_ **and Parag A. Pathak**, "Incentives and Stability in Large Two-Sided Matching Markets," *The American Economic Review*, 2009, *99* (3), 608–627.

Krishna, Aradhna and Utku Ünver, "Improving the Efficiency of Course Bidding at Business Schools: Field and Laboratory Studies," *Marketing Science*, 2008, *27* (2), 262–282.

Mennle, Timo and Sven Seuken, "Trade-offs in School Choice: Comparing Deferred Acceptance, the Naıve and the Adaptive Boston Mechanism," Technical Report, Mimeo, University of Zurich 2015.

Miralles, Antonio, "School Choice: The Case for the Boston Mechanism," in "Auctions, Market Mechanisms and Their Applications," Vol. 14 of *Lecture Notes of the Institute for Computer Sciences, Social Informatics and Telecommunications Engineering* AMMA Springer Berlin 2009, pp. 58–60.

Roberts, Donald J. and Andrew Postlewaite, "The Incentives for Price-Taking Behavior in Large Exchange Economies," *Econometrica*, 1976, *44* (1), 115–127.

Roth, Alvin E., "A Natural Experiment in the Organization of Entry Level Labor Markets: Regional Markets for New Physicians and Surgeons in the United Kingdom," *The American Economic Review*, 1991, *81* (3), 415–440.

Rustichini, Aldo, Mark A. Satterthwaite, and Steven R. Williams, "Convergence to Efficiency in a Simple Market with Incomplete Information," *Econometrica*, 1994, *62* (5), 1041–1063.

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

Swinkels, Jeroen M., "Efficiency of Large Private Value Auctions," *Econometrica*, 2001, *69* (1), 37–68.