

A Supply and Demand Framework for Two-Sided Matching Markets

Eduardo M. Azevedo

University of Pennsylvania and Microsoft Research

Jacob D. Leshno

Columbia University

This paper develops a price-theoretic framework for matching markets with heterogeneous preferences. The model departs from the Gale and Shapley model by assuming that a finite number of agents on one side (colleges) are matched to a continuum of agents on the other side (students). We show that stable matchings correspond to solutions of supply and demand equations, with the selectivity of each college playing a role similar to that of prices. We apply the model to an analysis of how competition induced by school choice gives schools incentives to invest in quality and to asymptotics of school choice mechanisms.

I. Introduction

A. Overview

In two-sided matching markets agents have preferences over whom they interact with on the other side of the market. For example, consulting firms competing for college graduates care about which workers they hire.

An earlier draft was circulated as “College Admissions with a Continuum of Students.” For comments on the earlier and present version we would like to thank Charles Angelucci, Itai Ashlagi, Dan Bjorkegren, Gabriel Carroll, Carlos da Costa, Drew Fudenberg, Oliver Hart, Thomas Kahle, Fuhito Kojima, Scott Kominers, Jose Montiel, Humberto Moreira, Alex Nichifor, Muriel Niederle, Phil Reny, Ilya Segal, Glen Weyl, Yosuke Yasuda, and seminar participants at Harvard, Chicago, Penn State, Kellogg, University of Southern California,

Such a market does not clear only through wages, as a college graduate cannot simply demand the firm she prefers; she must also be chosen by the firm. These are key features of many important markets, and matching markets have been extensively studied. Much of the literature is based on one of two classic frameworks, each with distinct advantages and limitations.¹

One strand of the literature follows Becker's (1973) marriage model. These models often assume simple preferences, with men and women being ranked from best to worst. Moreover, utility is transferable, so that a couple may freely divide the gains from marriage.² These assumptions lead to rich comparative statics that have been applied to diverse problems such as explaining sex differences in educational attainment, changes in chief executive officer wages, and the relationship between the distribution of talent and international trade (see Grossman 2004; Gabaix and Landier 2008; Terviö 2008; Chiappori, Iyigun, and Weiss 2009; Galichon, Kominers, and Weber 2016; Bojilov and Galichon 2016).

Another line of research follows Gale and Shapley's (1962) college admissions model. These models allow for complex heterogeneous preferences and (possibly) for limitations on how parties may split the surplus of a relationship. This model is a cornerstone of market design and has been applied to the study and design of market clearinghouses (e.g., the National Resident Matching Program, which matches 30,000 doctors and hospitals per year; the Boston and New York City public school matches, which match over 100,000 students per year), the use of signaling in labor markets, the relationship between matching and auctions, and supply chain networks.³ This framework has had less success in obtaining comparative statics results, especially results about the magnitude and not only the direction of an effect.⁴

This paper develops a framework to apply supply and demand analysis to matching markets. Our basic observation is that the standard equilibrium notion in matching, stability, corresponds to the solution of supply

Toulouse, University of California, Los Angeles, Columbia, and Wharton. We would like to thank specially Susan Athey, Eric Budish, and Alvin Roth for many helpful suggestions. Any remaining mistakes are our own.

¹ We focus on frictionless matching markets as opposed to markets with frictions as in the search literature.

² The basic properties of competitive and cooperative matching models with transferable utility were established by Koopmans and Beckmann (1957) and Shapley and Shubik (1971).

³ The redesign of the National Resident Matching Program is described in Roth and Peranson (1999). School choice was introduced as a mechanism design problem in the seminal paper of Abdulkadiroğlu and Sönmez (2003), and the redesign of the Boston and New York City matches is described in Abdulkadiroğlu et al. (2005) and Abdulkadiroğlu, Pathak, and Roth (2005, 2009). See also, respectively, Coles, Kushnir, and Niederle (2013), Hatfield and Milgrom (2005), Ostrovsky (2008), and Lee and Niederle (2015).

⁴ See, e.g., Gale and Sotomayor (1985) for comparative statics results on the direction of effects.

and demand equations where the selectivity of market participants plays a role similar to that of prices in standard competitive equilibrium theory. Moreover, we propose a new model of matching markets based on Aumann's (1964) insight that markets with a continuum of traders may be considerably simpler than those with a finite number of traders. Like the Gale and Shapley (1962) framework, the model allows for rich heterogeneous preferences and (possible) restrictions on transfers. However, like the Becker (1973) model, it permits straightforward derivation of comparative statics.

The basic features of our model follow the standard Gale and Shapley college admissions model. The main departure is that a finite number of agents on one side (colleges or firms) are to be matched to a continuum of agents on the other side (students or workers). As such, we model a situation in which each agent in the discrete side is matched to a large number of agents. Throughout most of the paper we consider the case in which there are no transfers. As in the standard model, the solution concept is stability. A matching between students and colleges is stable if no pair of a student and a college would like to break away from a match partner and match to each other.

Our basic results demonstrate how to use supply and demand to analyze stable matchings. We show that, even in the discrete model, stable matchings have a simple structure, given by admission thresholds P_c at each college c . We term such a threshold a cutoff, and colleges are said to accept students who are ranked above the cutoff. Given a vector of cutoffs P , a student's demanded college is defined as her favorite college that would accept her. We show that, for every stable matching, there exists a vector of cutoffs such that each student demands the college she is matched to. Moreover, at any vector of cutoffs P that clears supply of and demand for colleges, the demand function yields a stable matching. Therefore, finding stable matchings is equivalent to solving market clearing equations⁵

$$D(P) = S.$$

The continuum assumption considerably simplifies the analysis. The continuum model typically has a unique stable matching. This stable matching corresponds to the unique solution of the market clearing equations and varies continuously with the underlying fundamentals. Thus, comparative statics may be derived from the market clearing equations, as in standard price-theoretic arguments.

We give convergence results to clarify when the continuum model is a good approximation to real markets. A sequence of discrete economies is

⁵ As we discuss formally below, the formula $D(P) = S$ holds only when there is excess demand for all colleges.

said to converge to a continuum economy if the empirical distribution of student types converges to the distribution in the continuum economy and the number of seats per student in each college converges. Whenever the continuum economy has a unique stable matching, all of the stable matchings of the discrete economies converge to this unique stable matching of the continuum economy. In particular, all stable matchings of the large discrete economies become similar. Therefore, even in a large discrete economy, stability clears the market in a way that is essentially unique.

We consider two applications. The first application is a price-theoretic analysis of the effects of school competition. We consider how competition among schools induced by flexible school choice gives incentives for schools to invest in quality. This problem has been studied in the discrete Gale and Shapley framework by Hatfield, Kojima, and Narita (2016) and in a simplified model by Hoxby (1999). We consider a setting in which schools compete for students and determine how much schools benefit from investing in quality in terms of attracting a stronger entering class. The continuum model clarifies how incentives depend on the types of students catered to, the distribution of preferences, and market structure. Schools have muted, or possibly even negative, incentives to perform quality improvements that target lower-ranked students. Moreover, these concerns are exacerbated when schools have market power. Therefore, while school choice might give schools incentives to improve, such improvements may disproportionately benefit top students.

In the second application we characterize the asymptotics of a large class of matching mechanisms. In particular, we characterize the asymptotics of school choice mechanisms used to match students to schools, such as deferred acceptance with single tie breaking. Che and Kojima (2010) characterized the limit of the random serial dictatorship mechanism. This is a particular case of deferred acceptance in which schools treat all students equally. We extend their asymptotic results to the case in which schools give priority to subsets of students. As a corollary, we show that the deferred acceptance mechanism may produce Pareto-inefficient outcomes with high probability. This application is based on an extension of our convergence result to randomly drawn preferences. We report convergence rates and simulations, which clarify for what market sizes the approximation is useful.

Proofs of all results are in the online appendix.

B. Related Literature

Our paper is related to several active lines of research. First, it is related to the literature on large matching markets. Roth and Peranson (1999) observed that, even though stable matching mechanisms are not strategy-proof, they are difficult to manipulate in large markets. Subsequent pa-

pers have justified this finding theoretically (Immorlica and Mahdian 2005; Kojima and Pathak 2009; Lee 2014). Our work differs from this literature in two aspects. The first is that previous work has focused on showing approximate incentive compatibility of stable mechanisms. In contrast, we characterize the limit of the set of stable matchings in large matching markets. Second, the type of limit we take is different. While papers in this literature consider the limit in which both sides of the market grow, we consider the case in which there is a fixed, finite number of colleges and the number of students grows.

Another contribution of our paper is the characterization of stable matchings in terms of cutoffs clearing supply and demand. We highlight several related results in the literature. An early result by Roth and Sotomayor (1989) shows that different entering classes in a college at different stable matchings are ordered in the sense that, except for students who are in both entering classes, all students of an entering class are better than those in the other entering class. This suggests that parameterizing the set of stable matchings by the lowest-ranked student is possible, though their result does not describe such a parameterization. Balinski and Sönmez (1999) give a characterization of fair allocations based on threshold scores. Sönmez and Ünver (2010) propose a mechanism for the course allocation problem in which students place bids for courses and report preferences. Their proposition 1 shows that a deferred acceptance algorithm using bids as preferences for courses leads to thresholds such that students are matched to their preferred affordable bundle of courses. Biró (2007) studies the algorithm used for college admissions in Hungary. The algorithm, while similar to Gale and Shapley's, uses scores. Biró states that a definition of stability in terms of cutoffs is equivalent to the standard definition. Fleiner and Jankó (2014) propose generalized notions of stability based on scores. Abdulkadiroğlu, Che, and Yasuda (2015) use a cutoff characterization in a particular case of our model.

Adachi (2000) gives a characterization of stable matchings in terms of fixed points of an operator over prematchings.⁶ This insight has been widely applied in the matching with contracts literature (see, e.g., Hatfield and Milgrom 2005; Ostrovsky 2008; Echenique 2012). This characterization is different from the one in terms of cutoffs, as prematchings are considerably more complex than cutoffs. In fact, a prematching specifies a college for each student and a set of students for each college so that the dimensionality of the set of prematchings is much higher than the set of cutoffs. As such, this characterization is more useful for deriving theoretical results as opposed to simple comparative statics. Online ap-

⁶ These ideas have been extended to many-to-one and many-to-many matching by Echenique and Oviedo (2004, 2006). See also Hatfield and Milgrom (2005), Echenique and Yenmez (2007), Ostrovsky (2008), Hatfield and Kominers (2012, 2016), and Echenique (2012).

pendix E details the connection between our work and Adachi's characterization.

Our characterization lemma 2 is analogous to the fundamental theorems of welfare economics. Segal (2007) shows that these theorems may be stated for a wide class of social choice problems: namely, socially optimal outcomes can be decentralized with a notion of price equilibrium that is appropriate for the problem. Furthermore, he characterizes the prices that verify a problem's solutions with minimal communication (Segal's theorems 2 and 3). Applied to stable many-to-one matching, his characterization yields prematchings as the appropriate prices (Segal's proposition 5). In our model, where colleges' preferences are defined by students' scores, the minimally informative prices in Segal's theorem 3 coincide with our notion of market clearing cutoffs.

Like our paper, Bodoh-Creed and Hickman (2015), Chade, Lewis, and Smith (2014), and Abdulkadiroğlu et al. (2015) consider models with a finite number of colleges and a continuum of students. Abdulkadiroğlu et al. study the deferred acceptance with a multiple tie-breaking mechanism in a setting in which schools do not give priorities to different students (the assignment problem). They use a continuum of students and noted that outcomes of the mechanism can be described with cutoffs. As such, outcomes of this mechanism are essentially the same as in our paper in the case in which (1) scores are uniformly distributed in $[0, 1]^C$ (in particular, colleges' preferences are uncorrelated with each other), (2) student preferences have full support and are uncorrelated with college preferences, and (3) all colleges have the same capacity, which is exactly sufficient to accommodate all students. Unfortunately, it is not possible to directly use their results in our applications because using a specific preference structure precludes establishing comparative statics or general limit results.

The proof of Abdulkadiroğlu et al.'s lemma 4 has an argument demonstrating that cutoffs that equate supply and demand are unique. This argument introduces important technical ideas that we use in our proof. The commonalities and differences in the proofs are as follows. Both proofs compare the mass of students matched to undesirable schools under the greatest and smallest market clearing cutoffs and use a fixed-point map to establish that market clearing cutoffs exist and form a lattice. The differences are that, without their assumptions 1–3, described in the previous paragraph, we need a generalization of their map to establish the lattice result (see the working paper version of this article [Azevedo and Leshno 2013] for an alternative argument based on standard matching theory) and that we use the rural hospitals theorem to compare the mass of students matched to different schools under different stable matchings, while in their model each school is matched to the same number of students in every stable matching by definition.

II. The Continuum Model

A. Definitions

A finite set of colleges $C = \{1, 2, \dots, C\}$ is to be matched to a continuum mass of students. A student is described by her type $\theta = (\succ^\theta, e^\theta)$. The student's strict preference ordering over colleges is \succ^θ . The vector $e^\theta \in [0, 1]^C$ describes the colleges' ordinal preferences for the student. We refer to e_c^θ as student θ 's score at college c , and colleges prefer students with higher scores. That is, college c prefers student θ over θ' if $e_c^\theta > e_c^{\theta'}$. Colleges' preferences over sets of students are responsive (Roth 1985). To simplify notation we assume that all students and colleges are acceptable. Let \mathcal{R} be the set of all strict preference orderings over colleges. We denote the set of all student types by $\Theta = \mathcal{R} \times [0, 1]^C$.

A continuum economy is given by $E = [\eta, S]$, where η is a probability measure over Θ and $S = (S_1, S_2, \dots, S_C)$ is a vector of strictly positive capacities for each college.⁷ We make the following assumption on η , which corresponds to colleges having strict preferences over students in the discrete model.

ASSUMPTION 1 (Strict preferences). Every college's indifference curves have η measure 0. That is, for any college c and real number x we have $\eta(\{\theta : e_c^\theta = x\}) = 0$. The set of all economies satisfying this assumption is denoted by \mathcal{E} .

A matching μ describes an allocation of students to colleges. Formally, a matching is a function $\mu : C \cup \Theta \rightarrow 2^\Theta \cup (C \cup \Theta)$ such that

1. for all $\theta \in \Theta$, $\mu(\theta) \in C \cup \{\theta\}$;
2. for all $c \in C$, $\mu(c) \subseteq \Theta$ is measurable and $\eta(\mu(c)) \leq S_c$;
3. $c = \mu(\theta)$ iff $\theta \in \mu(c)$;
4. (open on the right) for any $c \in C$, the set $\{\theta \in \Theta : \mu(\theta) \leq^c c\}$ is open.

The definition of a matching is analogous to that in the discrete model. Conditions 1–3 mirror those in the discrete model. Condition 1 states that each student is matched to a college or to herself, which represents being unmatched. Condition 2 states that colleges are matched to sets of students with measure not exceeding its capacity. Condition 3 is a consistency condition, requiring that a college is matched to a student iff the student is matched to the college.

Condition 4 is imposed because in the continuum model it is always possible to add a measure 0 set of students to a college without exceeding capacity. This would generate multiplicities of stable matchings that

⁷ We must also specify a σ -algebra in which η is defined. We take the Borel σ -algebra of the product topology in Θ . We will also use this topology in Θ .

differ only in sets of measure 0. Condition 4 rules out such multiplicities. The intuition is that the condition implies that a stable matching always allows an extra measure 0 set of students into a college when this can be done without compromising stability.

A student-college pair (θ, c) *blocks* a matching μ at economy E if the student θ prefers c to her match and either (i) college c does not fill its quota or (ii) college c is matched to another student who has a strictly lower score than θ . Formally, (θ, c) blocks μ if $c \succ^\theta \mu(\theta)$ and either (i) $\eta(\mu(c)) < S_c$ or (ii) there exists $\theta' \in \mu(c)$ with $e_c^{\theta'} < e_c^\theta$.

DEFINITION 1. A matching μ for a continuum economy E is *stable* if it is not blocked by any student-college pair.

A stable matching always exists (see corollary A.1 in the appendix). The simplest proof is similar to Gale and Shapley's (1962) classic existence proof in the discrete case and works by showing that a deferred acceptance procedure converges to a stable matching. This was shown in a particular case by Abdulkadiroğlu et al. (2015).

We refer to the stable matching correspondence as the correspondence associating each economy in \mathcal{E} with its set of stable matchings. In some sections in the paper the economy is held fixed. Whenever there is no risk of confusion we will omit dependence of certain quantities on the economy.

B. The Supply and Demand Characterization of Stable Matchings

Throughout this subsection, we fix an economy E and abuse notation by omitting dependence on E . A *cutoff* is a minimal score $P_c \in [0, 1]$ required for admission at a college c . We say that a student θ can afford college c if $P_c \leq e_c^\theta$, that is, c would accept θ . A student's *demand* given a vector of cutoffs is her favorite college among those she can afford. If no colleges are affordable, define $D^\theta(P) = \emptyset$, meaning that the student demands being unmatched. *Aggregate demand* for college c is the mass of students who demand it,

$$D_c(P) = \eta(\{\theta : D^\theta(P) = c\}).$$

The aggregate demand vector $(D_c(P))_{c \in C}$ is denoted $D(P)$.

A market clearing cutoff is a vector of cutoffs that clears supply of and demand for colleges.

DEFINITION 2. A vector of cutoffs P is a *market clearing cutoff* if it satisfies the following market clearing equations:

$$D_c(P) \leq S_c$$

for all c and

$$D_c(P) = S_c$$

if $P_c > 0$.

There is a natural one-to-one correspondence between stable matchings and market clearing cutoffs, described by the following operators. Note that the operators are defined only for market clearing cutoffs and stable matchings, as opposed to all matchings and cutoffs. Given a market clearing cutoff P , define the associated matching $\mu = \mathcal{M}P$ with the demand function

$$\mu(\theta) = D^\theta(P).$$

Conversely, given a stable matching μ , define the associated cutoff $P = \mathcal{P}\mu$ by the score of marginal students matched to each college:

$$P_c = \inf_{\theta \in \mu(c)} e_c^\theta. \quad (1)$$

LEMMA 1 (Supply and demand lemma). If μ is a stable matching, then $\mathcal{P}\mu$ is a market clearing cutoff. If P is a market clearing cutoff, then $\mathcal{M}P$ is a stable matching. In addition, the operators \mathcal{P} and \mathcal{M} are inverses of each other.

The lemma shows that stable matchings all have a special structure. Given any stable matching μ , there exist corresponding cutoffs such that each student is matched to $\mu(\theta) = D^\theta(P)$. Therefore, any stable matching corresponds to each student choosing her favorite college conditional on being accepted at a vector of cutoffs P .

Intuitively, the lemma says that the selectivity of each college works similarly to prices in the following sense. In a competitive market, the price of each good is determined in equilibrium so that supply equals demand. In a matching market, however, students not only choose colleges but also must be chosen back. The lemma says that, instead of prices adjusting, the selectivity of each college is endogenously determined to clear the market.⁸

C. Example: Stable Matchings and Convergence of Discrete Economies to Continuum Economies

There are two colleges $c = 1, 2$, and the distribution of students η is uniform. That is, there is a mass $1/2$ of students with each preference list $1, 2$ or $2, 1$, and each mass has scores distributed uniformly over $[0, 1]^2$. If both colleges had capacity $1/2$, the unique stable matching would have each student matched to her favorite college. To make the example interesting, assume $S_1 = 1/4$ and $S_2 = 1/2$. That is, college 2 has enough

⁸ Online app. D extends this result to a model with flexible wages.

seats for all students who prefer college 2, but college 1 has capacity for only half of the students who prefer it.

A familiar way of finding stable matchings is the student-proposing deferred acceptance algorithm. Abdulkadiroğlu et al. (2015) formally define the algorithm and prove that it converges. Here, we informally follow the algorithm for this example to build intuition on the structure of stable matchings. At each step of the algorithm, unassigned students propose to their favorite college out of the ones that have not rejected them yet. If a college has more students than the capacity assigned to it, it rejects the lower-ranked students to stay within capacity. Figure 1 displays the trace of the algorithm in our example. In the first step, all students apply to their favorite college. Because college 1 has only capacity $1/4$ and each square has mass $1/2$, it then rejects half of the students who applied. The rejected students then apply to their second choice, college 2. But this leaves college 2 with $1/2 + 1/4 = 3/4$ students assigned to it, which is more than its quota. College 2 then rejects its lower-ranked students, and the process continues. Although the algorithm does not reach a stable matching in a finite number of steps, it converges, and its pointwise limit (shown in fig. 2) is a stable matching.

Figures 1 and 2 illustrate the cutoff structure of stable matchings. Indeed, we could have computed the stable matching by solving the market clearing equations. Consider first demand for college 1. The fraction of students in the left square of figure 2 demanding college 1 is $1 - P_1$ and in the right square it is $P_2(1 - P_1)$. Therefore, $D_1(P) = (1 + P_2)(1 - P_1)/2$. Demand D_2 has an analogous formula, and the market clearing equations are

$$S_1 = 1/4 = (1 + P_2)(1 - P_1)/2,$$

$$S_2 = 1/2 = (1 + P_1)(1 - P_2)/2.$$

Solving this system, we get $P_1 = (\sqrt{17} + 1)/8$ and $P_2 = (\sqrt{17} - 1)/8$. In particular, because the market clearing equations have a unique solution, the economy has a unique stable matching. Theorem 1 below shows that this is a much more general phenomenon.

We show below that the cutoff characterization is also valid in the discrete Gale and Shapley (1962) model, except for the fact that in a discrete model each stable matching may correspond to more than one market clearing cutoff. Figure 3 illustrates a stable matching in a discrete economy with 1,000 students. The 1,000 students were assigned random types, drawn from the distribution η used in the continuum example. In this sense, this discrete economy approximates the continuum economy. Note that the cutoffs in the discrete economy are approximately the same as the cutoffs in the continuum economy. Theorem 2 shows that,

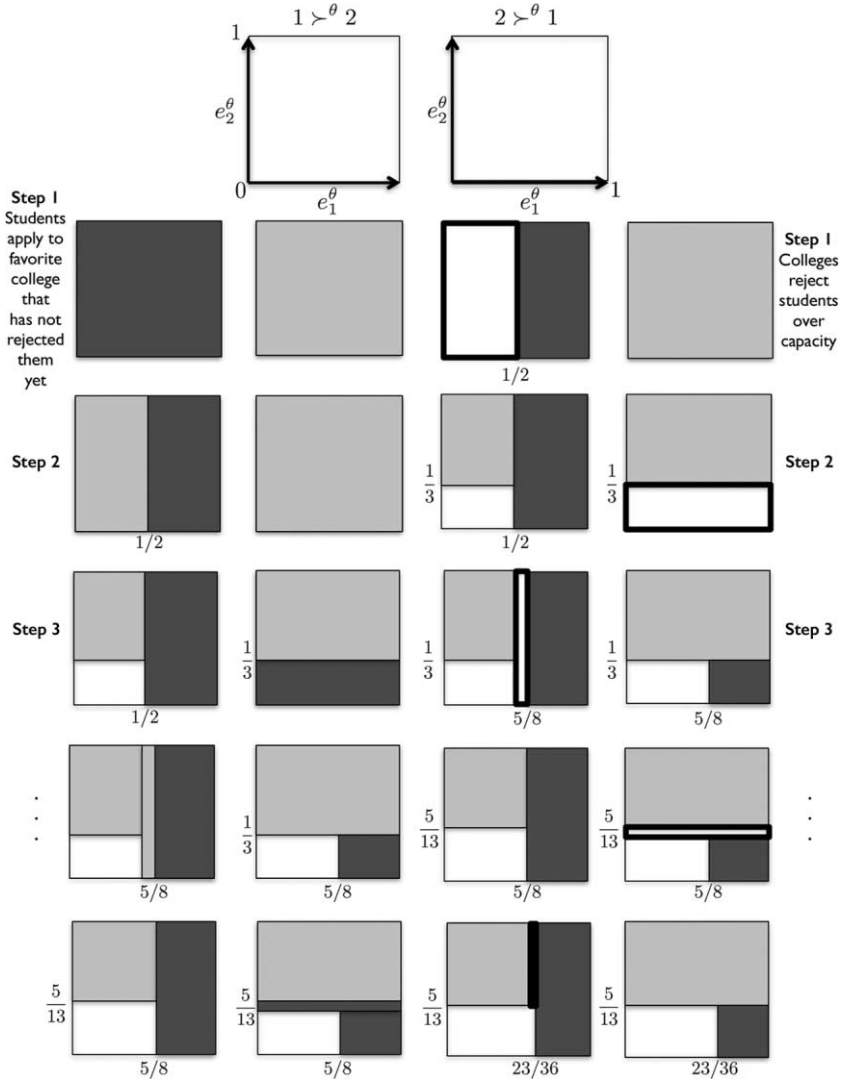


FIG. 1.—The Gale and Shapley algorithm. The set of student types Θ is represented by the two squares on the top panel. The left square represents students who prefer college 1 and the right square students who prefer college 2. Scores at each college are represented by the (x, y) coordinates. The lower panels show the first five steps of the Gale-Shapley student-proposing algorithm. In each line, students apply to their favorite colleges that have not yet rejected them in the left panel, and colleges reject students to remain within capacity in the right panel. Students in dark gray are tentatively assigned to college 1 and in light gray tentatively assigned to college 2.

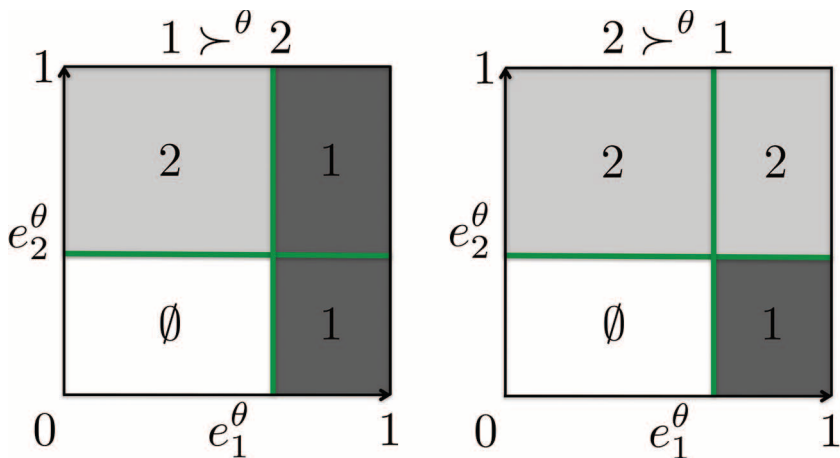


FIG. 2.—A stable matching in a continuum economy with two colleges. The two squares represent the set of student types Θ . The left square represents students who prefer college 1 and the right square students who prefer college 2. Scores at each college are represented by the (x, y) coordinates. The white area represents unmatched students, dark gray are matched to college 1, and light gray to college 2.

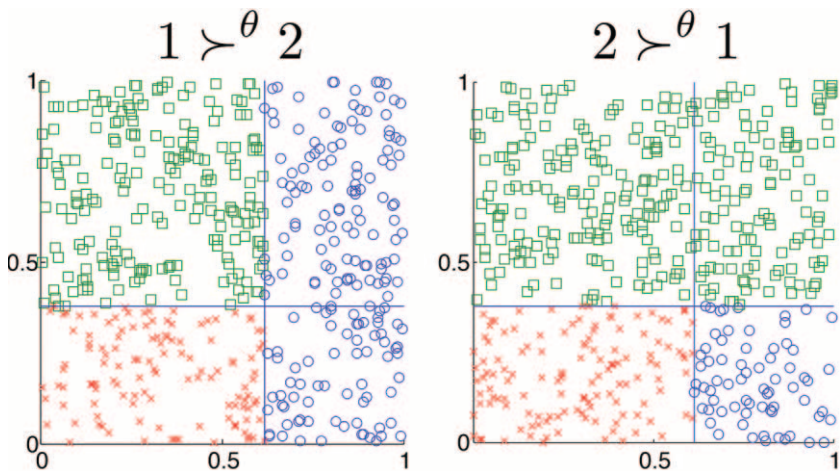


FIG. 3.—A stable matching in a discrete economy approximating the example. The two squares represent the set of student types Θ . The left square represents students who prefer college 1 and the right square students who prefer college 2. Scores at each college are represented by the (x, y) coordinates. There are two colleges, with capacities 250 and 500: 500 students have preferences $\succ^{\theta} = 1, 2$, \emptyset and 500 students have preferences 2, 1, \emptyset . Scores e^{θ} were drawn independently according to the uniform distribution in $[0, 1]^2$. The figure depicts the student-optimal stable matching. Circles represent students matched to college 1, squares represent students matched to college 2, and Xs represent unmatched students.

generically, the market clearing cutoffs of approximating discrete economies approach market clearing cutoffs of the continuum economy.

III. Convergence and Uniqueness

This section establishes conditions for the continuum model to have a unique stable matching and for it to correspond to the limit of the discrete Gale and Shapley (1962) model. Readers purely interested in applications may skip to Section IV, where the key ideas can be understood without Section III. We begin with definitions necessary to state the results.

A. The Discrete Gale and Shapley Model

The set of colleges is again C . A *finite economy* $F = [\tilde{\Theta}, \tilde{S}]$ specifies a finite set of students $\tilde{\Theta} \subset \Theta$ and an integer vector of capacities $\tilde{S}_c > 0$ for each college.⁹ We assume that no college is indifferent between two students in $\tilde{\Theta}$. A *matching* for finite economy F is a function $\tilde{\mu} : C \cup \tilde{\Theta} \rightarrow 2^{\tilde{\Theta}} \cup (C \cup \tilde{\Theta})$ such that

1. for all θ in $\tilde{\Theta}$, $\mu(\theta) \in C \cup \{\theta\}$;
2. for all $c \in C$, $\mu(c) \in 2^{\tilde{\Theta}}$ and $\#\mu(c) \leq \tilde{S}_c$;
3. for all $\theta \in \tilde{\Theta}$, $c \in C$, $\mu(\theta) = c$ iff $\theta \in \mu(c)$.

These conditions may be interpreted as follows. (1) Each student is matched to a college or to herself, (2) each college is matched to at most \tilde{S}_c students, and (3) in the consistency condition, a college is matched to a student iff the student is matched to the college.

The definition of a *blocking pair* is the same as in Section II.A. A matching $\tilde{\mu}$ is said to be *stable* for finite economy F if it has no blocking pairs.

Given a finite economy $F = [\tilde{\Theta}, \tilde{S}]$, we may associate with it the empirical distribution of types

$$\eta = \sum_{\theta \in \tilde{\Theta}} \frac{1}{\#\tilde{\Theta}} \delta_\theta,$$

where δ_θ denotes the probability distribution placing probability one on the point θ . The supply of seats per student is given by $S = \tilde{S}/\#\tilde{\Theta}$. Note that $[\eta, S]$ uniquely determine a discrete economy $F = [\tilde{\Theta}, \tilde{S}]$ as $\tilde{\Theta} = \text{support}(\eta)$ and $\tilde{S} = S \cdot \#\tilde{\Theta}$. Therefore, either pair $[\tilde{\Theta}, \tilde{S}]$ or $[\eta, S]$ uniquely determine a finite economy F . Throughout the remainder of the text we will abuse notation and refer to finite economies simply as

⁹ We use the notation F for finite economies to avoid confusion over results involving sequences of finite and continuum economies.

$F = [\eta, S]$. This convention will be useful for stating our convergence results below, as it makes finite economies F comparable to continuum economies E .

Cutoffs.—In the remainder of Section III.A, we fix a finite economy $F = [\eta, S]$ and will omit dependence on F in the notation. A cutoff is a number P_c in $[0, 1]$ specifying an admission threshold for college c . Given a vector of cutoffs P , a student's individual demand $D^\theta(P)$, the aggregate demand function $D(P)$, and market clearing cutoffs are defined as in Section II.B.

In the discrete model we define the operators $\tilde{\mathcal{M}}$ and $\tilde{\mathcal{P}}$, which have essentially the same definitions as \mathcal{M} and \mathcal{P} . Given market clearing cutoffs P , $\tilde{\mu} = \tilde{\mathcal{M}}P$ is the matching such that for all $\theta \in \tilde{\Theta}$, $\tilde{\mu}(\theta) = D^\theta(P)$. Given a stable matching $\tilde{\mu}$, $P = \tilde{\mathcal{P}}\tilde{\mu}$ is given by $P_c = 0$ if $\eta(\tilde{\mu}(c)) < S_c$ and $P_c = \min_{\theta \in \tilde{\mu}(c)} e_c^\theta$ otherwise. We have the following analogue of the supply and demand lemma.

LEMMA 2 (Discrete supply and demand lemma). In a discrete economy, the operators $\tilde{\mathcal{M}}$ and $\tilde{\mathcal{P}}$ take stable matchings into market clearing cutoffs, and vice versa. Moreover, $\tilde{\mathcal{M}}\tilde{\mathcal{P}}$ is the identity.

B. Convergence Notions

To describe our convergence results, we must define notions of convergence for economies and stable matchings. We will say that a sequence of continuum economies $\{E^k\}_{k \in \mathbb{N}}$, $E^k = [\eta^k, S^k]$, converges to a continuum economy $E = [\eta, S]$ if the measures η^k converge in the weak sense to η and if the vectors S^k converge to S .

Throughout the paper, we use the sup norm $\|\cdot\|$ whenever considering vectors in Euclidean space. We take the distance between stable matchings to be the distance between their associated cutoffs. That is, the distance between two stable matchings μ and μ' is

$$d(\mu, \mu') = \|\mathcal{P}\mu - \mathcal{P}\mu'\|.$$

A sequence of finite economies $\{F^k\}_{k \in \mathbb{N}}$, $F^k = [\eta^k, S^k]$, converges to a continuum economy $E = [\eta, S]$ if the empirical distribution of types η^k converges to η in the weak sense and the vectors of capacity per student S^k converge to S . Given a stable matching of a continuum economy μ and a stable matching of a finite economy $\tilde{\mu}$, we define

$$d(\tilde{\mu}, \mu) = \sup_P \|P - \mathcal{P}\mu\|$$

over all vectors P with $\tilde{\mathcal{M}}P = \tilde{\mu}$. The sequence of stable matchings $\{\tilde{\mu}^k\}_{k \in \mathbb{N}}$ with respect to finite economies F converges to stable matching μ of continuum economy E if $d(\tilde{\mu}^k, \mu)$ converges to 0.

Finally, we will show that the set of stable matchings of large finite economies becomes small under certain conditions. To state this, define the *diameter* of the set of stable matchings of a finite economy F as

$$\sup\{\|P - P'\| : P \text{ and } P' \text{ are market clearing cutoffs of } F\}.$$

C. Convergence and Uniqueness Results

We are now ready to state the uniqueness and convergence results. The first result shows that continuum economies typically have a unique stable matching. In this section, because we consider sequences of economies, we will explicitly denote the dependence of demand functions on measures as $D(\cdot|\eta)$ and on economies as $D(\cdot|E)$ or $D(\cdot|F)$. We begin by defining a notion of smoothness of measures.

DEFINITION 3. The distribution of student types η is *regular* if the image under $D(\cdot|\eta)$ of the closure of the set

$$\{P \in (0, 1)^c : D(\cdot|\eta) \text{ is not continuously differentiable at } P\}$$

has Lebesgue measure 0.

This condition is satisfied, for example, if $D(\cdot|\eta)$ is continuously differentiable or if η admits a continuous density. However, it also includes cases in which demand is not differentiable, as in the analysis of matching mechanisms in Section IV.C. The next result gives conditions for the continuum model to have a unique stable matching.

THEOREM 1. Consider an economy $E = [\eta, S]$.

1. If η has full support, then E has a unique stable matching.
2. If η is any regular distribution, then for almost every vector of capacities S with $\sum_c S_c < 1$, the economy E has a unique stable matching.

The theorem has two parts. First, whenever η has full support, a continuum economy has a unique stable matching. Therefore, whenever the set of students is rich enough, an economy has a unique stable matching. Moreover, even if the full support assumption does not hold, in a very general setting for almost every S , there exists a unique stable matching. Therefore, the typical case is for the continuum model to have a unique stable matching.

The proof is based on extensions of classic results in matching theory provided in the online appendix. The lattice theorem shows that there exist smallest and largest vectors of market clearing cutoffs. Moreover, the rural hospitals theorem guarantees that the number of unmatched students is the same under the highest and lowest market clearing cutoffs. This implies that demand is constant among cutoffs between the smallest and largest market clearing cutoffs. The first part of the theo-

rem follows because, with full support, demand is strictly decreasing. The intuition for the second part is that supply and demand almost never intersect in a region where demand is constant. This intuition is similar to the case of a supply and demand model with a single good, where it is a knife-edge case for demand to be vertical at the equilibrium price. The proof is based on a result in analysis known as Sard's theorem. This approach was introduced by Debreu (1970) in general equilibrium theory. Varian (1975) later applied similar ideas in differential topology to uniqueness of competitive equilibrium.

The next theorem connects the continuum model and the discrete Gale and Shapley model. It shows that when an economy E has a unique stable matching, which is the generic case, (1) it corresponds to the limit of stable matchings of approximating finite economies, (2) approximating finite economies have a small set of stable matchings, and (3) the unique stable matching varies continuously with fundamentals.

THEOREM 2. Assume that the continuum economy E admits a unique stable matching μ . We then have the following statements:

1. For any sequence of stable matchings $\{\tilde{\mu}^k\}_{k \in \mathbb{N}}$ of finite economies in a sequence $\{F^k\}_{k \in \mathbb{N}}$ converging to E , we have that $\tilde{\mu}^k$ converges to μ .
2. Moreover, the diameter of the set of stable matchings of $\{F^k\}_{k \in \mathbb{N}}$ converges to 0.
3. The stable matching correspondence is continuous at E within the set of continuum economies \mathcal{E} .

Part 1 justifies using the simple continuum model as an approximation of the Gale and Shapley (1962) model. Formally, the unique stable matching of a continuum economy is the limit of any sequence of stable matchings of approximating finite economies. We emphasize that, for a sequence of finite economies $\{F^k\}_{k \in \mathbb{N}}$ to converge to a continuum economy E , it is necessary that the empirical distribution of student types converges. Therefore, the economies F^k have an increasing number of students and a fixed number of colleges. Section IV.B gives convergence rate results and discusses when the continuum approximation is appropriate.

Part 2 states that the diameter of the set of stable matchings of any such sequence of approximating finite economies converges to 0. This means that, as economies in the sequence become sufficiently large, the set of stable matchings becomes small. More precisely, even if such an economy has several stable matchings, cutoffs are very similar in any stable matching.

Finally, part 3 states that the unique stable matching of E varies continuously in the set of all continuum economies \mathcal{E} . That is, the stable matching varies continuously with the fundamentals of economy E . Part 3 is of significance for studies that use data and simulations to inform market

design (Abdulkadiroğlu et al. 2009). It implies that, in large matching markets, the conclusions of such simulations are not sensitive to small changes in fundamentals.

We now provide an example of a continuum economy with multiple stable matchings.¹⁰ This shows that existence of a unique stable matching cannot be guaranteed in general. Moreover, we show that when a continuum economy has multiple stable matchings, very generally none of them are robust to small perturbations of fundamentals. This implies that the conclusions of theorem 2, linking the discrete and continuum models, are not valid when there are multiple stable matchings.

Example (Multiple stable matchings).—There are two colleges with capacity $S_1 = S_2 = 1/2$. Students differ in their height. While college 1 prefers taller students, college 2 prefers shorter students. To model this we assume that scores are uniformly distributed in the segment $[(1, 0), (0, 1)]$. Student preferences are uncorrelated with height, and half of the students prefer each college.

Under these assumptions, $P = (0, 0)$ clears the market, with demand $1/2$ for each college. Likewise, $P = (1/2, 1/2)$ corresponds to a stable matching, with the taller half of the population going to college 1 and the shorter half to college 2. These are, respectively, the student-optimal and college-optimal stable matchings. For P_1 and P_2 in $[0, 1/2]$, the demand functions are

$$D_c(P) = 1/2 - (P_c - P_c). \quad (2)$$

Therefore, any cutoff vector in the segment $[(0, 0), (1/2, 1/2)]$ corresponds to a stable matching.

Note that none of these stable matchings is robust to small perturbations of fundamentals. Consider adding a small amount of capacity to each college. If this is done, at least one of the colleges must be in excess supply and have a cutoff of 0 in equilibrium. This implies that the other college will have a cutoff of 0. Therefore, the only stable matching would correspond to $P = (0, 0)$. Likewise, in an economy in which each college had a slightly smaller capacity, any market clearing cutoff involves cutoffs greater than $1/2$. Otherwise, the demand equation (2) would imply that there is excess demand for at least one of the colleges.

The following proposition generalizes the example. It shows that, when the set of stable matchings is large, none of the stable matchings are robust to small perturbations.

PROPOSITION 1 (Nonrobustness). Consider an economy E with $\sum_c S_c < 1$. Let P be a market clearing cutoff. Assume that there exists another market clearing cutoff that is either strictly larger or strictly smaller than P . Let N be a sufficiently small neighborhood of P . Then there exists a

¹⁰ We would like to thank Ted Bergstrom for suggesting this example.

sequence of economies E^k converging to E without any market clearing cutoffs in N .

IV. Applications and Extensions

A. Competition and School Quality

This section considers the classic question of whether competition between public schools improves school quality and illustrates the derivation of comparative statics in our model. Hatfield et al. (2016) consider an important aspect of this problem, namely, whether competition for the best students gives schools incentives to improve.¹¹ They study the incentives for public schools to invest in quality in a city where there is school choice, so that schools compete for students, using the standard discrete Gale and Shapley framework. They show that, in large markets, the incentives for schools to invest in quality are nonnegative, but they are silent about their magnitude and about what types of investments schools pursue.¹² To address these issues, we approach the problem from a price-theoretic perspective.

Consider a city with a number of public schools $c = 1, \dots, C$, each with capacity S_c . Students and schools are matched according to a stable matching. This is a stark description of the institutional arrangements in cities where a centralized clearinghouse assigns students using a stable matching mechanism. Students are denoted as i , in a set of students I . Note that I is a set of students, distinct from the set Θ of student types. Schools' preferences over students are given by scores e_c^i . We assume that the vectors e^i are distributed according to a distribution function $G(\cdot)$ in $[0, 1]^C$, with a continuous density $g > 0$.

Students' preferences depend on the quality $\delta_c \in \mathbb{R}$ of each school. Quality δ_c should be interpreted as a vertical quality measure, in that all students prefer higher δ_c . However, different students may be affected differently by δ_c . If, for example, δ_c measures the quality of a school's calculus course, then students of high academic caliber, or with a focus in math, will be more sensitive to changes in δ_c . One of the advantages of

¹¹ The effect of competition on the provision of services by public schools, and local government services in general, is a classic topic in the economics of education and the public sector. Tiebout (1956) has pointed out that competition between locations allows agents to sort efficiently into places that offer the level of public services they prefer. With respect to schools, the literature has emphasized the importance of competition and choice to private efficiency of allocations and spillover effects (see Hoxby [2000] and references therein). More closely related to Hatfield et al. (2016) are papers that consider whether competition gives school administrators incentives to perform better, such as Hoxby's (1999) model of moral hazard in which families may move and school districts compete for resources.

¹² In contrast, Hatfield et al. (2016) provide sharp results comparing different school choice mechanisms, which we do not pursue as the present paper deals exclusively with stable matchings.

our approach is that it predicts which groups of students a school would like to target with improvements in quality. Student i has utility $u_i^c(\delta_c) > 0$ of attending school c , increasing in δ_c , and utility 0 of being unmatched. The measure of students who are indifferent between two schools is 0, for any value of the vector δ . Given δ , preferences induce a distribution η_δ over student types Θ , which we assume to have a density $f_\delta > 0$, smooth in δ and θ .

Under these assumptions, given δ , there exists a unique stable matching μ_δ . Let $P^*(\delta)$ be the unique associated market clearing cutoffs. Dependence on δ will be omitted when there is no risk of confusion.

For concreteness, define the aggregate quality of a school's entering class as

$$Q_c(\delta) = \int_{\mu_\delta(c)} e_c^\theta d\eta_\delta(\theta),$$

that is, the integral of scores e_c^θ over all students matched to the school. We consider how a school's quality δ_c affects the quality of its entering class Q_c . The motivation is that, following Hatfield et al. (2016), if schools are concerned about Q_c , then a direct link between δ_c and Q_c gives school administrators incentives to improve quality δ_c . Note that, because we are not performing an equilibrium analysis, it is not assumed that the quality of the entering class Q_c is the sole objective of a school. Even if schools have complex objectives, the impact of δ_c on Q_c isolates how investments benefit schools purely on the dimension of competing for a strong entering class, that is, the selection benefits of investment, driven by schools competing for students as opposed to being assigned a fixed entering class.

The effect of a school investing in quality can be written in terms of average characteristics of students who are marginally choosing or being chosen by schools, much as the effect of a demand shift in markets with prices is a function of characteristics of marginal consumers (Spence 1975; Veiga and Weyl 2016). To highlight the intuition behind the effect of investing in quality, we define the following quantities.

- The number N_c of additional students attracted to school c by a marginal increase in quality:

$$N_c \equiv \frac{d}{d\delta_c} D_c(P) \Big|_{P=P^*(\delta)} = \int_{\{\theta: D^\theta(P^*(\delta))=c\}} \frac{d}{d\delta_c} f_\delta(\theta) d\theta.$$

- The average quality of the attracted students:

$$\bar{e}_c \equiv \frac{1}{N_c} \int_{\{\theta: D^\theta(P^*(\delta))=c\}} e_c^\theta \cdot \frac{d}{d\delta_c} f_\delta(\theta) d\theta.$$

- The set of students who are marginally accepted to school c' and would go to school c otherwise:

$$\begin{aligned} \tilde{M}_{c'c} &\equiv \{\theta : c' \succ^\theta c, P_{c'} = e_{c'}^\theta, P_c \leq e_c^\theta, \\ &P_{c'} > e_{c'}^\theta \ \forall c'' \neq c' : c' \succ^\theta c\}. \end{aligned}$$

- The $C - 1$ -dimensional mass of students in this margin, and their average scores:

$$\begin{aligned} M_{c'c} &\equiv \sum_{\succ \in \mathbb{R}} \int_{(\succ, (P_{c'}, e_{-c'})) \in \tilde{M}_{c'c}} f_b(\succ, (P_{c'}, e_{-c'})) de_{-c'}, \\ \bar{P}_{c'c} &\equiv \frac{1}{M_{c'c}} \sum_{\succ \in \mathbb{R}} \int_{(\succ, (P_{c'}, e_{-c'})) \in \tilde{M}_{c'c}} e_c^{\succ, (P_{c'}, e_{-c'})} \\ &\cdot f_b(\succ, (P_{c'}, e_{-c'})) de_{-c'}. \end{aligned}$$

The effect of school quality δ_c on the quality of the entering class Q_c is as follows.

PROPOSITION 2. Assume that $P^*(\delta) > 0$ and that P is differentiable in δ_c . Then the quality of the entering class Q_c is differentiable in school quality δ_c , and its derivative can be decomposed as

$$\frac{dQ_c}{d\delta_c} = \underbrace{[\bar{e}_c - P_c^*] \cdot N_c}_{\text{Direct Effect}} - \underbrace{\sum_{c' \neq c} [\bar{P}_{c'c} - P_c^*] \cdot M_{c'c}}_{\text{Market Power Effect}} \cdot \left(-\frac{dP_{c'}^*}{d\delta_c}\right). \quad (3)$$

The direct effect term is weakly positive, always giving incentives to invest in quality. The market power terms increase (decrease) the incentives to invest in quality if an increase (decrease) in the quality of school c increases the market clearing cutoff of school c' , that is, $dP_{c'}^*(\delta)/d\delta_c > 0 (< 0)$.

The proposition states that the effect of an increase in quality can be decomposed into two terms. The direct effect is the increase in quality, with cutoffs P held fixed, due to students with $e_c \geq P_c^*$ choosing school c with higher frequency. Note that this term is proportional to $\bar{e}_c - P_c^*$. Since the total number of students that the school is matched to is fixed at S_c , the gain is only a change in composition. As the school attracts more students with average score \bar{e}_c , it must give up marginal students with scores P_c^* . The change in quality $\bar{e}_c - P_c^*$ is multiplied by N_c , the number of students who change their choices.

The market power effect measures how much the school loses because of its higher quality decreasing the equilibrium cutoffs of other schools. It is (the sum over all other schools c' of) the product of the change in cutoffs of the other school $(-dP_{c'}^*/d\delta_c)$ times the quantity of students in

the margin that change schools because of a small change in cutoffs, $M_{c'c}$, times the difference in the average quality of these students and the quality of a marginal student $\bar{P}_{c'c} - P_c^*$. The market power effect from school c' has the same sign as $dP_{c'}^*/d\delta_c$. It reduces the incentives to invest in quality if increasing δ_c reduces the selectivity of school c' . However, it can be positive in the counterintuitive case in which improving the quality of school c increases the selectivity of school c' . The latter case is possible only if $C \geq 3$.¹³ The intuition for the direction of the market power effect is that improvements in quality help if they induce competing schools to become more selective but harm in the more intuitive case in which improving quality makes other schools less selective and therefore compete more aggressively for students.

Hatfield et al.'s (2016) main result is that, in a large thick market, where each school makes up a negligible fraction of the market, the incentives to invest in quality are weakly positive. Within our framework this can be interpreted as saying that in such markets the market power term becomes small, and therefore, $dQ_c/d\delta_c \geq 0$.

Note that the decomposition of incentives in equation (3) gives conditions in which schools have muted incentives to invest in quality improvements for lower-ranked students. If δ_c is a dimension of quality such that $df_b(\theta)/d\delta_c \approx 0$ unless $e_c^\theta \leq P_c$ or $e_c^\theta \approx P_c$, then the direct effect

$$[\bar{e}_c - P_c^*] \cdot N_c \approx 0.$$

Consider the case in which the effect of the quality of school c on the cutoffs of other schools either is small or has the intuitive sign $dP_{c'}^*/d\delta_c \leq 0$. Then the small direct effect and weakly negative market power effect imply $dQ_c/d\delta_c \leq 0$. Therefore, by allowing schools to compete, school choice gives incentives to invest in improvements benefiting the best students but not the marginal accepted students. An example would be that a school has incentives to invest in a better calculus teacher and assign counselor time to advise students in applying to top colleges and, at the

¹³ To see this, write the aggregate demand function conditional on δ as $D(P, \delta)$. Then $D(P^*(\delta), \delta) = S$. By the implicit function theorem, we have $\partial_\delta P^* = -(\partial_P D)^{-1} \cdot \partial_\delta D$. If $C = 2$, solving this system implies $dP_{c'}^*/d\delta_c \leq 0$ for $c \neq c'$. With $C = 3$, an example of $dP_{c'}^*/d\delta_c > 0$ for $c \neq c'$ is given by

$$\partial_P D = \begin{pmatrix} -10 & 1 & 1 \\ 4 & -10 & 1 \\ 4 & 1 & -10 \end{pmatrix}, \quad \partial_\delta D = (10, -9, -1).$$

In this example the effect of increasing the quality of college 1 on cutoffs is $\partial_\delta P^* = (.98, -.49, .24)$, so that the cutoff of college 3 goes up with an increase in δ_1 . The intuition is that an increase in quality of college 1 takes more students from college 2 than from college 3, and the decrease in the selectivity of college 2 induces college 3 to become more selective.

same time, small or negative incentives to improve the quality of classes for lower-ranked students or invest counselor time in helping students with low grades. The logic of this result is that, since the quantity of students S_c that are matched to school c is fixed, for every student of score e_c that a school gains by improving quality, it must shed a marginal student with score P_c^* . The direct effect can be profitable only if $e_c - P_c^*$ is appreciably greater than zero. The argument is completed by the observation that the market power term is weakly negative if $dP_c^*/d\delta_c \leq 0$. Note that marginal students with scores $e_c \approx P_c^*$ are not necessarily “bad.” For an elite high school, cutoffs P_c^* are high, in the sense that a type e_c student is very desirable.¹⁴ Yet, because changes in quality shift the composition of only the entering class, it is still the case that the incentives to invest in attracting such students are small. Another way to frame this discussion is that the only scenario in which the incentives to invest in marginal students may be positive is the case in which a school does have market power, in the sense that it can affect the cutoffs of other schools, and for at least one of these other schools $dP_{c'}^*/d\delta_c > 0$.

The model displays an additional distortion. Even though quality affects $u_c^i(\delta_c)$ for all students, schools are concerned only with the impact on students who are indifferent between different schools, as equation (3) depends only on changes in f_δ . This is the familiar Spence (1975) distortion of a quality-setting monopolist. Its manifestation in our setting is that schools’ investment decisions take into account marginal but not inframarginal students.

Finally, if we assume that schools are symmetrically differentiated, it is possible to gain further intuition on the market power effect.¹⁵ If the function $f_\delta(\theta)$ is symmetric over schools and all $S_c = S_{c'}$, $\delta_c = \delta_{c'}$, then the market power term reduces to

$$-\frac{M_{c,c}}{M_{c,\emptyset} + C \cdot M_{c,c}} \cdot [\bar{P}_{c,c} - P_c^*] \cdot N_c,$$

where

$$M_{c,\emptyset} = \int_{\{\theta: e_c = P_{c'}, e_c < P_{c'} \forall c' \neq c\}} f_\delta(\theta) d\theta$$

is the $(C - 1)$ -dimensional mass of agents who are marginally accepted to school c and not accepted to any other schools. In the symmetric case, the market power effect is negative and proportional to the quality wedge $\bar{P}_{c,c} - P_c^*$ times the number of students whom school c attracts with

¹⁴ For example, Stuyvesant High School’s Scholastic Aptitude Test scores are in the 99.9th percentile in the state of New York (Abdulkadiroğlu et al. 2014).

¹⁵ We would like to thank Glen Weyl for the suggestion to consider the symmetric case.

improvements in quality N_c . Ceteris paribus, the absolute value of the market power effect grows with $M_{c,c}$, the mass of students on the margin between schools c' and c . These are the students whom school c may lose to c' if c' competes more aggressively. The absolute value of the market power effect also decreases with the number of schools C , and with the other quantities held fixed, it converges to 0 as the number of schools grows. The expression suggests conditions under which competition reduces the incentives for schools to invest in quality improvements for marginal applicants. This is the case in which a small number of schools compete for densely populated margins $\tilde{M}_{c,c}$. An example would be a city with a small number of elite schools that compete mostly with each other for the best students but are horizontally differentiated, so that many students are in the margins $\tilde{M}_{c,c}$.

This effect might help explain puzzling findings from regression discontinuity studies of elite schools. Abdulkadiroğlu et al. (2014) and Dobbie and Fryer (2014) find that marginally accepted students to the top three exam schools in Boston and New York do not attain higher SAT scores, despite the better peers and large amount of resources invested in these schools. This is consistent with the prediction that competition gives elite schools incentives to compete for the best students, but not to invest in improvements that benefit marginal students. At the same time, marginally rejected students at elite schools are likely to be among the best students in the schools they eventually go to. Therefore, the non-elite schools have incentives to tilt investments toward these students, which helps to explain the absence of a large difference in outcomes.

The analysis in this section could be extended in a number of ways. If the model specified costs for schools to invest, and possibly more complex objectives, it would be possible to derive first-order conditions for equilibrium play of schools. By specifying social welfare, the equilibrium conditions could be compared with optimization by a social planner. As the goal of this section is simply to illustrate the derivation of comparative statics in the continuum framework, in the interest of space we leave these extensions for future research and discuss related applications in the conclusion. We do note that the expression derived in proposition 2 is valid for general demand shocks in matching markets, so that the methodology can be readily applied to other markets. Moreover, this type of comparative static leads to straightforward equilibrium analysis in a market in which firms are assumed to play strategically, as illustrated in Azevedo (2014).

B. Random Economies and Convergence Rates

This section extends the convergence results to randomly generated finite economies and bounds the speed of convergence. Many mecha-

nisms used in practical market design explicitly incorporate randomly generated preferences, so that the results imply new characterizations of the asymptotics of such mechanisms, which we explore in Section IV.C.

We extend the convergence in theorem 2 to economies in which agents are randomly drawn, with types independently and identically distributed. The following proposition implies convergence of the sets of stable matchings and bounds the speed of convergence. Denote by $[x]$ the nearest integer to a real number x , rounding down in case of a tie.

PROPOSITION 3. Assume that the continuum economy $E = [\eta, S]$ admits a unique stable matching μ , associated market clearing cutoff P^* , and $\sum_c S_c < 1$. Let $F^k = [\eta^k, S^k]$ be a randomly drawn finite economy, with k students drawn independently according to η and the vector of capacity per student S^k defined as $S^k k = [S k]$. Let $\{\tilde{\mu}^k\}_{k \in \mathbb{N}}$ be a sequence of random variables, such that each $\tilde{\mu}^k$ is a stable matching of F^k . We have the following results:

1. Economy F^k converges almost surely to E , and $\tilde{\mu}^k$ converges almost surely to μ . Moreover, convergence is fast in the following sense. Assume that $D(\cdot|\eta)$ is C^1 and $\partial D(P^*)$ is invertible. Fix $\epsilon > 0$. We then have the following result:
2. There exist constants $\alpha, \beta > 0$ such that the probability that F^k has a market clearing cutoff P^k with $\|P^k - P^*\| \geq \epsilon$ is bounded by

$$\Pr\{F^k \text{ has a market clearing cutoff } P^k \text{ with } \|P^k - P^*\| > \epsilon\} \leq \alpha \cdot e^{-\beta k}.$$

3. Moreover, if η can be represented by a continuous density, let G^k be the fraction of students in economy F^k that receive a match different from that in the continuum economy E , that is, $D^\theta(P^k) \neq D^\theta(P^*)$ for some market clearing cutoff P^k of F^k . Then G^k converges to 0 almost surely, and there exist $\alpha', \beta' > 0$ such that the probability that $G^k > \epsilon$ is bounded by

$$\Pr\{G^k > \epsilon\} \leq \alpha' \cdot e^{-\beta' k}.$$

Part 1 says that the stable matchings of the randomly drawn economies converge almost surely to stable matchings of the limit approximation. This justifies using the continuum model as an approximation of the discrete model in settings in which preferences are random, such as in mechanisms that rely on tie-breaking lotteries. Parts 2 and 3 bound the speed of convergence. Given $\epsilon > 0$, the probability that market clearing cutoffs in F^k deviate from those in E by more than ϵ converges to 0 exponentially. Moreover, it guarantees that the fraction of students who may receive different matches in the continuum and finite economy is lower than ϵ with probability converging to 1 exponentially.

We performed simulations to gauge the applicability of the model in realistic market sizes. Student scores were drawn as the average of a student-specific component and a college student component, each drawn independently from the uniform distribution. As such, the correlation of e_c^b and $e_{c'}^b$ for $c \neq c'$ is $1/2$. Student preferences were drawn uniformly at random, and all colleges have equal capacity. The total number of seats per capita is $1/2$, so that half of the students are unmatched in equilibrium.

Figure 4 reports the results of 1,000 simulations for various market sizes. In each simulation we draw an economy and calculate the student-optimal stable matching and associated cutoffs. The simulations show that, even with as few as 10 seats per college, the average fraction of mismatched students G^k is not too large, around 15 percent. More interestingly, if there are at least 100 seats per college, then the error is quite small, with the average value of G^k lower than 5 percent. Moreover, this fraction does not increase substantially with the number of colleges. With 100 seats per college, the average number of misplaced students does not exceed 5 percent for any number of colleges between two and 500. Finally, the bottom panels show that, with at least 100 seats per college, realized cutoffs are close to the continuum cutoffs with very high probability and that their mean is virtually identical to the continuum cutoffs. Online appendix G reports results on the asymptotic distributions of cutoffs.

C. Market Design Applications

We now apply our convergence results to market design. Since many matching and assignment mechanisms use lotteries to break ties, our results on convergence of random economies readily imply asymptotic characterizations and large market properties of these mechanisms. Specifically, we give a simple derivation of results by Che and Kojima (2010) for the canonical random serial dictatorship mechanism and generalize them with novel results for a state-of-the-art mechanism used in real school choice systems.

1. The Random Serial Dictatorship Mechanism

The assignment problem consists in allocating indivisible objects to a set of agents. No transfers of a numeraire or any other commodity are possible. The most well-known solution to the assignment problem is the random serial dictatorship (RSD) mechanism. In the RSD mechanism, agents are first ordered randomly by a lottery. They then take turns picking their favorite object out of the ones that are left. Recently, Che and Kojima (2010) have characterized the asymptotic limit of the RSD

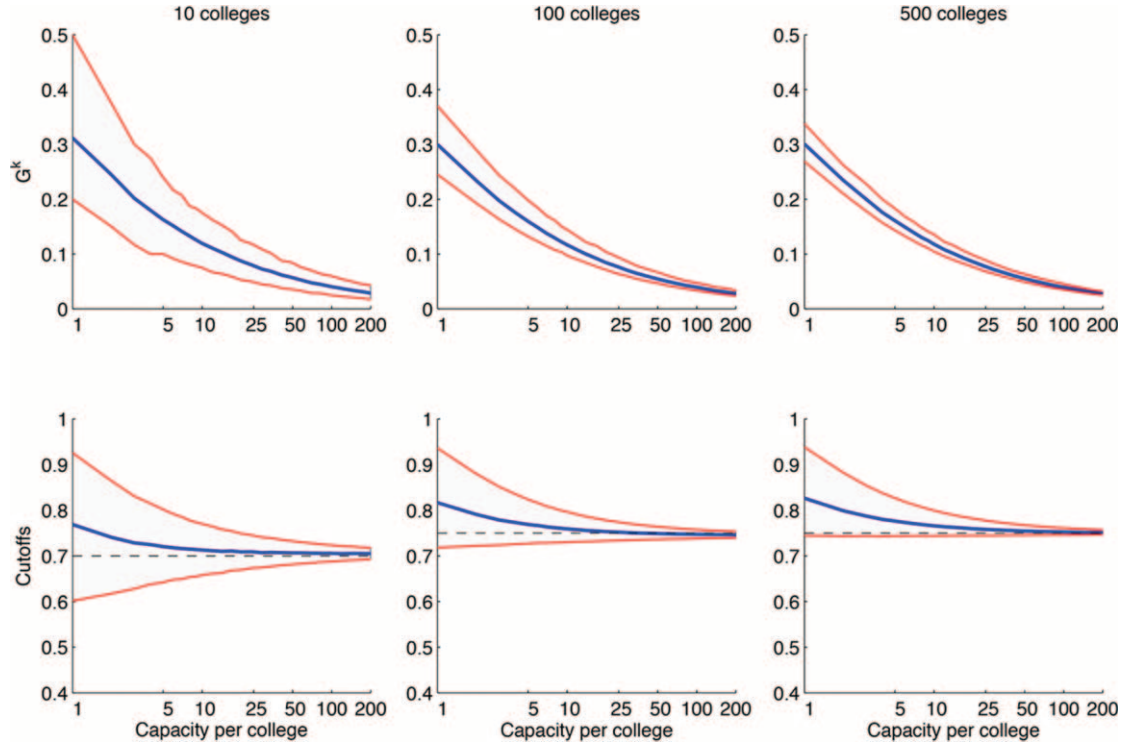


FIG. 4.—Speed of convergence. The figure depicts statistics of student-optimal stable matchings in 1,000 simulations for each market size, with the distribution of preferences as in the text. The top panels display the fraction G^k of students who receive different matches in the discrete economy, with lines corresponding to the mean, 5th, and 95th percentiles across simulations. The bottom panels report cutoffs, with lines representing the mean, 5th, and 95th percentiles, across all colleges and simulations. The dashed line represents the continuum cutoffs. Capacity per college, in the horizontal axes, is depicted in a log scale.

mechanism.¹⁶ In their model, the number of object types is fixed, and the number of agents and copies of each object grows. Their main result is that RSD is asymptotically equivalent to the probabilistic serial mechanism proposed by Bogomolnaia and Moulin (2001). This is a particular case of our results, as the serial dictatorship mechanism is equivalent to deferred acceptance when all colleges have the same ranking over students. This section formalizes this point.

In the assignment problem there are C object types $c = 1, 2, \dots, C$. An instance of the assignment problem is given by $AP = (k, m, S)$, where k is the number of agents, m is a vector with m_{\succ} representing the fraction of agents with preferences \succ for each $\succ \in \mathcal{R}$, and S is a vector with S_c being the number of copies of object c available per capita. An allocation specifies for each agent $i \in \{1, 2, \dots, k\}$ a probability $x_c^i(AP)$ of receiving each object c . Because we will consider only allocations that treat ex ante equal agents equally, we denote by $x_c^{\succ}(AP)$ the probability of an agent with preferences \succ receiving object c , for all preferences \succ present in the assignment problem.

We can describe RSD as a particular case of the deferred acceptance mechanism in which all colleges have the same preferences. First, we give agents priorities based on a lottery l , generating a random finite college admissions problem $F(AP, l)$, where agents correspond to students and colleges to objects. Formally, given assignment problem AP , randomly assign each agent i a single lottery number l^i uniformly in $[0, 1]$ that gives her a score in all colleges (i.e., objects) of $e_c^i = l$. Associate with this agent a student type $\theta_i = (\succ^i, e^i)$. This induces a random discrete economy $F(AP, l)$ as in the previous section. That is, as l is a random variable, $F(AP, l)$ is a random finite economy, and for particular draws of l it is a finite economy. For almost every draw of the economy $F(AP, l)$ has strict preferences. Each agent i 's allocation $x_c^i(AP)$ under RSD is then equal to the probability of receiving object c in her allocation in $F(AP, l)$.

Consider now a sequence of finite assignment problems $\{AP^k\}_{k \in \mathbb{N}}$, $AP^k = (k, m^k, S^k)_{k \in \mathbb{N}}$. Assume that (m^k, S^k) converges to some (m, S) with $S > 0$, $m > 0$. Let each l^k be a lottery consisting of k draws, one for each agent, uniformly distributed in $[0, 1]^k$. For each k , the assignment problem and the lottery induce a random economy $F(AP^k, l^k)$.

Note that the finite economies $F(AP^k, l^k)$ converge almost surely to a continuum economy E with a vector S of quotas, a mass m_{\succ} of agents with each preference list \succ , and scores e^{θ} uniformly distributed along the diagonal of $[0, 1]^C$. This continuum economy has a unique market clear-

¹⁶ This asymptotic characterization has been generalized by Liu and Pycia (2013) to any uniform randomization over Pareto-efficient mechanisms under an equicontinuity condition.

ing cutoff $P(m, S)$. By proposition 3, cutoffs in large finite economies are converging almost surely to $P(m, S)$. We have the following characterization of the limit of the RSD mechanism.

PROPOSITION 4. Under the RSD mechanism the probability $x_c^\succ(AP^h)$ that an agent with preferences \succ will receive object c converges to

$$\int_{l \in [0,1]} \mathbf{1}_{(c = \operatorname{argmax}_{\{c \in C | P_c(m,S) \leq l\})} dl .$$

That is, the cutoffs of the continuum economy describe the limit allocation of the RSD mechanism. In the limit, agents are given a lottery number uniformly drawn between 0 and 1 and receive their favorite object out of the ones with cutoffs $P_c(m, S)$ below the lottery number l . Note that this implies the Che and Kojima (2010) result on the asymptotic equivalence of RSD and the probabilistic serial mechanism. Namely, in the probabilistic serial mechanism, agents simultaneously eat probability shares of their favorite object that is still available. The market clearing equations imply that object c runs out at time $1 - P_c(m, s)$. Hence, the probability that an agent with preferences \succ is assigned her favorite object c_1 is $1 - P_{c_1}(m, s)$. The probability of being assigned her second choice c_2 or better equals $\max_{c \geq c_2} 1 - P_c(m, s)$, and so on.

2. School Choice Mechanisms

We now derive new results for deferred acceptance mechanisms used by actual clearinghouses that allocate seats in public schools to students. These mechanisms generalize RSD, as in school choice some students are given priority to certain schools.

The school choice problem consists of assigning seats in public schools to students while observing priorities some students may have to certain schools. It differs from the assignment problem because schools give priorities to subsets of students. It differs from the classic college admissions problem in that often schools are indifferent between large sets of students (Abdulkadiroğlu and Sönmez 2003). For example, a school may give priority to students living within its walking zone but treat all students within a priority class equally. In Boston and New York City, the clearinghouses that assign seats in public schools to students were recently re-designed by academic economists (Abdulkadiroğlu, Pathak, and Roth 2005; Abdulkadiroğlu et al. 2005). The chosen mechanism was deferred acceptance with single tie breaking (DA-STB). DA-STB first orders all students using a single lottery, which is used to break indifferences in the schools' priorities, generating a college admissions problem with strict preferences. It then runs the student-proposing deferred acceptance al-

gorithm given the refined preferences (Abdulkadiroğlu et al. 2009; Kesten and Ünver 2015).

We can use our framework to derive the asymptotics of the DA-STB mechanism. Fix a set of schools $C = \{1, \dots, C\}$. Students are described as $i = (\succ^i, e^i)$ given by a strict preference list \succ^i and a vector of scores e^i . However, to incorporate the idea that schools have only very coarse priorities, corresponding to a small number of priority classes, we assume that all e_c^i are integers in $\{0, 1, 2, \dots, \bar{e}\}$ for $\bar{e} \geq 0$. Therefore, the set of possible student types is finite. We denote by $\bar{\Theta}$ the set of possible types. A school choice problem $SC = (k, m, S)$ is given by a number of students k , a fraction m_i of students of each of the finite types $\bar{\theta} \in \bar{\Theta}$, and a vector of capacity per capita of each school S .

We can describe the DA-STB mechanism as first breaking indifferences through a lottery l , which generates a finite college admissions model $F(SC, l)$, and then giving each student the student-proposing deferred acceptance allocation. Assume that each student $i \in \{0, 1, \dots, k\}$ receives a lottery number l^i independently uniformly distributed in $[0, 1]$. The student's refined score in each school is given by her priority, given by her type, plus lottery number, $e_c^{\theta_i} = e_c^i + l_i$. The refined type is defined as $\theta_i = (\succ^i, e^{\theta_i})$. Therefore, the lottery yields a randomly generated finite economy $F(SC, l)$, as defined in proposition 3. The DA-STB mechanism then assigns each student i in F to her match in the unique student-optimal stable matching. For each type $\bar{\theta} \in \bar{\Theta}$ in the school choice problem, denote by $x_c^{\bar{\theta}}(SC|DA)$ the probability that a student with type $\bar{\theta}$ receives school c if type $\bar{\theta}$ is present in the economy.

Consider now a sequence of school choice problems $SC^k = (k, m^k, S^k)$, each with $k \rightarrow \infty$ students. Problem k has a fraction $m_{\bar{\theta}}^k$ of students of each type, and school c has capacity S_c^k per student. Assume that (m^k, S^k) converges to some (m, S) with $S > 0, m > 0$.

Analogously to the assignment problem, as the number of agents grows, the aggregate randomness generated by the lottery disappears. The randomly generated economies $F(SC^k, l^k)$ are converging almost surely to a continuum economy, given as follows. For each of the possible types $\bar{\theta} \in \bar{\Theta}$, let the measure $\eta_{\bar{\theta}}$ over $\bar{\Theta}$ be uniformly distributed in the line segment $\succ^{\bar{\theta}} \times [e^{\bar{\theta}}, e^{\bar{\theta}} + (1, 1, \dots, 1)]$, with total mass 1. Let $\eta = \sum_{\bar{\theta} \in \bar{\Theta}} m_{\bar{\theta}} \cdot \eta_{\bar{\theta}}$. The limit continuum economy is given by $E = [\eta, S]$. We have the following generalization of the result in the previous section.

PROPOSITION 5. Assume that the continuum economy E has a unique market clearing cutoff $P(m, S)$. Then the probability $x_c^{\bar{\theta}}(SC^k|DA)$ that DA-STB assigns a student with type $\bar{\theta} \in \bar{\Theta}$ to school c converges to

$$\int_{l \in [0,1]} \mathbf{1}_{(c = \operatorname{argmax}_{\{c \in C | P_c(m,S) \leq e^{\bar{\theta}} + l\})} dl .$$

Moreover, the realized fraction of agents of type $\bar{\theta}$ that are assigned to school c converges almost surely to this value.

The proposition says that the asymptotic limit of the DA-STB allocation can be described using cutoffs. The intuition is that, after breaking ties, a discrete economy with a large number of students is very similar to a continuum economy in which students have lottery numbers uniformly distributed in $[0, 1]$. The main limitation of the proposition is that it requires the continuum economy to have a unique market clearing cutoff. Although we know that this is valid for generic vectors of capacities S , online appendix F shows that it is not always the case.

The second part of the result shows that as the market grows the aggregate randomness of the DA-STB mechanism disappears. Although the mechanism depends on random lottery draws, the fraction of agents with the same priority and preferences going to each school converges almost surely to that in the asymptotic limit. Therefore, while the allocation of an individual agent depends on the lottery, the aggregate allocation is unlikely to change with different draws. This limit result is consistent with data from the New York City match. Abdulkadiroğlu et al. (2009) report that in multiple runs of the algorithm the average number of applicants assigned to their first choice is 32,105.3, with a standard deviation of only 62.2. The proposition predicts that in increasingly large markets, with a similar distribution of preferences and seats per capita, this standard deviation divided by the total number of students (in the New York City data $62.2/78,728 \approx 0.0008$) converges to 0.

The proposition has an important consequence for the efficiency of DA-STB. Che and Kojima (2010) show that, while the RSD mechanism is ordinally inefficient, the magnitude of this inefficiency goes to 0 as the number of agents grows. Similarly, the DA-STB mechanism is ex post inefficient, having a positive probability of its outcome being Pareto dominated by other stable matchings (Erdil and Ergin 2008). In contrast to RSD, this inefficiency does not go away in a large market. We give an example in online appendix F in which the probability that DA-STB produces a Pareto-dominated outcome converges to 1 as the market grows.

Finally, the proposition generalizes the result in the previous section that describes the asymptotic limit of the RSD mechanism. RSD corresponds to DA-STB in the case in which all students have equal priorities. Therefore, the market clearing equations provide a unified way to understand asymptotics of RSD, the probabilistic serial mechanism, and DA-STB. Moreover, one could easily consider other ways in which the lottery l is drawn and derive asymptotics of other mechanisms, such as deferred acceptance with multiple tie breaking discussed by Abdulkadiroğlu et al. (2009).

V. Conclusion

This paper applies standard supply and demand analysis to matching markets. Although the supply and demand characterization is valid in both discrete and continuum cases, the paper focuses on the model with a continuum of agents on one side of the market. This approach permits simple derivation of comparative statics and of large-market results. We now highlight four points that were not addressed in the analysis.

First, analysis of matching markets has typically taken one of two polar perspectives: either focusing on assortative matching, where rich comparative statics can be derived, or using models based on Gale and Shapley, where a limited set of such results are possible. We view our approach as a middle ground that complements these analyses. This is illustrated by our school competition application, where we derive key comparative statics as functions of the distribution of preferences and competitive structure in a market.

Second, the tractability of the continuum model may be useful in applications, such as Azevedo's (2014) analysis of imperfect competition and Agarwal and Somaini's (2014) derivation of identification results and estimators for preferences in matching mechanisms. It is, of course, important to determine whether the continuum approximation is appropriate in any given application.

Third, the matching literature has explored frameworks that are more general than the Gale and Shapley model. In these models, the existence of stable matchings depends on restrictions on preferences, such as substitutability. It would be interesting to understand to what extent the continuum assumption obviates the need for such restrictions, as in Che, Kim, and Kojima (2013) and Azevedo and Hatfield (2015).¹⁷

The common theme in our analysis is the application of basic ideas from competitive equilibrium to matching markets. We hope that this underlying idea will prove useful in the analysis of other problems and broaden the applicability of Gale and Shapley's (1962) notion of stability, yielding insights in specific markets in which Becker's (1973) assumptions of vertical preferences and assortative matching do not hold.

References

- Abdulkadiroğlu, A., J. D. Angrist, and P. A. Pathak. 2014. "The Elite Illusion: Achievement Effects at Boston and New York Exam Schools." *Econometrica* 82 (1): 137–96.
- Abdulkadiroğlu, A., Y. K. Che, and Y. Yasuda. 2015. "Expanding 'Choice' in School Choice." *American Econ. J.: Microeconomics* 7 (1): 1–42.

¹⁷ See Hatfield and Milgrom (2005), Hatfield and Kojima (2008, 2010), and Hatfield and Kominers (2012, 2016) for existence and maximal domain results.

- Abdulkadiroğlu, A., P. A. Pathak, and A. E. Roth. 2005. "The New York City High School Match." *A.E.R.* 95 (2): 364–67.
- . 2009. "Strategy-Proofness versus Efficiency in Matching with Indifferences: Redesigning the NYC High School Match." *A.E.R.* 99 (5): 1954–78.
- Abdulkadiroğlu, A., P. A. Pathak, A. E. Roth, and T. Sönmez. 2005. "The Boston Public School Match." *A.E.R.* 95 (2): 368–71.
- Abdulkadiroğlu, A., and T. Sönmez. 2003. "School Choice: A Mechanism Design Approach." *A.E.R.* 93 (3): 729–47.
- Adachi, H. 2000. "On a Characterization of Stable Matchings." *Econ. Letters* 68 (1): 43–49.
- Agarwal, Nikhil, and Paulo Somaini. 2014. "Demand Analysis Using Strategic Reports: An Application to a School Choice Mechanism." Working paper, NBER, Cambridge, MA.
- Aumann, R. J. 1964. "Markets with a Continuum of Traders." *Econometrica* 32 (1–2): 39–50.
- Azevedo, Eduardo M. 2014. "Imperfect Competition in Two-Sided Matching Markets." *Games and Econ. Behavior* 83:207–23.
- Azevedo, E. M., and J. W. Hatfield. 2015. "Existence of Equilibrium in Large Matching Markets with Complementarities." Manuscript, Wharton School, Univ. Pennsylvania.
- Azevedo, E. M., and J. D. Leshno. 2013. "A Supply and Demand Framework for Two-Sided Matching Markets." Manuscript, Wharton School, Univ. Pennsylvania.
- Balinski, M., and T. Sönmez. 1999. "A Tale of Two Mechanisms: Student Placement." *J. Econ. Theory* 84 (1): 73–94.
- Becker, G. S. 1973. "A Theory of Marriage: Part I." *J.P.E.* 81:813–46.
- Biró, P. 2007. "Higher Education Admission in Hungary by a Score-Limit Algorithm." Paper presented at the 18th International Conference on Game Theory, Stony Brook, NY. <http://old.gtcenter.org/Archive/Conf07/Downloads/Conf/Biro453.pdf>.
- Bodoh-Creed, Aaron, and Brent R. Hickman. 2015. "College Assignment as a Large Contest." Manuscript, Univ. Chicago.
- Bogomolnaia, A., and H. Moulin. 2001. "A New Solution to the Random Assignment Problem." *J. Econ. Theory* 100 (2): 295–328.
- Bojilov, Raicho, and Alfred Galichon. 2016. "Matching in Closed-Form: Equilibrium, Identification, and Comparative Statics." *Econ. Theory* 61 (4): 587–609.
- Chade, H., G. Lewis, and L. Smith. 2014. "Student Portfolios and the College Admissions Problem." *Rev. Econ. Studies* 81 (3): 971–1002.
- Che, Yeon-Koo, Jinwoo Kim, and Fuhito Kojima. 2013. "Stable Matching in Large Economies." Manuscript, Stanford Univ.
- Che, Yeon-Koo, and Fuhito Kojima. 2010. "Asymptotic Equivalence of Probabilistic Serial and Random Priority Mechanisms." *Econometrica* 78 (5): 1625–72.
- Chiappori, P.-A., M. Iyigun, and Y. Weiss. 2009. "Investment in Schooling and the Marriage Market." *A.E.R.* 99 (5): 1689–1713.
- Coles, P., A. Kushnir, and M. Niederle. 2013. "Preference Signaling in Matching Markets." *American Econ. J.: Microeconomics* 5 (2): 99–134.
- Debreu, G. 1970. "Economies with a Finite Set of Equilibria." *Econometrica* 38:387–92.
- Dobbie, W., and R. G. Fryer. 2014. "The Impact of Attending a School with High-Achieving Peers: Evidence from the New York City Exam Schools." *American Econ. J.: Appl. Econ.* 6 (3): 58–75.
- Echenique, Federico. 2012. "Contracts vs. Salaries in Matching." *A.E.R.* 102:594–601.

- Echenique, Federico, and J. Oviedo. 2004. "Core Many-to-One Matchings by Fixed-Point Methods." *J. Econ. Theory* 115 (2): 358–76.
- . 2006. "A Theory of Stability in Many-to-Many Matching Markets." *Theoretical Econ.* 1 (2): 233–73.
- Echenique, Federico, and M. Bumin Yenmez. 2007. "A Solution to Matching with Preferences over Colleagues." *Games and Econ. Behavior* 59 (1): 46–71.
- Erdil, A., and H. Ergin. 2008. "What's the Matter with Tie-Breaking? Improving Efficiency in School Choice." *A.E.R.* 98 (3): 669–89.
- Fleiner, Tamás, and Zsuzsanna Jankó. 2014. "Choice Function–Based Two-Sided Markets: Stability, Lattice Property, Path Independence and Algorithms." *Algorithms* 7 (1): 32–59.
- Gabaix, X., and A. Landier. 2008. "Why Has CEO Pay Increased So Much?" *Q. J.E.* 123 (1): 49–100.
- Gale, David, and Lloyd S. Shapley. 1962. "College Admissions and the Stability of Marriage." *American Math. Monthly* 69 (1): 9–15.
- Gale, David, and Marilda Sotomayor. 1985. "Some Remarks on the Stable Matching Problem." *Discrete Appl. Math.* 11 (3): 223–32.
- Galichon, Alfred, Scott Duke Kominers, and Simon Weber. 2016. "Costly Concessions: An Empirical Framework for Matching with Imperfectly Transferable Utility." Working paper, Soc. Sci. Res. Network. <http://dx.doi.org/10.2139/ssrn.2535394>.
- Grossman, G. M. 2004. "The Distribution of Talent and the Pattern and Consequences of International Trade." *J.P.E.* 112:209–39.
- Hatfield, John William, and Fuhito Kojima. 2008. "Matching with Contracts: Comment." *A.E.R.* 98 (3): 1189–94.
- . 2010. "Substitutes and Stability for Matching with Contracts." *J. Econ. Theory* 145 (5): 1704–23.
- Hatfield, John William, Fuhito Kojima, and Yusuke Narita. 2016. "Improving Schools through School Choice: A Market Design Approach." *J. Econ. Theory*, electronically published July 20.
- Hatfield, John William, and Scott Duke Kominers. 2012. "Matching in Networks with Bilateral Contracts." *American Econ. J.: Microeconomics* 4 (1): 176–208.
- . 2016. "Contract Design and Stability in Many-to-Many Matching." *Games and Econ. Behavior*, electronically published February 12.
- Hatfield, John William, and Paul R. Milgrom. 2005. "Matching with Contracts." *A.E.R.* 95 (4): 913–35.
- Hoxby, C. M. 1999. "The Productivity of Schools and Other Local Public Goods Producers." *J. Public Econ.* 74 (1): 1–30.
- . 2000. "Does Competition among Public Schools Benefit Students and Taxpayers?" *A.E.R.* 90 (5): 1209–38.
- Immorlica, N., and M. Mahdian. 2005. "Marriage, Honesty, and Stability." In *SODA '05 Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms*. Philadelphia: Soc. Indus. and Appl. Math.
- Kesten, O., and M. U. Ünver. 2015. "A Theory of School Choice Lotteries." *Theoretical Econ.* 10 (2): 543–95.
- Kojima, F., and P. A. Pathak. 2009. "Incentives and Stability in Large Two-Sided Matching Markets." *A.E.R.* 99 (3): 608–27.
- Koopmans, Tjalling C., and Martin Beckmann. 1957. "Assignment Problems and the Location of Economic Activities." *Econometrica* 25 (1): 53–76.
- Lee, SangMok. 2014. "Incentive Compatibility of Large Centralized Matching Markets." Manuscript, Univ. Pennsylvania.

- Lee, Soohyung, and Muriel Niederle. 2015. "Propose with a Rose? Signaling in Internet Dating Markets." *Experimental Econ.* 18 (4): 731–55.
- Liu, Q., and M. Pycia. 2013. "Ordinal Efficiency, Fairness, and Incentives in Large Markets." Manuscript, Univ. California, Los Angeles.
- Ostrovsky, M. 2008. "Stability in Supply Chain Networks." *A.E.R.* 98 (3): 897–923.
- Roth, A. E. 1985. "The College Admissions Problem Is Not Equivalent to the Marriage Problem." *J. Econ. Theory* 36 (2): 277–88.
- Roth, A. E., and E. Peranson. 1999. "The Redesign of the Matching Market for American Physicians: Some Engineering Aspects of Economic Design." *A.E.R.* 89 (4): 748–80.
- Roth, A. E., and M. Sotomayor. 1989. "The College Admissions Problem Revisited." *Econometrica* 57 (3): 559–70.
- Segal, I. 2007. "The Communication Requirements of Social Choice Rules and Supporting Budget Sets." *J. Econ. Theory* 136 (1): 341–78.
- Shapley, Lloyd S., and Martin Shubik. 1971. "The Assignment Game I: The Core." *Internat. J. Game Theory* 1 (1): 111–30.
- Sönmez, T., and M. U. Ünver. 2010. "Course Bidding at Business Schools." *Internat. Econ. Rev.* 51 (1): 99–123.
- Spence, A. M. 1975. "Monopoly, Quality, and Regulation." *Bell J. Econ.* 6 (2): 417–29.
- Terviö, M. 2008. "The Difference That CEOs Make: An Assignment Model Approach." *A.E.R.* 98 (3): 642–68.
- Tiebout, C. M. 1956. "A Pure Theory of Local Expenditures." *J.P.E.* 64 (5): 416–24.
- Varian, Hal R. 1975. "A Third Remark on the Number of Equilibria of an Economy." *Econometrica* 43 (5–6): 985–86.
- Veiga, André, and E. Glen Weyl. 2016. "Product Design in Selection Markets." *Q.J.E.* 131 (2): 1007–56.