The College Admissions Problem With a Continuum of Students*

Eduardo M. Azevedo[†] and Jacob D. Leshno[‡]

January 26, 2011

Abstract

In many two-sided matching markets, agents on one side are matched to a large number of agents on the other side (e.g. college admissions). Yet little is known about the structure of stable matchings when there are many agents on one side. To approach this question we propose a variation of the Gale and Shapley (1962) college admissions model where a finite number of colleges is matched to a continuum of students. It is shown that, generically (though not always) (i) there is a unique stable matching, (ii) this stable matching varies continuously with the underlying economy, and (iii) it is the limit of the set of stable matchings of approximating large discrete economies.

Keywords: Stable matchings, large markets.

*We would like to thank Nikhil Agarwal, Susan Athey, Eric Budish, Gabriel D. Carroll, Carlos E. da Costa, Drew Fudenberg, Felipe Iachan, Scott Kominers, Mihai Manea, Alex Nichifor, Parag Pathak, Alvin Roth, Robert Townsend, Yosuke Yasuda and seminar participants at Harvard, MIT, TAU, Technion and USP for helpful comments and discussions. All mistakes are our own.

[†]Harvard University, azevedo@fas.harvard.edu.

[‡]Harvard University and Harvard Business School, jleshno@hbs.edu.

1 Introduction

In several two-sided matching markets, agents on one side are matched to a large number of agents on the other side. For example, Princeton, Harvard, Yale, Stanford, and MIT all have incoming classes with over 1,000 freshmen. Even CalTech, which has a relatively small entering class accepts around 250 freshmen yearly.[1] In some of these markets matching is decentralized.[2] One example is college admissions in the US. Another is the market for junior associates at top law firms. Most of the top American law firms hire around 50-150 associates from each cohort, mostly from the nation's most prestigious law schools.[3] Other markets also have a larger number of agents on one side, but are organized around a centralized clearinghouse, where agents report their preferences, and receive a match based on a mechanism. This is the case of public schools in several American cities, in Hungary, and of college admissions in Hungary and Turkey.[4] These markets are usually modeled using the Gale and Shapley (1962) college admissions model. Moreover, the centralized clearinghouses often employ variations of their deferred acceptance mechanism. An extensive literature considers the design and properties of these markets.[5] However, there is little work understanding matching markets with a large number of agents on one side, although this is the case in many applications.[6]

[1] Forbes ranking of "America's Best Colleges 2008".

[2] For models of two-sided matching in decentralized markets, see Adachi (2003); Niederle and Yariv (n.d.). These papers outline conditions under which decentralized matching process lead to stable allocations.

[3] See Avery et al. (2004) for details on the college admission market, and (Ginsburg and Wolf, 2003) for a description of the American and Canadian markets for junior law associates.

[4] A discussion of school choice mechanisms used in various cities is given in the seminal work of Abdulkadiroglu and Sonmez (2003), which introduced the problem of designing school choice mechanisms in the literature. Accounts of the redesign of the matching systems in Boston and NYC are given by Abdulkadiroglu et al. (2005a,b). College admissions in Turkey are described by Balinski and Sonmez (1999). Biró (2007) describes the centralized clearinghouses in Hungary.

[5] See Roth (2008) for a survey.

[6] Some interesting papers have investigated strategic properties of stable mechanisms in markets where the number of agents on both sides grows. The conclusion typically is that, as agents become insignificant, stable mechanisms become approximately strategy-proof (Roth and Peranson (1999); Immorlica and Mahdian (2005); Kojima and Pathak (2009)). This is different from the direction we pursue, in which the number of colleges is fixed, and its the number of students and the quotas of each college that grow. Our model is more similar in spirit with a literature on asymptotics of the assignment problem where the number of object types remains constant and the market grows (Che and Kojima (Forthcoming); Kojima and Manea (2009); Manea (2009)), and on large markets in the course allocation problem, where the number of courses is fixed (Budish and Cantillon (Forthcoming); Budish (2008)).

In this paper, we propose a variation of the Gale and Shapley (1962) college admissions model, where a finite number of colleges is matched to a continuum of students. Although we use the colleges and students terminology, the model can represent other matching markets, and we extend it to allow for matching with contracts. Our model allows for tractable analysis of markets where agents on one side are matched to a large number of agents on the other side. Our main results are as follows. Generically (though not always), (i) the continuum model admits a unique stable matching, (ii) this stable matching varies continuously with the underlying economy, and (iii) it is the limit of the set of stable matchings of approximating discrete economies. These results provide foundations to continuum matching models considered in the literature ((Abdulkadiroglu et al., 2008; Miralles, 2008)), imply new results on the size of the set of stable matchings in discrete models (complementing those in (Roth and Peranson, 1999; Immorlica and Mahdian, 2005; Kojima and Pathak, 2009)), and generalizes characterizations of the asymptotic behavior of commonly used mechanisms ((Che and Kojima, Forthcoming)). Besides contributing to understand markets with a large number of agents in one side, the tractability of the model makes it useful in exploring problems which are too complex in the discrete setting.

To fix ideas, say colleges preferences rank students according a number, which we term the score. Different colleges may rank students differently. A unifying idea in our analysis is considering the *score of the marginal student* accepted to each college, in a given stable matching. We denote the score of a marginal student accepted as the *cutoff*[7] at each college. This means students with scores above the cutoff are accepted, and those with lower scores are rejected. We offer a new lemma, in both the discrete and continuum models, that shows that stable matchings are associated with cutoffs that clear the market.[8] That is, such that when each student points to her favorite college that would accept her, demand for colleges equals supply.[9] Since cutoffs characterize stable matchings in both the continuum and discrete model, this Lemma is the key idea linking continuum and discrete economies. This gives, first, a

[7]This term was introduced by Abdulkadiroglu et al. (2008), who consider the case where all colleges have the same preferences.

[8]As we detail below, this was observed by Biró (2007). Yet, the particular bijection between market clearing cutoffs and stable matchings given in our lemma was is new, as are the version with a continuum of students and of matching with contracts. As discussed below, this is also related to an important result by (Roth and Sotomayor, 1989), although the results are independent.

[9]More precisely, a set of cutoffs clears the market if the demand for each school does not exceed its quota, and equals the quota if the cutoff is strictly positive.

tractable characterization of stable matchings in the continuum model. And, second, allows us to prove convergence results without relying on combinatorial arguments. Therefore the arguments used to establish our limit results differ markedly from those used in other papers that considered large markets in matching and the assignment problem (Immorlica and Mahdian (2005); Che and Kojima (Forthcoming); Manea (2009); Kojima and Manea (2009); Kojima and Pathak (2009)).

Albeit very simple, the Lemmas relating stable matchings to cutoffs are of independent interest, and among our main results. In the discrete case, the cutoff Lemma can be described informally as follows. Given a stable matching, we can define admission thresholds at each college such that, if each student points to her favorite college that would accept her, the result is the original stable matching. Moreover, the lemma implies that any vector of thresholds that clears the market induces a stable matching.

The model has implications to several strands of the matching literature. We show that a generic continuum economy has a unique stable matching, which is the limit of the sets of stable matchings of any sequence of approximating discrete economies. Therefore, large discrete economies with many agents on one side may have several stable matchings, but they will often be very similar. This complements results by (Immorlica and Mahdian, 2005; Kojima and Pathak, 2009) who give conditions under which the set of stable matchings of large discrete economies is small, and seems to be consistent with data from the redesign of the National Resident Matching Program (NRMP) (Roth and Peranson (1999)).[10]

Another important implication for empirical work and simulations is that we should expect the set of stable matchings in actual markets to be robust with respect to small perturbations of the economy. This is important, in light of examples we give in the text where the set of stable matchings can change discontinuously with respect to small perturbations of the economy. Even though such cases do exist, they only arise for a measure 0 set of economies, and therefore are not likely to arise in empirical settings. This is important if data and simulations are to be used to evaluate the impact of alternative mechanisms.[11] If stable matching mechanisms were very sensitive with respect to the underlying economy, these exercises would have little value.

[10] While this is interesting, the number of doctors hired by each hospital is small, rendering the continuum model a very coarse approximation.

[11] Budish and Cantillon (Forthcoming) for example use data from the Harvard Business School course allocation mechanism to evaluate different mechanisms.

The continuity result is consistent with empirical results reported by Abdulkadiroglu et al. (2009). They consider preference data from the New York City school choice mechanism. Students are given priorities to schools based on some criteria, such as the area where they live and where their siblings go to school, and ties are broken using a lottery. Seats are then assigned according to the student-proposing deferred acceptance mechanism. Interestingly, in several different runs of the algorithm, many aggregate statistics of the match do not vary much. For example, on average 32,105.3 students receive their first choice, with a standard deviation of only 62.2. The average number of students receiving their 7th choice is 1,732.7 with a standard deviation of 26.0. It seems remarkable at first that aggregate statistics of the match are so stable, as the allocation depends on the results of a lottery. However, this is consistent with the fact that, for a typical draw, the economy after tie-breaking does not vary too much, and the result that stable matchings generically depend continuously on the primitives.

The convergence results give foundations to some interesting recent work that applies continuum models to school choice problems (Abdulkadiroglu et al. (2008); Miralles (2008)). These papers have considered the particular case of our model where all colleges have the same preferences over students. Miralles (2008) uses the continuum model to compare deferred acceptance with the Boston Mechanism. Abdulkadiroglu et al. (2008) evaluate mechanisms where agents can express the intensity of their preferences. Our results show that, generically, the stable matchings in these continuum models correspond to limits of discrete economies. In addition, we generalize the models to encompass the case where school preferences are not the same.

Our results also imply a characterization of the limit of deferred acceptance mechanisms. In particular, it includes as a special case the state of the art mechanism used in school choice, which is deferred acceptance where ties are broken according to a single lottery (DA-STB). In a related paper, Che and Kojima (Forthcoming) consider the limit of the widely used random serial dictatorship mechanism. They show that, in the limit, it corresponds to the probabilistic serial mechanism proposed by Bogomolnaia and Moulin (2001). Because serial dictatorship is equivalent to deferred acceptance in the particular case where all colleges have the same preferences, their result is also a particular case of ours. Therefore, our model gives a unified description of the limit behavior or random serial dictatorship, deferred acceptance, and the probabilistic serial mechanism.

Finally, we pursue additional applications of the model in two companion papers. Azevedo and Leshno (2010) evaluate the equilibrium performance of the stable improvement cycles mechanism, proposed by Erdil and Ergin (2008). Azevedo (2010) investigates strategic behavior of firms in matching markets.[12]

Section 2 presents the model, some preliminary results, and gives an example illustrating the results. Section 3 describes the main results, and Section 4 concludes. The appendix provides all omitted proofs, and also covers additional results on matching with a continuum of students, asymptotics of commonly used mechanisms, and extends the continuum model to matching with contracts.

2 Model

2.1 College admissions with a continuum of students

The model follows closely the Gale and Shapley (1962) college admissions problem. The main departure is that a finite number of colleges $C = \{1, 2, \dots, n\}$ is matched to a continuum mass of students. A student is described by $\theta = (\succ^\theta, e^\theta)$. $\succ^\theta$ is the student's strict preference ordering over colleges. The vector $e^\theta \in [0,1]^n$ describes the colleges' ordinal preferences for the student. We refer to e^θ_s as student θ's score or rank at college s. Colleges prefer students with higher scores. That is, college c prefers[13] student θ over θ' if $e^\theta_c > e^{\theta'}_c$. To simplify notation we assume that all students and colleges are acceptable. Let $\mathcal{S}$ be the set of all strict preference orderings over colleges. We denote the set of all student types by $\Theta = \mathcal{S} \times [0,1]^n$.

A continuum economy is given by $E = [\eta, q]$, where η is a probability measure[14]

[12] Since Roth (1985) it has been known that no stable matching mechanism is strategyproof for the colleges in the college admissions model. This is in contrast to the marriage model, where the men have no incentives to manipulate the men-optimal stable mechanism. Sonmez (1997) has shown that they may always gain by manipulating reported capacity. Konishi and Unver (2006) have then introduced games of capacity manipulation, which were also studied by Ehlers (2010); Kesten (2008); Kojima (2006); Mumcu and Saglam (2009); Romero-Medina and Triossi (2007). Azevedo (2010) also focuses on quantity manipulations, and uses the continuum model to derive equilibrium predictions in matching markets, with firms acting strategically.

[13] We take college's preferences over students as primitives, rather than preferences over sets of students. It would have been equivalent to start with preferences over sets of students that were responsive to the preferences over students, as in Roth (1985).

[14] We must also specify a σ-algebra where η is defined. The set Θ is the product of $[0,1]^n$ and the finite set of all possible orderings. We take the Borel σ-algebra of the product topology (the normal topology for $\mathbb{R}^n$ times the discrete topology for the set of orderings) .

over Θ and $q = (q_1, q_2, \ldots, q_n)$ is a vector of strictly positive capacities for each college. 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.[15]

The set of all economies satisfying Assumption 1 is denoted by $\mathcal{E}$.

A **matching** for a continuum economy $E = [\eta, q]$ is a function $\mu : C \cup \Theta \to 2^{\Theta} \cup C$, such that[16]

1. Each student is matched to a college or to herself.

2. Each college c is matched to a subset of students of measure of at most q_c.

3. A college is matched to a student iff the student is matched to the college.

4. The matching is right-continuous.[17]

This is the standard definition, with the addition of the last technical requirement, which eliminates multiplicities of matchings that coincide in a measure 0 set. 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 that has a stricly lower score than θ.[18]

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

[15] That is, for any college c and real number x we have $\eta(\{\theta | e^{\theta}_c = x\}) = 0$.

[16] Mathematically, these properties are:

1. For all $\theta \in \Theta$: $\mu(\theta) \in C \cup \{\theta\}$.

2. For all $c \in C$: $\mu(c) \subset \Theta$, and $\eta(\mu(c)) \leq q_c$.

3. $c = \mu(\theta)$ iff $\theta \in \mu(c)$.

4. For any sequence of students $\theta^k = (\succ, e^k)$, with e^k converging to e, and all $e^k \geq e$ (in every coordinate), we can find some large K so that $\mu(\theta^k) = \mu(\theta)$ for $k > K$.

[17] See the previous footnote for a precise definition.

[18] That is, (θ, c) blocks μ if $c \succ^{\theta} \mu(\theta)$ and either (i) $\eta(\mu(c)) < q_c$ or (ii) there exists $\theta' \in \mu(c)$ with $e^{\theta'}_c < e^{\theta}_c$.

We will 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 kept fixed. Whenever there is no risk of confusion we will omit dependence of certain variables on the economy, to make the notation less cumbersome.

2.1.1 Cutoffs

Scores of marginal accepted students in each college will play a large role in the analysis. This subsection shows that the score of the marginal accepted student at a college, which we term the college's cutoff following Abdulkadiroglu et al. (2008), parametrizes the set of stable matchings. This idea is closely related to a result by Roth and Sotomayor (1989). They show that the entering classes a college receives in any two stable matchings are ordered by first order stochastic dominance. This suggests the possibility of parametrizing the set of stable matchings using the score of the worst student in each college's entering class.

Throughout this subsection, we fix an economy E, and abuse notation by omitting dependence on E when there is no risk of confusion. 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 $p_c \leq e^\theta_c$, that is c would accept θ. A student's **demand** given a vector of cutoffs is her favorite college among those that would accept her. That is,

$$D^\theta(p) = \arg\max_{\succ^\theta}\{c|p_c \leq e^\theta_c\}. \tag{1}$$

Aggregate demand for college c is the mass of students that demand it,

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

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

Definition 2. A vector of cutoffs p is a **market clearing cutoff** if satisfies the market clearing equations: for all c

$$D_c(p) \leq q_c$$

and $D_c(p) = q_c$ if $p_c > 0$.

Market clearing cutoffs can be used to parametrize stable matchings. To describe this parametrization, we define two operators. Given a market clearing cutoff p, we define the associated matching $\mu = \mathcal{M}p$ using the demand function:

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

Conversely, for a stable matching μ, we define the associated cutoff $p = \mathcal{P}\mu$ by:

$$p_c = \inf_{\theta \in \mu(c)} e_c^\theta. \tag{2}$$

The operators $\mathcal{M}$ and $\mathcal{P}$ give a bijection between stable matchings and market clearing cutoffs.

Lemma 1. ***(Cutoff Lemma)***[19] *If μ is 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.*[20]

Intuitively, the Lemma says that stable matchings can be described by cutoff scores at each college. Given a stable matching, define its corresponding cutoffs for each college as the lowest score of all students matched to the college. We then have that if each student points to her favorite affordable college, the result is the stable matching. This means we could have defined stability directly in terms of

[19] To our knowledge, the discrete and continuum versions of the cutoff lemma are new, but they have some precursors in the literature. Abdulkadiroglu et al. (2008) use cutoffs extensively, in a model where all colleges rank students in the same order, and introduced the term cutoff. Biró (2007) describes the algorithm used for college admissions in Hungary. In the algorithm, colleges start with a low cutoff score. At each step, students apply to their favorite college that would accept them, and each college increase the cutoff score up to the point where its quota is filled exactly. With strict preferences, the outcome is the same as student proposing deferred acceptance. Biró (2007) terms this a "score limit algorithm", and remarks that a definition of stability similar to market clearing cutoffs is equivalent to the standard definition, although he does not offer a proof. Cutoffs are also related, but different, to a very interesting characterization of stable matchings due to Adachi (2000), in terms of what he calls pre-matchings. The main difference is that pre-matchings assign a "cutoff" to each man and each woman, while cutoffs only have to be assigned for one side of the market. Similar ideas have been successfully applied to a series of matching problems (Adachi (2003); Hatfield and Milgrom (2005); Echenique and Oviedo (2004, 2006); Ostrovsky (2008)).

[20] This lemma relates to the results by Roth and Sotomayor (1989), they prove that the entering classes a college may receive, in any stable matching, are always ordered by first order stochastic dominance.

cutoffs. That is, a matching μ is stable if and only if for some market clearing cutoff p we have $\mu = \mathcal{M}p$. In addition, it implies that the structure of stable matchings is simple, and stable matchings can be described by a vector of one real number per school. Moreover, the Lemma guarantees that any cutoff that clears supply and demand corresponds to a stable matching. We defer the proof to the Appendix, but for the reader interested in the intuition of the proof the next section gives a proof of the counterpart of this result in the discrete model, which is simpler and contains similar ideas.

The lemma shows that we could have equivalently defined stability by using cutoffs, instead of the standard definition, given in section 2.1. That is, a matching μ is stable if and only if for some market clearing cutoff p we have $\mu = \mathcal{M}p$. In addition, the Lemma specifies a natural bijection between stable matchings and market clearing cutoffs. If one could compute the cutoffs related to a stable matching, and have each student points to her favorite college that would accept her, the result would be the stable matching.

Note that demand functions depend on the economy E. When there is no risk of confusion, we will omit this dependence, as above. However, when we consider different economies, we will write $D(p|E)$or $D(p|\eta)$.

2.2 College admissions with a finite number of students

We use the standard definition of the college admissions model with a finite number of students. The set of colleges is again C. A finite economy $F = [\tilde{\Theta}, \tilde{q}]$ specifies a finite set of students $\tilde{\Theta} \subset \Theta$, and a vector of integer quotas $q_c > 0$ for each college. We assume that no college is indifferent between two students in $\tilde{\Theta}$. A **matching** for finite economy F is a function $\tilde{\mu} : \tilde{\Theta} \cup C \to C \cup 2^{\tilde{\Theta}}$ such that[21]

1. Each student is matched to a college or to herself.

2. Each college is matched to a at most $\tilde{q}_c$ students.

[21] Formally, these conditions are:

1. For all θ in $\tilde{\Theta}$ we have $\mu(\theta) \in \{\theta\} \cup C$.
2. For all $c \in C$ we have that $\mu(c) \in 2^{\tilde{\Theta}}$ and $\#\mu(c) \leq q_c$.
3. For all $\theta \in \tilde{\Theta}, c \in C$, we have $\mu(\theta) = c$ iff $\theta \in \mu(c)$.

3. 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 2.1.1. A matching $\tilde{\mu}$ is said to be **stable** for finite economy F if it has no blocking pairs.

2.2.1 Cutoffs

In this section we fix a finite economy F, and will omit dependence on F in the notation. A cutoff is a number $\tilde{p}_i$ in $[0, 1]$ specifying an admission threshold for college i. Given a vector of cutoffs p, a student's demand is defined as in section 2.1.1. Demand for a college c is defined as

$$\tilde{D}_c(\tilde{p}) = \#\{\theta \in \tilde{\Theta} : D^\theta(\tilde{p}) = c\}.$$

$\tilde{p}$ is a **market clearing cutoff** for economy F if for all colleges

$$\tilde{D}_c(\tilde{p}) \leq \tilde{q}_c,$$

with equality if $\tilde{p}_c > 0$.

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 $\mathcal{P}$; we only adjust the definition of $\tilde{\mathcal{P}}$ in that if a school has empty spots we assign it a cutoff of 0. In the discrete case, we have an analogue of the cutoff lemma. The only difference is that, in the discrete model, each matching can have many corresponding market clearing cutoffs, so we don't get a bijection.

Lemma 2. ***(Discrete Cutoff 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.*

Proof. Consider a stable matching $\tilde{\mu}$, and let $\tilde{p} = \tilde{\mathcal{P}}\tilde{\mu}$. Any student θ can afford $c = \tilde{\mu}(\theta)$, as $e^\theta_c \geq \tilde{p}_c$. It also can't afford any other college $c' \succ^\theta c$: if it could, then there would be another student θ' matched to c' with $e^{\theta'}_{c'} < e^\theta_{c'}$, which would contradict $\tilde{\mu}$ being stable. Consequently, we must have $D^\theta(\tilde{p}) = \tilde{\mu}(\theta)$. This proves both that $\tilde{\mathcal{M}}\tilde{\mathcal{P}}$ is the identity, and that p is a market clearing cutoff.

In the other direction, let p be a market clearing cutoff, and $\tilde{\mu} = \tilde{\mathcal{M}}\tilde{p}$. By the definition of the operator and the market clearing conditions it is a matching, so we only have to show there are no blocking pairs. Assume by contradiction that (θ, c) is a blocking pair. If c has empty slots, then $\tilde{p}_c = 0 \leq e_c^\theta$. If c is matched to a student θ' that it likes less than θ, then $\tilde{p}_c \leq e_c^{\theta'} \leq e_c^\theta$. Hence, we must have $\tilde{p}_c \leq e_c^\theta$. But then by the definition of μ we have $c \prec^\theta \tilde{\mu}(\theta)$, so it can't be a blocking pair, reaching a contradiction. □

Intuitively, the Lemma says that given a stable matching we can find cutoffs at each college, such that the matching is given by all students pointing to their favorite college that would accept them. This means that, even in the discrete model, stable matchings have a very simple structure. This was previously pointed out by Biró (2007),[22] although he does not provide a proof. He points out that in Hungary college admissions are made through a clearinghouse, that uses an algorithm similar to the Gale and Shapley deferred acceptance algorithm but that uses cutoffs. In addition, Lemma 2 guarantees that any set of cutoffs that clears the market corresponds to a stable matching. The only respect in which the discrete cutoff Lemma is weaker than the continuum version, is that each stable matching can correspond to several different market clearing cutoffs, while in the continuum model we have a bijection.

2.3 Convergence notions

To describe our convergence results, we must define notions of convergence for economies and stable matchings. On the set of continuum economies $\mathcal{E}$ we take the product topology given by the weak-* topology over measures and the Euclidean topology over vectors of capacities. We take the distance between stable matchings to be the distance between their associated cutoffs in the supremum norm in $\mathbb{R}^n$. That is, the distance between two stable matchings μ and μ' is

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

There is a natural way to define what it means for a sequence of discrete economies to converge to a continuum economy. Consider a discrete economy $F = [\tilde{\Theta}, \tilde{q}]$, with m students. An equivalent notation to describe it is using a measure $\eta[F]$ that gives weight $1/m$ to each point in $\tilde{\Theta}$, and a vector of quotas $q[F] = \tilde{q}/m$. Note that the

[22] The model and definitions used by Biró (2007) are slightly different. However, he states without proof that the usual definition of stability is equivalent to a definition very similar to a matching being associated with market clearing cutoffs. More substantially, our result differs from his in that we outline specific operators associating stable matchings to equilibrium cutoffs.

measure $\eta[F]$ gives positive weight to some points in Θ, so that this pair could not be a continuum economy as defined before, as it violates assumption 1. But it is normalized so that $\eta[F](\Theta) = 1$, as in the definition of a continuum economy.

Definition 3. A sequence of finite economies F^k converges to a limit economy $E = [\eta, q]$ if $\eta[F^k]$ converges to η in the weak-* topology and $q[F^k]$ converges to q in $\mathbb{R}^n$.

Given a stable matching of a continuum economy μ, and a stable matching of a finite economy $\tilde{\mu}$, we define

$$d(\tilde{\mu}, \mu) = \sup_{\tilde{p}} ||\tilde{p} - \mathcal{P}\mu||_\infty$$

over all vectors $\tilde{p}$ with $\tilde{\mathcal{M}}\tilde{p} = \tilde{\mu}$.

Definition 4. The sequence of stable matchings $\tilde{\mu}^k$ 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, given a finite economy F, we define the radius of the set of stable matchings of F as

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

2.4 A simple example

This simple example illustrates the main results. 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 $[0, 1]^2$ (figure 1). If both colleges had capacity $1/2$, the unique stable matching would have each student matched to her favorite school. To make the example interesting, assume $q_1 = 1/4$, $q_2 = 1/2$.

A familiar way of finding stable matchings is using the student-proposing deferred acceptance algorithm. At each step, unassigned students propose to their favorite college out of the ones that still haven't rejected them. If a college has more students than its capacity assigned to it, it rejects the lower ranked students it has assigned to it, to stay below its capacity. Figure 1 displays the trace of the algorithm in our example. In the first step, all students apply to their favorite school. Because school 1 only has capacity $1/4$, and each square has mass $1/2$, it then rejects half of the

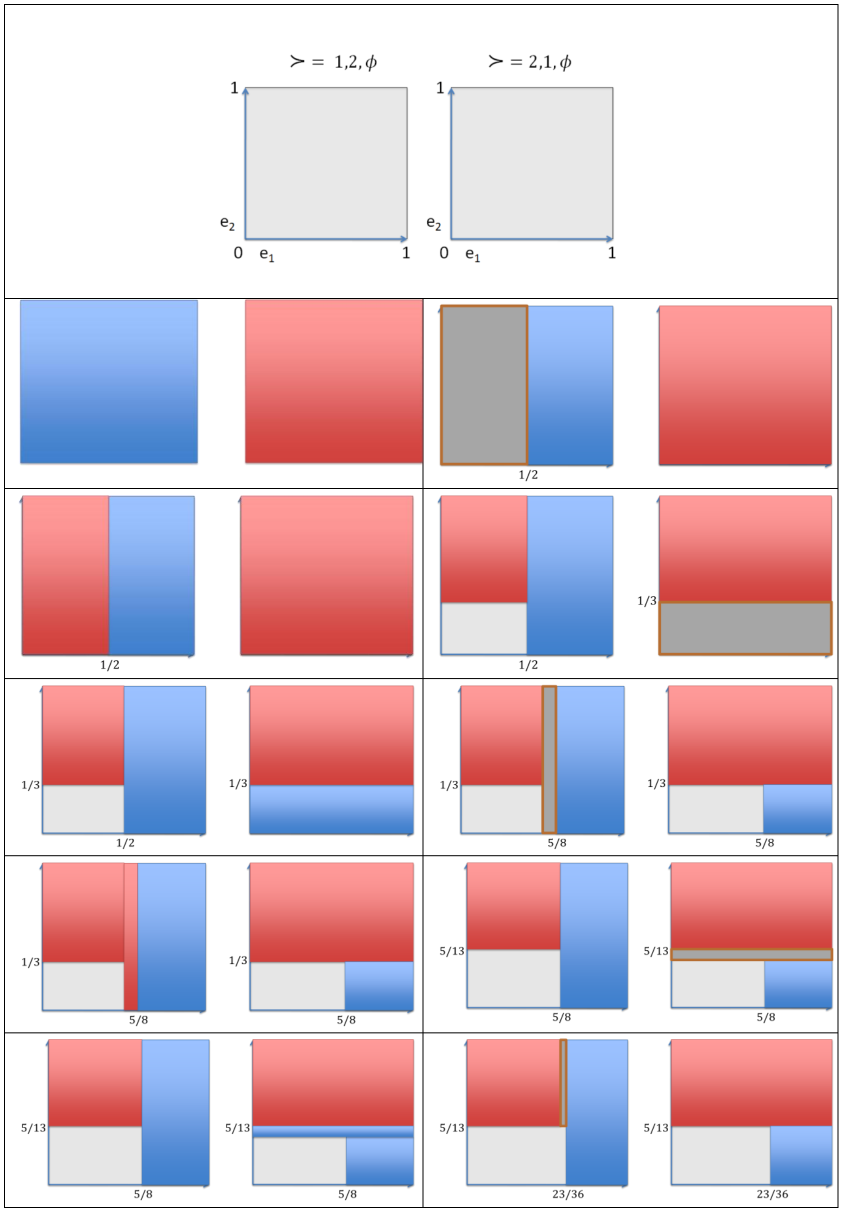

Figure 1: In the example types are uniformly distributed in the two squares on the top panel. The lower panels show the first 10 steps of the Gale-Shapley student-proposing algorithm.

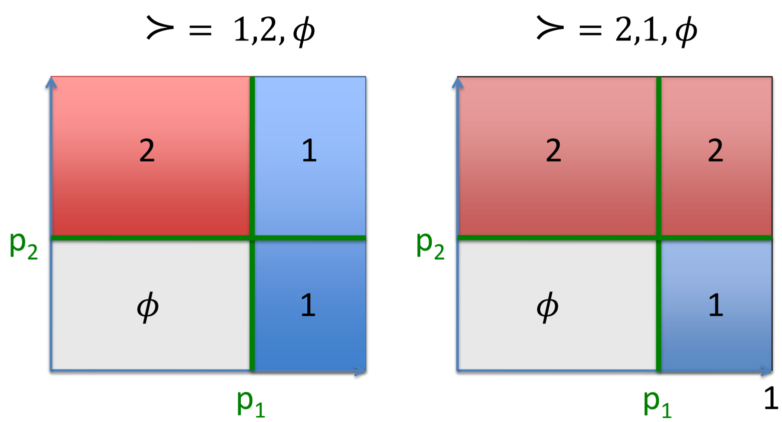

Figure 2: The outcome of the GS algorithm.

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 worse ranked students. Those who had already been rejected stay unmatched. But those who hadn't been rejected by college 1 apply to it, leaving it with more students than capacity, and the process continues. Although the algorithm does not finish, it always converges, and the outcome (figure 2) is a stable matching (see Appendix A). Figure 1 hints at this, as the measure of students getting rejected in each round is becoming smaller and smaller.

However, figures 1 and 2 give much more information than simply convergence of the deferred acceptance mechanism. We can see that cutoffs yield a simpler decentralized way to compute the matching. Note that all students accepted to college 1 are those with a score above a cutoff of $p_1 \approx .640$. And those accepted to college 2 are those with a score above some cutoff $p_2 \approx .390$. Hence, had we known these numbers in advance, it would not be necessary to run the deferred acceptance algorithm. All we would have to do is assign each student to her favorite college such that her score is above the cutoff, $e_c^\theta \geq p_c$ (Cutoff Lemma 1).

Indeed, solving for market clearing cutoffs is much simpler than running the de-

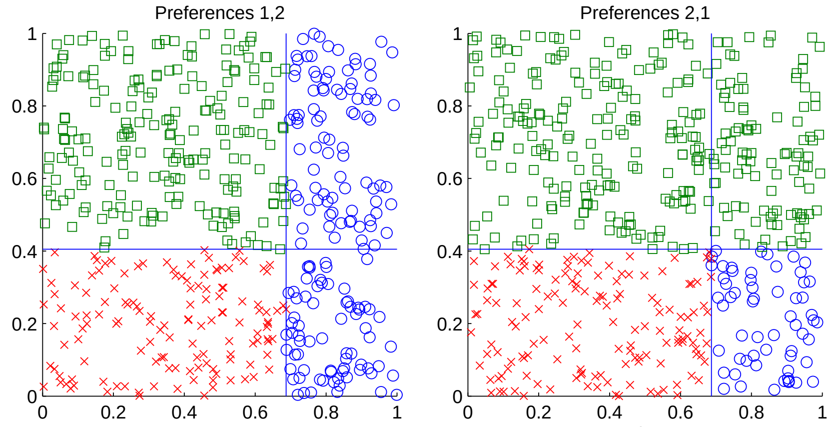

Figure 3: Cutoffs of a stable matching in a discrete economy, approximating the continuum economy in the example. There are 2 colleges, with capacities $q_1 = 250$, $q_2 = 500$. 500 students have preferences $\succ^\theta = 1, 2, \emptyset$ and 500 students have preferences $2, 1, \emptyset$. Scores e^θ were drawn independently according to the uniform distribution in $[0, 1]^2$. The figure depicts the student-optimal stable matching. Balls represent students matched to college 1, squares to college 2, and Xs represent unmatched students.

ferred acceptance algorithm. For example, 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 $p_2(1 - p_1)$. Market clearing cutoffs must satisfy the pair of equations

$$\begin{aligned} q_1 &= 1/4 = & (1 + p_2)(1 - p_1)/2 \\ q_2 &= 1/2 = & (1 + p_1)(1 - p_2)/2. \end{aligned}$$

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 shows this is a more general phenomenon).

Note that the cutoff lemma is also valid in the discrete college admissions model, save for the fact that in discrete models each stable matching may correspond to more than one market clearing cutoff (Discrete Cutoff Lemma 2). Figure 3 illustrates cutoffs for a stable matching in a discrete economy with $1,000$ students, analogous to the continuum economy in the example. Note that the cutoffs in the discrete economy are approximately the same as the cutoffs in the continuum economy. Theorem 2 shows that, generically, the market clearing cutoffs of approximating discrete economies approach market clearing cutoffs of the limit economy.

3 Results

We are now ready to state the main results in the paper. The first result shows that, typically, continuum economies have a unique stable matching.

Definition 5. Measure η is regular if the closure of the set of points

$$\{p \in [0,1]^n : D(\cdot|\eta) \text{ is not continuosuly differentiable at } p\}$$

and its image under $D(\cdot|\eta)$ have Lebesgue measure 0.

In particular, if $D(\cdot|\eta)$ is continuously differentiable then η is regular.We then have:

Theorem 1. *The economy $E = [\eta, q]$ has a unique stable matching:*

i) For any η with full support.

ii) For any regular measure η and almost every q such that $\sum_i q_i < 1$.

This result shows that, for typical parameter values, the continuum model has a unique stable matching. This is important because the convergence results depend on uniqueness, and Theorem 1 guarantees that these results apply broadly. It also shows that typically the notion of stability is enough to uniquely determine the market's allocation in the continuum model.

Proof. (Proof sketch) Here we outline the main ideas in the proof, which is deferred to Appendix B. The proof depends crucially on two results which we develop in Appendix A, which extend classic results of matching theory to the continuum model. The first is the Lattice Theorem, which guarantees that for any economy E the set of market clearing cutoffs is a complete lattice. In particular, this implies that there exist smallest and largest vectors of market clearing cutoffs. In the proof we will denote these cutoffs p^- and p^+. The other result is the Rural Hospitals Theorem, which guarantees that the measure of unmatched students in any two stable matchings is the same.

Part (i).

Note that the set of unmatched students at p^+ contains the set of unmatched students at p^-, and their difference is

$$\{\theta \in \Theta : e^\theta < p^+, e^\theta \not< p^-\}.$$

By the Rural Hospitals Theorem, this set must have η measure 0. Since η has full support, this implies that $p^- = p^+$, and therefore there is a unique stable matching.

Part (ii).

For simplicity, consider the case where for all c we have $p_c^- < p_c^+$, and where the function $D(p|\eta)$ is continuously differentiable. The general case is covered in Appendix B.

We begin by applying Sard's Theorem.[23] The Theorem states that, given a continuously differentiable function $f : \mathbb{R}^n \to \mathbb{R}^n$ we have that for almost every $q_0 \in \mathbb{R}^n$ the derivative $\partial f(p_0)$ is nonsingular at every solution p_0 of $f(p_0) - q_0 = 0$. The intuition for this result is easy to see in one dimension. It says that if we randomly perturb the graph of a function with a small vertical translation, all roots will have a non-zero derivative with probability 1.

Given q, as we assumed that there is excess demand for colleges, the market clearing cutoffs are the set of roots p of the equation

$$D(p|\eta) = q.$$

By Sard's Theorem, we have that for almost every q, the derivative $\partial_p D(\cdot|\eta)$ is invertible at every market clearing cutoff associated with $[\eta, q]$. Henceforth, we will restrict attention to an economy $E = [\eta, q]$ where this is the case.

To reach a contradiction assume that $E = [\eta, q]$ has more than one market clearing cutoff. By the Lattice Theorem we can write $p^- \neq p^+$ for the smallest and largest cutoffs. For any p in the cube $[p^-, p^+]$, the measure of unmatched students

$$1 - \sum_c D_c(p|\eta) \tag{3}$$

must be higher than the measure of unmatched students at p^- but lower than the measure at p^+. However, by the Rural Hospitals Theorem, this measure must be the same at p^- and p^+. Therefore, the expression in equation 3 must be constant in the cube $[p^-, p^+]$. This implies that the derivative of D at p^- must satisfy

$$\sum_c \partial_{p_c} D(p^-|\eta) = 0.$$

[23]See (Guillemin and Pollack, 1974).

However this implies that this derivative is not invertible, contradicting Sard's Theorem. □

The next result shows that in the case where uniqueness holds, stable matchings of the continuum model correspond to limits of stable matchings of approximating finite economies.

Theorem 2. *Assume that the continuum economy E admits a unique stable matching μ. We then have*

i) The stable matching correspondence is continuous at E.

ii) For any sequence of stable matchings $\tilde{\mu}^k$ of finite economies F^k converging to E, we have that $\tilde{\mu}^k$ converges to μ.

iii) Moreover the diameter of the set of stable matchings of F^k converges to 0.

Taken together, Theorems 1 and 2 Part (ii) imply that typically the continuum model admits a unique stable matching. In addition any sequence of stable matchings of approximating finite economies is converging to this stable matching. This shows that the continuum model has an intimate link to the discrete model, and justifies using the continuum model, under appropriate circumstances, as a simplified market model. By Part (iii), it is also the case that the set of stable matchings of the economies F^k is shrinking. This is a form of core convergence result, which says that all stable matchings of large economies become very similar as the economy grows. Roth and Peranson (1999); Immorlica and Mahdian (2005); Kojima and Pathak (2009) had shown results in this line, in markets where both the number of doctors and hospitals goes to infinity. However, their results depend on very specific stochastic processes generating preferences, and on agents having short preference lists.

Theorem 2 part (i) guarantees that the stable matchings of the limit economy vary continuously with respect to fundamentals. This validates using empirical data and simulations to study matching markets, as it shows that small measurement errors do not radically alter the set of stable matchings.

An immediate implication of Theorem 2 is that the stable matchings of an economy of agents randomly drawn according to η converge almost surely to a stable matching of the continuum model.

Corollary 1. *Assume that the continuum economy $E = [\eta, q]$ admits a unique stable matching μ. Let $F^k = [\tilde{\Theta}^k, \tilde{q}^k]$ be a randomly drawn finite economy, with k students* *drawn independently according to η and the vector of capacity per student $\tilde{q}^k/k$ converging almost surely to q. Let $\tilde{\mu}^k$ be a stable matching of F^k. Then almost surely we have that F^k converges to F, and $\tilde{\mu}^k$ converges to μ.*

This corollary follows from a direct application of the Glivenko–Cantelli theorem. Its importance is twofold. First, for a general class of random processes generating large finite economies, all sequences of stable matchings will converge to the unique stable matching given by the continuum model. Second, this can be used to characterize the asymptotics of mechanisms used in practice. One particular case is the random serial dictatorship (RSD) mechanism, which is used to allocate a number of objects (the colleges in our model correspond to object types) among agents (which correspond to the students). Agents are randomly ordered in a queue, and take turns selecting their favorite object. In a recent paper, Che and Kojima Forthcoming show that the RSD mechanism is asymptotically equivalent to the probabilistic serial mechanism proposed by Bogomolnaia and Moulin 2001. Because RSD is a particular case of the deferred acceptance mechanism when all colleges have the same preference ordering over students, their result is a particular case of ours, where the measure η has all its weight on the diagonal $\{\theta \in \Theta : e_1^\theta = \cdots = e_n^\theta\}$. In addition, Corollary 1 can be used to characterize the asymptotics of other mechanisms used in school choice, such as deferred acceptance with single tie-breaking. Appendix C provides details of these constructions.

3.1 Multiple stable matchings and robustness

Section 3 shows that most continuum economies have a unique stable matching, and that there is a close connection between the stable matchings of the continuum and discrete model in that case. The reason why uniqueness is an important requirement is that, when the continuum economy admits more than one stable matching, these matchings may not be robust with respect to small perturbations in the economy. The following example illustrates this point.

Example 1. (School Choice)

This example is based on a school choice problem. In Boston and New York City, academic economists have redesigned the centralized clearinghouse that matches students to public schools. The algorithm chosen was to start by breaking ties between students using a single lottery, and then run the student-proposing deferred acceptance algorithm.[24]

A city has two schools, $c = 1, 2$, with a quota of $q_1 = q_2 = 1$. Students have priorities to schools according to the walk zones where they live in. A mass 1 of students lives in the walk zone of each school. To break ties, the city gives each student a single lottery number l uniformly distributed in $[0, 1]$. The student's score is

$$l + I(\theta \text{ is in } c\text{'s walk zone}).$$

First note that market clearing cutoffs must be in $[0, 1]$, as the mass of students with priority to each school is only large enough to exactly fill each school. Consequently, the market clearing equations can be written

$$\begin{aligned} 1 &= q_1 = (1 - p_1) + p_2 \\ 1 &= q_2 = (1 - p_2) + p_1. \end{aligned}$$

The first equation describes demand for school 1. $1 - p_1$ students in the walk zone of 2 are able to afford it, and that is the first term. Also, p_2 students in the walk zone of 1 would rather go to 2, but don't have high enough lottery number, so they have to stay in school 1. The market clearing equation for school 2 is the same.

Note that these equations are equivalent to

$$p_1 = p_2.$$

Hence any point in the line $\{p = (x, x) | x \in [0, 1]\}$ is a market clearing cutoff - the lattice of stable matchings has infinite points, ranging from a student optimal stable matching, $p = (0, 0)$ to a school optimal stable matching $p = (1, 1)$.

Now modify the economy by adding a small mass of agents that have no priority, so that the new mass has e^θ uniformly distributed in $[(0, 0), (1, 1)]$. It's easy to see that in that case the unique stable matching is $p = (1, 1)$. Therefore adding this small mass unravels all stable matchings except for $p = (1, 1)$. In addition it is also possible to find perturbations that undo the school optimal stable matching $p = (1, 1)$.

[24] The example is a continuum version of an example used by Erdil and Ergin (2008) to show a shortcoming of deferred acceptance with single tie-breaking: it may produce matchings which are ex post inefficient with respect to the true preferences, before the tie-breaking. That is, the algorithm often produces allocations which are Pareto dominated by other stable allocations.

If we add a small amount ϵ of capacity to school 1, the unique stable matching is $p = (0, 0)$. And if we reduce the capacity of school 1 by ϵ, the unique stable matching is $p = (1 + \epsilon, 1)$, which is close to $p = (1, 1)$.

The following Proposition generalizes the example. It shows that, when the set of stable matchings is large, then none of the stable matchings are robust to small perturbations. The statement uses the fact, proven in Appendix A, that for any economy E there exists a smaller and a largest market clearing cutoff, in the sense of the usual partial ordering of $\mathbb{R}^n$.

Proposition 1. ***(Instability)*** *Consider an economy E with more than one stable matching and $\sum_c q_c < 1$. Let p be one of its market clearing cutoffs. Assume p is either strictly larger than the smallest market clearing cutoff p^-, or strictly smaller than the largest p^+. Let N be a sufficiently small neighborhood of p. Then there exists a sequence of economies E^k converging to E without any market clearing cutoffs in N.*

Proof. Suppose $p > p^-$; the case $p < p^+$ is analogous. Assume N is small enough such that all points $p' \in N$ satisfy $p' > p^-$. Denote $E = [\eta, q]$, and let $E^k = (\eta, q^k)$, where $q_c^k = q_c + 1/nk$. Consider a sequence p^k of market clearing cutoffs of E^k. Then

$$\sum_{c \in C} D_c(p^k|\eta) = \frac{1}{k} + \sum q_c.$$

However, for all points p' in N,

$$\sum_{c \in C} D_c(p'|\eta) \leq \sum_{c \in C} D_c(p^-|\eta) = \sum q_c < \sum q_c^k.$$

However, for large enough k, $\sum q_c^k < 1$, which means that for any market clearing cutoff p^k of E^k we must have $D(p^k|\eta) = q_c^k$, and therefore there are no market clearing cutoffs in N. $\square$

4 Conclusion

As market design tackles ever more sophisticated problems, it becomes increasingly common that exact analytic results in discrete models are not available. Several recent contributions have focused on obtaining results that are valid only asymptotically, as markets become large in some sense (Roth and Peranson (1999); Immorlica and Mahdian (2005); Budish (2008); Che and Kojima (Forthcoming); Kojima and Pathak (2009); Kojima and Manea (2009); Manea (2009)). In this paper we consider the case, ubiquitous in practice, of matching markets where agents on one side are matched to several agents on the other side. We propose a variation of the Gale and Shapley (1962) college admissions problem, where a finite number of colleges is matched to a continuum of students that captures this setting.

The main results are, first, the convergence results outlining the close connection between stable matchings of the continuum model and of approximating discrete economies. This lays foundations for continuum models that have been used in the case of perfectly correlated college preferences, and permits extending their analysis to more complex settings (Abdulkadiroglu et al. (2008); Miralles (2008)). Second, we find that generically the set of stable matchings depends continuously on the underlying economy. This justifies the use of empirical data and simulations in the study and design of matching markets (Roth and Peranson (1999); Abdulkadiroglu et al. (2009); Budish and Cantillon (Forthcoming)). Third, our model implies that generically the continuum model has a unique stable matching. Coupled with the convergence results, this implies that large discrete economies close to a given generic limit tend to have stable matchings which are all very similar. This complements previous results showing that large economies have few stable matchings (Roth and Peranson (1999); Immorlica and Mahdian (2005)). Fourth, we use the framework to derive new results on the asymptotics of commonly used mechanisms, generalizing previous findings (Che and Kojima (Forthcoming)).

Another innovation is the use of the score of marginal accepted students (cutoffs) as a centerpiece of our analysis. One of our contributions is the cutoff lemma, which characterizes stable matchings in terms of market clearing cutoffs, and describes a natural relationship between the two. The fact that this relationship holds both in the discrete and continuum setting is the driving force behind our convergence results, and allows us to sidestep the more conventional combinatorial arguments.

The usefulness of the continuum model will depend on whether it can be fruitfully applied to new problems in matching theory and market design. In two companion papers, we use the model to tackle open questions. In Azevedo and Leshno (2010), we apply the continuum framework to study how deferred acceptance mechanisms compare with student-optimal stable mechanisms, in equilibrium. Azevedo (2010) applies the framework to understand equilibrium behavior in stable mechanisms, and the equilibrium of imperfectly competitive matching markets. In future research, it would be interesting to explore further applications of the model, and use it to derive results which, although not feasible in the discrete model, help us understand real-life matching markets.

References

Abdulkadiroglu, A. and T. Sonmez, "School Choice: A Mechanism Design Approach," *American economic review*, 2003, *93* (3), 729.

_ **, P.A. Pathak, A.E. Roth, and T. Soenmez**, "The Boston Public School Match," *The American economic review*, 2005, *95* (2).

_ **,** _ **, and** _ , "The New York city high school match," *American Economic Review*, 2005, *95* (2), 364–367.

_ **,** _ **, and** _ , "Strategy-proofness versus efficiency in matching with indifferences: Redesigning the NYC high school match," *American Economic Review*, 2009, *99* (5), 1954–1978.

_ **, Y.K. Che, and Y. Yasuda**, "Expanding "Choice" in School Choice," *Unpublished manuscript*, 2008.

Adachi, H., "On a characterization of stable matchings," *Economics Letters*, 2000, *68* (1), 43–49.

_ , "A search model of two-sided matching under nontransferable utility," *Journal of Economic Theory*, 2003, *113* (2), 182–198.

Avery, C., A. Fairbanks, and R. Zeckhauser, *The early admissions game: Joining the elite*, Harvard Univ Press, 2004.

Azevedo, E.M., "Imperfect competition in two-sided matching markets," 2010.

_ **and J. Leshno**, "Can we make school choice more efficient? An incentives approach," 2010.

Balinski, M. and T. Sonmez, "A Tale of Two Mechanisms: Student Placement," *Journal of Economic Theory*, 1999, *84* (1), 73–94.

Biró, P., "Higher education admission in Hungary by a "score-limit algorithm"," in "in" The 18th International Conference on Game Theory at Stony Brook University 2007.

Bogomolnaia, A. and H. Moulin, "A new solution to the random assignment problem," *Journal of Economic Theory*, 2001, *100* (2), 295–328.

Budish, E.B., "The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes," *Harvard University, Unpublished mimeo*, 2008.

_ **and E. Cantillon**, "The Multi-Unit Assignment Problem: Theory and Evidence from Course Allocation at Harvard," *American Economic Review*, Forthcoming.

Che, Y. and F. Kojima, "Asymptotic Equivalence of Probabilistic Serial and Random Priority Mechanisms," *Econometrica: Journal of the Econometric Society*, Forthcoming.

Echenique, F. and J. Oviedo, "Core many-to-one matchings by fixed-point methods," *Journal of Economic Theory*, 2004, *115* (2), 358–376.

_ **and** _ , "A theory of stability in many-to-many matching markets," *Theoretical Economics*, 2006.

Ehlers, L., "Manipulation via capacities revisited," *Games and Economic Behavior*, 2010.

Erdil, A. and H. Ergin, "What's the Matter with Tie-Breaking? Improving Efficiency in School Choice," *The American Economic Review*, 2008, *98* (3), 669–689.

Gale, D. and LS Shapley, "College admissions and the stability of marriage," *American Mathematical Monthly*, 1962, *69* (1), 9–15.

Ginsburg, T. and J.A. Wolf, "The Market for Elite Law Firm Associates," *Florida State University Law Review*, 2003, *31*, 909.

Guillemin, V. and A. Pollack, *Differential topology*, Prentice-Hall Englewood Cliffs, NJ, 1974.

Hatfield, J.W. and P.R. Milgrom, "Matching with contracts," *American Economic Review*, 2005, *95* (4), 913–935.

Immorlica, N. and M. Mahdian, "Marriage, honesty, and stability," in "Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms" Society for Industrial and Applied Mathematics 2005, p. 62.

Kelso, A.S. and V.P. Crawford, "Job matching, coalition formation, and gross substitutes," *Econometrica: Journal of the Econometric Society*, 1982, *50* (6), 1483–1504.

Kesten, O., "On Two Kinds of Manipulation for School Choice Problems," *Tepper School of Business*, 2008, p. 224.

Knuth, D.E., *Mariages stables et leurs relations avec d'autres problèmes combinatoires: introduction à l'analyse mathématique des algorithmes*, Presses de l'Université de Montréal, 1976.

Kojima, F., "Mixed strategies in games of capacity manipulation in hospital–intern markets," *Social choice and Welfare*, 2006, *27* (1), 25–28.

_ **and M. Manea**, "Incentives in the probabilistic serial mechanism," *Journal of Economic Theory*, 2009.

_ **and P.A. Pathak**, "Incentives and stability in large two-sided matching markets," *American Economic Review*, 2009, *99* (3), 608–627.

Konishi, H. and M.U. Unver, "Games of Capacity Manipulation in Hospital-Intern Markets," *Social choice and Welfare*, 2006, *27* (1), 3–24.

Manea, M., "Asymptotic ordinal inefficiency of random serial dictatorship," *Theoretical Economics*, 2009, *4* (2), 165–197.

Milnor, J.W., *Topology from the differentiable viewpoint*, Princeton Univ Pr, 1997.

Miralles, A., "School choice: The case for the Boston mechanism," 2008.

Mumcu, A. and I. Saglam, "Games of capacity allocation in many-to-one matching with an aftermarket," *Social Choice and Welfare*, 2009, *33* (3), 383–403.

Niederle, M. and L. Yariv, "Matching through decentralized markets," Technical Report, mimeo, Stanford University, 2007.

Ostrovsky, M., "Stability in supply chain networks," *American Economic Review*, 2008, *98* (3), 897–923.

Romero-Medina, A. and M. Triossi, "Games of Capacities: A (Close) Look to Nash Equilibria," *Carlo Alberto Notebooks*, 2007.

Roth, A.E., "The college admissions problem is not equivalent to the marriage problem," *Journal of Economic Theory*, 1985, *36* (2), 277–288.

_ , "Deferred acceptance algorithms: history, theory, practice, and open questions," *International Journal of Game Theory*, 2008, *36* (3), 537–569.

_ **and E. Peranson**, "The Redesign of the Matching Market for American Physicians: Some Engineering Aspects of Economic Design," *The American Economic Review*, 1999, *89* (4), 748–780.

_ **and M. Sotomayor**, "The college admissions problem revisited," *Econometrica: Journal of the Econometric Society*, 1989, *57* (3), 559–570.

Sonmez, T., "Manipulation via Capacities in Two-Sided Matching Markets," *Journal of Economic Theory*, 1997, *77* (1), 197–204.

Appendix

Guide to the appendix

This appendix includes proofs of the results in the text, as well as additional results. The Appendix is organized as follows. Appendix A extends some classic results of classic matching theory to the continuum model. It provides a proof of the continuum cutoff lemma 1, and of Theorem 1. Appendix B derives results on the continuity of the stable matching correspondence, and on the convergence of stable matchings of discrete economies. It provides proofs of Theorem 2 and Corollary 1. Appendix C then discusses how to use the model to obtain results on the asymptotics of the RSD mechanism, and of some school choice mechanisms. Appendix D extends the model to matching with contracts.

A Basic Results

We begin the analysis by deriving some basic properties of the set of stable matchings in the continuum model. Besides being of independent interest, they will be useful in the derivation of the convergence results. Throughout this section we fix a continuum economy $E = [\eta, q]$, and omit dependence on E, η, and q in the notation.

First we will prove the continuum cutoff Lemma 1.

Proof. **(Lemma 1)** Let μ be a stable matching, and $p = \mathcal{P}\mu$. Consider a student θ with $\mu(\theta) = c$. By definition, $p_c \leq e^{\theta}_{c}$. Consider a college c' that θ prefers over c. By right continuity, there is a student θ_+ with slightly higher rank that is matched to c and prefers c'. By stability of μ all the students that are matched to c' have higher rank than θ_+, so $p_{c'} \geq e^{\theta_+}_{c'} > e^{\theta}_{c'}$. This means that c is better than any other college that θ can demand, so $D^{\theta}(p) = \mu(\theta)$. This shows that no school is over-demanded given p, and that $\mathcal{M}\mathcal{P}\mu = \mu$. To conclude that p is a market clearing cutoff we are observe that if $\eta(\mu(c)) < q_c$ stability implies that a student that most prefers c and has rank zero is matched to c, so $p_c = 0$.

Let p be a market clearing cutoff, and $\mu = \mathcal{M}p$. First, by the definition of $D^{\theta}(p)$, μ is right-continuous. Because p is a market clearing cutoff, μ respects quota constraints. To show that μ is stable, consider any potential blocking pair (θ, c) with $\mu(\theta) \prec^{\theta} c$. Since θ could not demand c it must be that $p_c > e^{\theta}_{c}$, so $p_c > 0$ and c has no empty slots. If $\theta' \in \mu(c)$ we have $e_c^{\theta'} \geq p_c > e_c^{\theta}$, so (θ, c) is not a blocking pair. Thus μ is stable. Let $p' = \mathcal{P}\mu$. If $\mu(\theta) = c$, then $e_c^{\theta} \geq p_c$. This implies that $p'_c \geq p_c$. But if θ is a student with $e_c^{\theta} = p_c$ that most prefers c, then $\mu(\theta) = c$. Therefore $p'_c \leq p_c$. Together $p' = p$, showing that $\mathcal{PM}p = p$. □

Now consider the sup ($\vee$) and inf ($\wedge$) operators on $\mathbb{R}^n$ as lattice operators on cutffs. That is, given two vectors of cutoffs

$$(p \vee p')_c = \sup\{p_c, p'_c\}.$$

We then have that the set of market clearing cutoffs forms a complete lattice with respect to these operators.

Theorem 3. ***(Lattice Theorem)*** *The set of market clearing cutoffs is a complete*[25] *lattice under* $\vee, \wedge$.

Proof. Consider two market clearing cutoffs p and p', and let $p^+ = p \vee p'$. Take a college c, and assume without loss of generality that $p_c \leq p'_c$. By the definition of demand, we must have that $D_c(p^+) \geq D_c(p')$, as $p_c^+ = p'_c$ and the cutoffs of other colleges are higher under p^+. Also, if $p'_c > 0$, then $D_c(p^+) \geq q_c \geq D_c(p)$. And if $p'_c = 0$, then $p_c = p'_c$, and $D_c(p^+) \geq D_c(p)$. Either way, we have that

$$D_c(p^+) \geq \max\{D_c(p), D_c(p')\}.$$

Moreover, the demand for staying unmatched must be higher under p^+ than under p or p'. Because demand for staying unmatched plus for all colleges always sums to 1, we have that for all colleges $D_c(p^+) = D_c(p) = D_c(p')$. In particular p^+ is a market clearing cutoff. The proof for the inf operator is analogous. □

This Theorem imposes a strict of structure in the set of stable matchings. It differs from the Conway lattice Theorem in the discrete setting (Knuth (1976)), as the set of stable matchings forms a lattice with respect to the operation of taking the sup of

[25] In a complete lattice, the operators must be defined and closed over any subset. In our case, the operators are defined over arbitrary sets of cutoffs as these are subsets of $\mathbb{R}^n$. For notational simplicity the proof only considers the sup of two elements, the proof for arbitrary sets is essentially the same.

A complete lattice also cannot be empty. The fact that at least one market clearing cutoff exists is proved in Corollary 2.

the associated cutoff vectors. In the discrete model, where the sup of two matchings is defined as the matching where each student gets her favorite college in each of the matchings. This statement is not valid in the continuum model.

As a direct corollary of the proof we have the following.

Theorem 4. ***(Rural Hospitals Theorem)*** *The measure of students matched to a given school is the same in any stable matching.*
Furthermore, if a college does not fill its capacity, it is matched to the same set of students in every stable matching, except for a set of students with η measure 0.

Proof. The first part was proved in the proof of Theorem 3. Let $\mu = \mathcal{M}p$, $\mu' = \mathcal{M}p'$, and $\mu'' = \mathcal{M}(p \wedge p')$. Now consider a college c such that $\eta(\mu(c)) < q_c$. Therefore $p_c = p'_c = \min\{p_c, p'_c\} = 0$. By the gross substitutes property of demand we have that $\mu(c) \subseteq \mu''(c)$. By the first part of the Theorem we have that $\eta(\mu''(c)\backslash\mu(c)) = 0$. Therefore $\eta(\mu(c)\backslash\mu'(c)) \leq \eta(\mu''(c)\backslash\mu'(c)) = 0$. Using the same argument we get that $\eta(\mu'(c)\backslash\mu(c)) = 0$. □

This result implies that a hospital that does not fill its quota in one stable matching does not fill its quota in any other stable matching. Moreover, the measure of unmatched students is the same in every stable matching, an observation that will be very useful when we prove results on uniqueness of a stable matching.

We now define the continuum version of the student-proposing deferred acceptance algorithm. The algorithm starts with all students unassigned and follows these two steps:

- Step 1: Each student that is unassigned is tentatively assigned to her favorite college that hasn't rejected her yet, if there are any.

- Step 2: If no college has more students assigned than its capacity, finish, and let the matching be for each student her currently assigned college. Otherwise, each college rejects all students strictly below a minimum threshold score such that the measure of students assigned to it is exactly q_c, and it is above its threshold score in the previous periods.

We have that, although the algorithm does not necessarily finish in a finite number of steps, the tentative assignments always converges to a stable matching.

Proposition 2. ***(Deferred Acceptance Convergence)*** *The student-proposing deferred acceptance algorithm converges pointwise to a stable matching.*

Proof. To see that the algorithm converges, note that each student can only be rejected at most n times. Consequently, for every student there exists k high enough such that in all rounds of the algorithm past k she is assigned the same college or matched to herself, so the pointwise limit exists. To see that the limit is a matching, we only have to prove that the measure of students assigned to each college is no more than its capacity. At each round k of the algorithm, let R_k be the measure of rejected students. Again, because no student can be rejected more than n times, we have $R_k \to 0$. But at round k, the excess of students assigned to each college has to be at most R_k, so in the limit each school is assigned at most its quota. Also, if the measure is less than the quota, then we know the school hasn't rejected any students throughout the algorithm.

Right continuity follows from the fact that sets of rejected students are always of the form $e_c^\theta < k$.

The proof that the outcome is stable is identical to the discrete case examined in Gale and Shapley (1962). Assume by contradiction that (θ, c) is a blocking pair. If $\eta(\mu(c)) < q_c$, by the argument above c does not reject anyone during the algorithm, which contradicts (θ, c) being a blocking pair. This implies that there is θ' in $\mu(c)$ with $e_c^{\theta'} < e_c^\theta$. At some step k of the algorithm, both agents are already matched to their final outcomes. But because θ was rejected by c in an earlier step, all agents matched to c at step k must have higher priority than e_c^θ, which is a contradiction. $\square$

This shows that the traditional way of finding stable matchings also works in the continuum model, although the algorithm converges, without necessarily finishing in a finite number of steps. An immediate corollary of this Proposition is that stable matchings always exist[26].

Corollary 2. ***(Existence)*** *There exists at least one stable matching.*

We can now prove Theorem 1.We denote the excess demand given a vector of

[26]Notice that our existence proof uses the deferred acceptance algorithm, following Gale and Shapley (1962). When we consider matching with contracts, we will give an alternative existence proof, using Tarski's fixed point Theorem. It is also possible to prove existence using the existence Theorem for finite economies, and our convergence results below.

cutoffs p and an economy $E = [\eta, q]$ by

$$z(p|E) = D(p|\eta) - q.$$

Proof. (Theorem 1) **Part 1:**

Assume by contradiction that there is more than one stable matching. By the lattice Theorem, there is a smallest p^- and highest p^+ equilibrium cutoffs. The set of students that are unmatched in p^+ but not in p^- is

$$\{\theta | e^\theta < p^+\} \backslash \{\theta | e^\theta < p^-\}.$$

This contains an open set, and so has positive measure. Therefore the measure of unmatched students is higher under p^+, contradicting the rural hospitals Theorem.

Part 2:

The proof is based on Sard's Theorem, from differential topology.[27] By Sard's Theorem, for generic q, every market clearing cutoff is a regular point of $z(\cdot|E)$.[28] That is, the derivative of z at that cutoff is invertible. We will reach a contradiction by showing that if E has multiple stable matchings, then at least one of them is not a regular point.

Formally, consider a capacity vector q such that market clearing cutoffs are regular points of z. By Sard's Theorem, this is the case for almost every q. To reach a contradiction, assume that the economy $[\eta, q]$ has more than one stable matching. Let $p^- \neq p^+$ be the minimum and maximum market clearing cutoffs. We will show p^- is not a regular point of z, contradicting Sard's Theorem.

First we consider the case $p_c^- < p_c^+$ for all $c = 1, \ldots, n$. Consider the cube $\{x \in [0,1]^n | p^- \leq x \leq p^+\}$. For any p in the cube, we have $p^- \leq p \leq p^+$. Therefore $0 = \sum_c z_c(p^-) \geq \sum_c z_c(p) \geq \sum_c z_c(p^+) = 0$. This implies that the sum $\sum_c z_c(p)$ is constant on the cube. Therefore the derivative of p^- satisfies $\partial_p z(p^-) \cdot \vec{1} = 0$, and p^-

[27] See Guillemin and Pollack (1974); Milnor (1997). Consider a C^1 function $f : \mathbb{R}^n \to \mathbb{R}^n$. Sard's Theorem says that, for generic q, all the roots of $f(x) = q$ have an invertible derivative. That is, if x_0 is a root, then $\partial_x f(x_0)$ is nonsingular.

[28] Here is a detailed argument. We have $z(p|E) = D(p|\eta) - q$. Consequently the roots of z are the points where $D(p|\eta) = q$. Denote by P_0 the set of points p where D is not continuously differentiable. P_0 is closed, because the set of points where a function is continuously differentiable is open. By our smoothness assumption, $D(P_0|\eta)$ has measure 0. Let P_1 be the set of critical points of D in $[0,1]^n \backslash P_0$. By Sard's Theorem, its image $D(P_1|\eta)$ has measure 0. Therefore, almost every q is not in the image of either P_0 nor P_1, and so it is a regular value of $D(p|\eta)$. Because D and z differ by a constant, 0 is a regular value of z for generic q.

is not a regular point.

We now turn to the case when $p_c^- < p_c^+$ for $c = 1, \dots, l$ and $p_d^- = p_d^+$ for $d = l+1, \dots, n$. Let F be the subspace of $\Re^n$ that spans dimensions 1 to l. Let us consider $z(\cdot|E)$ on the cube $\{x \in [0,1]^n | \, p^- \leq x \leq p^+\}$. By the definition of demand $\frac{\partial z_c(\cdot|E)}{\partial p_{c'}} \geq 0$ for $c \neq c'$, so $z_{l+1}(\cdot|E), \dots z_n(\cdot|E)$ are weakly increasing in all positive directions in the cube. As all those function are equal to 0 at both ends of the cube, p^-, p^+, they must be identically zero on the cube. Therefore the derivative $\partial_p z(p^-|E)$ must take the subspace F into itself. However, by argument in the previous paragraph, $\sum_c z_c(\cdot|E)$ is constant and equal to 0 in the cube. Therefore $\sum_{c=1,\dots,l} z_c(\cdot|E)$ is constant, and $\partial_p z(p^-|E)$ is singular. □

B Continuity and convergence

B.1 Continuity

In this section we ask when the stable matching correspondence is continuous.

Note that, by our definition of convergence, we have that if the sequence of continuum economies E^k converges to a continuum economy E, then we have that the functions $z(\cdot|E^k)$ converge pointwise to $z(\cdot|E)$. Moreover, using the assumption that firms' indifference curves have measure 0 at E, we have the following.

Lemma 3. *Consider a continuum economy $E = [\eta, q]$, a vector of cutoffs p and a sequence of cutoffs p^k converging to p. If η^k converges to η in the weak-* sense and q^k converges to q then*

$$z(p^k|[\eta^k, q^k]) = D(p^k|\eta^k) - q^k$$

converges to $z(p|E)$.

Proof. Let G^k be the set

$$\cup_i \{\theta \in \Theta : |e_i^\theta - p_i| \leq \sup_{k' \geq k} |p_i^{k'} - p_i|\}.$$

The set

$$\cap_k G^k = \cup_i \{\theta \in \Theta : e_i^\theta = p_i\},$$

has η-measure 0 by the strict preferences assumption 1. Since the G^k are nested, we have that $\eta(G^k)$ converges to 0.

Now take $\epsilon > 0$. There exists k_0 such that for all $k \geq k_0$ we have $\eta(G^k) < \epsilon/4$. Since the measures η^k converge to η in the weak sense, we may assume also that $\eta^k(G^{k_0}) < \epsilon/2$. Since the G^k are nested, this implies $\eta^k(G^k) < \epsilon/2$ for all $k \geq k_0$. Note that $D_\theta(p)$ and $D_\theta(p^k)$ may only differ in G^k. We have that

$$|D(p|\eta) - D(p^k|\eta^k)| = |D(p|\eta) - D(p|\eta^k)| + |D(p|\eta^k) - D(p^k|\eta^k)|.$$

As η^k converges to η, we may take k_0 large enough so that the first term is less than $\epsilon/2$. Moreover, since the measure $\eta(G^k) < \epsilon/2$, we have that for all $k > k_0$ the above difference is less than ϵ, completing the proof. □

Note that this Lemma immediately implies the following:

Lemma 4. *Consider a continuum economy $E = [\eta, q]$, a vector of cutoffs p a sequence of cutoffs p^k converging to p, and a sequence of continuum economies E^k converging to E. We have that $z(p^k|E^k)$ converges to $z(p|E)$.*

Upper continuity is easy to guarantee in general.

Proposition 3. ***(Upper Hemicontinuity)*** *The stable matching correspondence is upper hemicontinuous*

Proof. Consider a sequence (E^k, p^k) of continuum economies and associated market clearing cutoffs, with $E^k \to E$ and $p^k \to p$, for some continuum economy E and vector of cutoffs p. We have $z(p|E) = \lim_{k\to\infty} z(p^k, E^k) \leq 0$. If $p_c > 0$, for high enough k we must have $p_c^k > 0$ so that $z_c(p|E) = \lim_{k\to\infty} z_c(p^k, E^k) = 0$. □

With uniqueness, continuity also follows easily.

Lemma 5. ***(Continuity)*** *Let E be a continuum economy with a unique stable matching. Then the stable matching correspondence is continuous at E.*

Proof. Let p be the unique market clearing cutoff of E. Consider a sequence (E^k, p^k) of economies and associated market clearing cutoffs, with $E^k \to E$. Assume, by contradiction that p^k does not converge to p. Then p^k has a convergent subsequence that converges to another point $p' \in [0,1]^n$. By the previous Proposition, this must be a market clearing cutoff of E, reaching a contradiction. □

B.2 Convergence

We now consider the relationships between the stable matchings of a continuum economy, and stable matchings of a sequence of discrete economies that converge to it. First, we define the normalized demand function for a finite economy $F = [\tilde{\Theta}, q]$ as

$$D(\tilde{p}|F) = \tilde{D}(\tilde{p}|F)/\#\tilde{\Theta},$$

which we also denote $D(\tilde{p}|\eta[F])$. We will extend the notation of excess demand functions to include finite economies, denoting

$$z(\tilde{p}|F) = D(\tilde{p}|\eta[F]) - q[F].$$

Note that with this definition, $\tilde{p}$ is a market clearing cutoff for finite economy F iff $z(\tilde{p}|F) \leq q$, with $z_c(\tilde{p}|F) = 0$ whenever $\tilde{p}_c > 0$.

We make note of an useful particular case of Lemma 3 as the following Lemma.

Lemma 6. *Consider a limit economy E, sequence of cutoffs $\tilde{p}^k$ converging to p, and sequence of finite economies F^k converging to E. We then have that $z(\tilde{p}^k|F^k)$ converges to $z(p|E)$.*

Proposition 4. ***(Convergence)*** *Let E be a continuum economy, and $(F^k, \tilde{p}^k)$ a sequence of discrete economies and associated market clearing cutoffs, with $F^k \to E$ and $\tilde{p}^k \to p$. Then p is a market clearing cutoff of E.*

Proof. **(Proposition 4)** We have $z(p|E) = \lim_{k\to\infty} z(\tilde{p}^k|F^k) \leq 0$. If $p_c > 0$, then $\tilde{p}_c^k > 0$ for large enough k, and we have $z_c(p|E) = \lim_{k\to\infty} z_c(\tilde{p}^k|F^k) = 0$. □

When the continuum economy has a unique stable matching, we can prove the stronger result below.

Lemma 7. ***(Convergence with uniqueness)*** *Let E be a continuum economy with an unique market clearing cutoff p, and $(F^k, \tilde{p}^k)$ a sequence of discrete economies and associated market clearing cutoffs, with $F^k \to E$. Then $\tilde{p}^k \to p$.*

Proof. **(Proposition 7)** Assume, by contradiction that $\tilde{p}^k$ does not converge to p. Then $\tilde{p}^k$ has a convergent subsequence that converges to another point $p' \in [0,1]^n$. Then $z(p'|E) = \lim_{k\to\infty} z(\tilde{p}^k, F^k) \leq 0$. If $p'_c > 0$, we must have that $\tilde{p}_c^k > 0$ for all large enough k, and so $z_c(p'|E) = 0$. Therefore, $p' \neq p$ is a market clearing cutoff, a contradiction. □

Note that Theorem 2 and Corollary 1 follow from the previous results.

Proof. **(Theorem 2)** Part (i) follows from Lemma 5 and Part (ii) follows from Lemma 4. As for Part (iii), note first that given an economy F^k the set of market clearing cutoffs is compact, which follows easily from the definition of market clearing cutoffs. Therefore there exist market clearing cutoffs p^k and p'^k of F^k such that the diameter of F^k is $\|p^k - p'^k\|_\infty$. However, by Part (ii), both sequences p^k and p'^k are converging to p, and therefore the diameter of F^k is converging to 0. □

Proof. (Corollary) Given Theorem 2, it only remains to prove that the sequence of random economies F^k converges to E almost surely. It is true by assumption that $q[F^k]$ converges to q. Moreover, by the Glivenko-Cantelli Theorem, the realized measure $\eta[F^k]$ converges to η in the weak-* topology almost surely. Therefore, by definition of convergence, we have that F^k converges to E almost surely. □

C Asymptotics of mechanisms

C.1 Random serial dictatorship

The assignment problem consists of 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 mechanism (RSD). 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 (Forthcoming) have characterized the asymptotic limit of the RSD mechanism. They show 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 n object types $c = 1, 2, \ldots, n$, plus a null object $n+1$, which corresponds to not being assigned an object. Consider a sequence of finite assignment problems, with $k \to \infty$ agents. A fraction $m^k_\succ$ has each possible preference list $\succ$ over objects. There are q^k_c objects of each type per agent, plus an infinite number of copies of the null object. Assume (m^k, q^k) converges to some (m, q) with $q > 0, m > 0$.

We can describe RSD as a particular case of the deferred acceptance mechanism where all colleges have the same preferences. First, we give agents priorities based on a lottery, generating a random college admissions problem, where agents correspond to students, and colleges to objects. Given assignment problem k, randomly assign each agent a single lottery number l uniformly in $[0, 1]$, that gives her score in all colleges (that is, objects) of $e_c = l$. This induces a random discrete economy F^k defined as in Corollary 1. We have that the RSD outcome is the unique stable matching of F^k.[29]

Notice that the F^k converge almost surely to a continuum economy E with a vector q of quotas, a mass $m_\succ$ of agents with each preference list $\succ$, and scores e^θ uniformly distributed along the diagonal of $[0, 1]^n$. This limit economy has a unique market clearing cutoff $p(m, q)$. We have the following characterization of the limit of the RSD mechanism.

Proposition 5. *Under the RSD mechanism the probability that an agent with preferences $\succ$ will recieve object c converges to*

$$\int_{l\in[0,1]} \mathbf{1}_{(c=\arg\max_{\succ}\{c\in C|p_c(m,q)\leq l\})} dl.$$

That is, the cutoffs of the limit 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 below the lottery number. Inspection of the market clearing equations shows that cutoffs correspond to 1 minus the times where objects run out in the probabilistic serial mechanism. This yields the Che and Kojima (Forthcoming) result on the asymptotic equivalence of RSD and the probabilistic serial mechanism.

C.2 School choice mechanisms

The argument in the previous section can be extended to characterize the asymptotic behavior of actual school choice mechanisms used in practice. The school choice prob-

[29] Formally, we are using the known facts that for almost every drawing of the economy preferences are strict, and that when all colleges agree on the rankings of all students, there is a unique stable matching, and that this matching corresponds to the outcome of serial dictatorship where the most preferred students choose first.

lem consists in 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 have priorities. And differs from the classic college admissions problem in that often schools are indifferent between large sets of students (Abdulkadiroglu and Sonmez (2003)). For example, a school may give priority to students students living within its walking zone, but treat all students within a priority class equally. In Boston and NYC, the clearinghouses that assign seats in public schools to students were recently redesigned by academic economists (Abdulkadiroglu et al. (2005a,b)). 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 algorithm, given those refined preferences (Abdulkadiroglu et al. (2009)).

We can use our framework to derive the asymptotics of the DA-STB mechanism. Fix a set of schools $C = 1, \ldots, n, n+1$ (which correspond to the colleges in our framework). School $n+1$ is the null school that corresponds to being unmatched, and is the least preferred school of each student. Student types $\theta = (\succ^\theta, e^\theta)$ are again given by a strict preference list $\succ^\theta$ and a vector of scores e^θ. However, to incorporate the idea that schools only have very coarse priorities, corresponding to a small number of priority classes, we assume that all e^θ_c are integers in $\{0, 1, 2, \ldots, \bar{e}\}$ for $\bar{e} > 0$. Therefore the set of possible student types is finite. We denote by $\bar{\Theta}$ the set of possible types. Consider a sequence of school choice problems, each with $k \to \infty$ students. Problem k has a fraction m^k_θ of students of each type, and school c has capacity q_c per student. The null school has capacity $q_{n+1} = \infty$. Assume (m^k, q^k) converges to some (m, q) with $q > 0, m > 0$.

We can describe the DA-STB mechanism as first breaking indifferences through a lottery, which generates a college admissions model, and then giving each student the student-proposing deferred acceptance allocation. Assume each student receives a lottery number l 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, $e^\theta_c + l$. Therefore, for each k the lottery yields a randomly generated finite economy F^k, as the one defined in Corollary 1. The DA-STB mechanism then assigns each student in F^k to her match in the unique student-optimal stable matching in F^k. For each type θ in the original problem, denote by $x^k_{DA-STB}(\theta)$ in ΔC the random allocation she receives from the DA-STB mechanism.

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

Proposition 6. *Assume the limit economy E has a unique market clearing cutoff $p(m,q)$. Then the probability that DA-STB assigns student θ to school c converges to*

$$\int_{l \in [0,1]} \mathbf{1}_{(c = \arg\max_{\succ^\theta} \{c \in C \,|\, p_c(m,q) \leq e_c^\theta + l\})} dl.$$

The Proposition says that the asymptotic limit of the DA-STB allocation can be described using cutoffs. The intuition is that, after tie-breaks, a discrete economy with a large number of students is very similar to a continuum economy where 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 q, example 1 show that it is not always the case.

This result also suggests that the outcome of the DA-STB mechanism should display small aggregate randomness, even though the mechanism is based on a lottery. The Proposition suggests that, for almost every vector (m,q), the market clearing cutoffs of large discrete economies approach the unique market clearing cutoff of the continuum limit. Therefore, although the allocation a student receives depends on her lottery number, she faces approximately the same cutoffs with very high probability. This is consistent with simulations using data from the New York City match, reported by Abdulkadiroglu et al. (2009). For example, they report that in multiple runs of the algorithm, the average number of applicants who are assigned their first choice is 32,105.3, with a standard deviation of only 62.2.

Another important feature of the Proposition is that the asymptotic limit of DA-STB given by cutoffs is analytically simpler than the allocation in a large discrete economy. To compute the allocation of DA-STB in a discrete economy, it is in principle necessary to compute the outcome for all possible ordering of the students by a lottery. Therefore, to compute the outcome with ten students, it is necessary to consider $10! \approx 4 \cdot 10^6$ lottery outcomes, and for each one compute the outcome of the deferred acceptance algorithm. For an economy with 100 students, the number of possible lottery outcomes is $100! \approx 10^{156}$. Consequently, the continuum model can be applied to derive analytic results on the outcomes of DA-STB in large economies. Azevedo and Leshno (2010) apply this model to compare the equilibrium properties of deferred acceptance with student optimal stable mechanisms.

In addition, 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 where 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.

D Matching with contracts

D.1 The Setting

In many markets, agents must negotiate not only who matches with whom, but also wages and other terms of contracts. When hiring faculty most universities negotiate both in wages and teaching load. Firms that supply or demand a given production input may negotiate, besides the price, terms like quality or timeliness of the deliveries. This section extends the continuum model to include these possibilities.

Formally, we now consider a set of **doctors** Θ distributed according to a measure η, a finite set of **hospitals** H, and a set of **contracts** X. η is assumed to be defined over a σ-algebra Σ^Θ. Each contract x in X specifies

$$x = (\theta, h, w),$$

that is, a doctor, a hospital, and other terms of the contract w. A case of particular interest, to which we return to later, is when w is a wage, and agents have quasilinear preferences.

We also assume that X contains a null contract $\emptyset$, that corresponds to being unmatched. A **matching** is a function

$$\mu : \Theta \cup H \to X \cup 2^X$$

that associates each doctor (hospital) to a (set of) contract(s) that contain it, or to the empty contract. In addition, each doctor can be assigned to at most one hospital. Moreover, hospitals must be matched to a set of doctors of measure at most equal to its quota q_h. Finally, a matching has to be stable with respect to the product σ-algebra given by Σ^{Θ} in the set Θ and the σ-algebra 2^H in the set of hospitals.

Models of matching with contracts have been proposed by Kelso and Crawford (1982); Hatfield and Milgrom (2005). Those papers define stable matchings with respect to preferences of firms over sets of contracts. While this is an interesting direction, we focus on a simpler model, where stability is defined with respect to preferences of firms over single contracts. This corresponds to the approach that focuses on responsive preferences in the college admissions problem. This restriction considerably simplifies the exposition, as the same arguments used in the previous sections may be applied. Henceforth we assume that hospitals have **preferences** over single contracts and the empty contract $\succ^h$, and agents have preferences over contracts and over being unmatched $\succ^{\theta}$.

A single agent (doctor or hospital) **blocks** a matching μ if it is matched to a contract that is worse than the empty contract. A matching is **individually rational** if no single agent blocks it. A doctor-hospital pair θ, h is said to **block** a matching μ if they are not matched, there is a contract $x = (\theta, h, w)$ that θ prefers over $\mu(\theta)$ and either (i) hospital h did not fill its capacity $\eta(\mu(h)) < q_h$ and h prefers x to the empty contract, (ii) h is matched to a contract x' which it likes less than contract x .

Definition 6. A matching μ is **stable** if

- It is individually rational.
- There are no blocking pairs.

Assume also that doctor's preferences can be expressed by an utility function $u^{\theta}(x)$. And hospital's by an **utility function** $\pi_h(x)$. To get an analogue of the cutoff Lemma, we impose some additional restrictions. Let X_h^{θ} be the set of contracts that contain both a hospital h and a doctor θ.

Assumption 2. ***(Regularity Conditions)***

- **(Boundedness)** All $u^{\theta}(x)$ and $\pi_h(x)$ are contained in $[-M, M]$.

- **(Compactness)** For any doctor-hospital pair θ, h, the set of pairs

$$\{(u^\theta(x), \pi_h(x)) | x \in X_h^\theta\}$$

is compact.

- **(No Redundancy)** Given θ, h, no contract in X_h^θ weakly Pareto dominates, nor has the same payoffs as another.

- **(Richness)** Given h and $k \in [-M, M]$, there exists an agent $\theta \in \Theta$ whose only acceptable hospital is h, and $X_h^\theta = \{x\}$ with $\pi_h(x) = k$.

- **(Measurability)** The σ-algebra Σ^Θ contains all sets of the form

$$\{\theta \in \Theta | K \subseteq \{(u^\theta(x), \pi_h(x)) | x \in X_h^\theta\}\}.$$

D.2 Cutoffs

We can write an agent's maximum utility of working for a hospital h and providing the hospital with utility of at least a cutoff p_h as

$$\begin{aligned} \bar{u}_h^\theta(p_h) &= \sup u^\theta(x) \\ \text{s.t.} \quad & x \in X_h^\theta \\ & \pi_h(x) \geq p_h. \end{aligned}$$

We refer to this as the **reservation utility** that hospital offers the doctor. Note that the reservation utility may be $-\infty$, if the feasible set is empty. Moreover, whenever this sup is finite, it is attained by a contract x, due to the compactness assumption. We will also define $\bar{u}_\emptyset^\theta(\cdot) \equiv 0$.

Now we define a doctor's demand. Note that doctors demand hospitals, and not contracts. The **demand** of a doctor θ given a vector of cutoffs p is

$$D^\theta(p) = \arg \max_{H \cup \{\emptyset\}} \bar{u}_h^\theta(p_h),$$

Demand may not be uniquely defined, as an agent may have the same reservation utility in more than one hospital. Henceforth we make an assumption, analogous to the case without contracts, that indifferences only occur in sets of measure 0.

Assumption 3. ***(Strict Preferences)*** *For any vector p, and hospitals h, h', the set of agents with $\bar{u}^\theta_h = \bar{u}^\theta_{h'}$ has measure* 0.

From now on, we fix a selection from the demand correspondence, so that it is a function. The **aggregate demand** for a hospital is defined as

$$D_h(p) = \eta(\{D^\theta(p) = h\}).$$

notice that this does not depend on the demand of agents which are indifferent between more than one hospital, by the strict preferences assumption.

A **market clearing cutoff** is defined exactly as in definition 2. Given a stable matching μ, let $p = \mathcal{P}\mu$ be given by

$$p_h = \inf\{\pi_h(x) | x \in \mu(h)\},$$

if $\eta(\mu(h)) = q_h$ and $p_h = 0$ otherwise. Given a market clearing cutoff p, let $\mu = \mathcal{M}p$ be given by the demand function. Given θ, let the hospital to which θ is matched be denoted $h = D^\theta(p)$. If $h \in H$, let the contract that θ gets be

$$\begin{aligned} \mu(\theta) \quad &= \quad \arg\max_{x \in X^\theta_h} \pi_h(x) \\ &\text{s.t.} \quad u^\theta(x) \geq \bar{u}^\theta_{h'}(p) \text{ for all } h' \neq h, \end{aligned}$$

and $\mu(\theta) = \emptyset$ otherwise. Note that $\mu(\theta)$ is uniquely defined, by the compactness and no redundancy assumptions. We have the following extension of the cutoff Lemma.

Lemma 8. ***(Cutoff Lemma with Contracts)*** *If μ is a stable matching, then $\mathcal{P}\mu$ is a market clearing cutoff, and if p is a market clearing cutoff then $\mathcal{M}p$ is a stable matching. Also $\mathcal{PM}$ is the identity.*

Proof. Let μ be a stable matching, and $p = \mathcal{P}\mu$. Consider a doctor θ. Let x be any contract she strictly prefers to $\mu(\theta)$, in any hospital different than the one to which she is matched. By definition of stability, that hospital must be filling its quota, and for all contracts $x' \in \mu(h)$ we must have $\pi_h(x') \geq \pi_h(x)$. Because this is true for any such contracts, $\bar{u}^\theta_h(p_h) \leq p_h$. By the strict preferences assumption, except for a measure 0 set we have $\bar{u}^\theta_h(p_h) < p_h$ for all such agents θ. Hence, for almost every agent,

$$D^\theta(p) = \mu(\theta),$$

and so aggregate demand satisfies $D(p) \leq q$. By the completeness assumption and stability, we must have that if $D_h(p) < q$, then $p_h = 0$.

Now consider a market clearing cutoff p, and let $\mu = \mathcal{M}p$. It is immediate that μ respects the capacity constraints. It also is individually rational. Hence, we only have to show it has no blocking pairs. Assume, by contradiction, that (θ, h) is a blocking pair. Assume θ is matched to h', signing a contract x'. Then $u^\theta(x') = \bar{u}^\theta_{h'}(p)$. But since (θ, h) is a blocking pair there is a contract x giving utility larger than this to θ and profits of at least p_h to h. This is a contradiction with $\bar{u}^\theta_{h'}(p) > \bar{u}^\theta_h(p)$. If there were no such h', then θ would be unmatched, although $\bar{u}^\theta_h > 0$, a contradiction. □

In the case of matching with contracts, there is no longer a bijection between market clearing cutoffs and stable matchings.

D.3 Existence, Lattice Property, and Rural Hospitals

The proofs of the lattice Theorem and the rural hospitals Theorem only relied on the fact that aggregate demand has a gross substitutes property. Therefore these results extend to the case of matching with contracts using the same argument.

Corollary 3. *The lattice Theorem 3 and the rural hospitals Theorem 4 extend to the matching with contracts model.*

As for existence of a stable matching, we must modify the previous argument, which used the deferred acceptance algorithm. One easy modification is using a version of the algorithm that Biró 2007 terms a "score limit algorithm", which calculates a stable matching by progressivelly increasing cutoffs to clear the market. A straightforward application of Tarski's fixed point Theorem gives us existence in this case.

Proposition 7. *A stable matching always exists.*

Proof. Consider the operator $p' = Tp$ given by the smallest solution $p' \in [0, M]^n$ to the system of inequalities

$$D_h(p'_h, p_{-h}) \leq q_h.$$

T is weakly increasing in p. Moreover, it takes the cube $[0, M]^n$ in itself. By Tarski's fixed point Theorem, it has a fixed point, which must be a market clearing cutoff. □

D.4 The quasilinear case

A particularly interesting case of the model is when contracts only specify a wage w, and preferences are quasilinear. That is, the utility of a contract $x = (\theta, h, w)$ is just

$$\begin{aligned} u^\theta(x) &= u_h^\theta + w \\ \pi_h(x) &= \pi_h^\theta - w. \end{aligned}$$

and contracts include all possible ws, such that these values are in $[-M, M]$. Define the surplus of a doctor-hospital pair as

$$s_h^\theta = u_h^\theta + \pi_h^\theta.$$

If we assume that M is large enough so that, for all θ in the support of η we have $s_i^\theta < M$, doctors and hospitals freely divide the surplus of a relationship with each other. From the definition of reservation utility we get that for all doctors in the support of η

$$\bar{u}_h^\theta(p) = s_h^\theta - p_h.$$

Therefore, in any stable matching, doctors are sorted into the hospitals where $s_h^\theta - p_h$ is the highest, subject to it being positive. One immediate consequence is that doctors *do not go* necessarily to the hospital where they create the largest surplus. If $p_h \neq p_{h'}$, it may be the case that $s_h^\theta > s_{h'}^\theta$, but doctor θ is assigned to h'.