Programming Approaches to School Choice Problems

Aaron L. Bodoh-Creed
University of California, Berkeley *

Abstract

I develop a tractable continuum model that approximates the equilibria of finite matching economies. The limit model allows the market designer to build policy goals into the objective of a convex program with feasibility, incentive compatibility, and stability as constraints. I use data from the Boston Public School system to analyze the welfare-optimal match, the trade-offs between welfare and encouraging neighborhood schools, and the tension between welfare and school diversity. I find that the Gale-Shapley algorithm is approximately efficient, but that significant gains in terms of diversity or encouraging neighborhood schools are possible without significantly affecting student welfare.

1 Introduction

Matching problems are two sided markets wherein each side of the market has preferences that need to be accounted for in equilibrium. Matching markets without transfers have been the subject of intensive mechanism design efforts over the past 15 years. Boston and other large cities employ centralized matching procedures to allocate students to public schools (Abdulkadiroğlu and Sönmez [5]). Centralized matching mechanisms are used to place medical students into residencies in numerous

*I would like to thank Ilya Segal, Paul Milgrom, Doug Bernheim, Fuhito Kojima, Jon Levin, C. Lanier Benkard, Matthew Jackson, Eduardo Azevedo, Lisa Creed, Marcello Miccoli, Juuso Toikka, Alex Hirsch, Carlos Lever, Edward Augenblick, Isabelle Sin, John Lazarev, and Bumin Yenmez. I also thank the seminar participants at and faculty of the NYU Economics department and the Haas School of Business.

countries, a process that involves accommodating the preferences of tens of thousands of students and thousands of hospitals (Roth [39]). The U.S. military uses a centralized matching mechanism to assign cadets graduating from military academies to branches within the military (Sönmez and Switzer [46]).

The usual solution concept for matching economies, *stability*, reflects a notion of fairness that requires that all agents with a high priority for a good be allocated that good before a lower priority agent is provided the good. The challenge for matching market designers is to simultaneously accommodate stability, incentive compatibility, and other goals the designer might have (e.g., welfare). In school choice settings, the priority structure is often used as a tool to encourage desired policy outcomes. For example, to encourage a student to attend his or her neighborhood school, the Boston Public Schools (BPS) until recently gave higher priority for seats at a school to students in the school's walk zone.[1] In effect, the goals of the market designer have been encoded into the stability constraints.

I argue in this paper that school choice problems can be described as computationally tractable convex programs. The term *convex program* refers to a constrained optimization problem where the objective function and the constraint set are convex. Importantly, computational techniques exist to rapidly solve convex programs with a very large number of control variables and constraints. The programming approach has the advantage of allowing the market designer to explicitly describe the policy goals in the objective function of the program, consider trade-offs between the goals (e.g., welfare and school diversity), and find the globally optimal match. The programming approach also allows one to easily assess the benefits of weakening the constraints of the program by (for example) adding capacity to a popular school.

Throughout this paper I consider a setting where there is a finite set of large schools, and during my theoretical discussion I follow the literature in referring to the schools as *colleges*. Each student is characterized by a preference ranking over the colleges, verifiable traits, and nonverifiable traits. These traits may represent demographic attributes (e.g., residency location, language spoken at home) or administrative data (e.g., truancy rate, academic performance). Since I assume the set of

[1]Since many students share the same priority class (e.g., many students are in each school's walk zone), random tie-breaking is used to assign a unique priority number to each student at each school. Students typically will have different priorities at different schools. For example, a particular student will be in the walk-zone of some schools (and have relatively high priority at these schools) and be outside the walk-zone of others (and have relatively low priority).

traits is finite, in school choice settings with many students there will be many copies of each student-type in the economy.

I prove that one can approximate the game between a large, but finite, set of agents with a model in which each college has a continuum of available seats and there is a continuum of students of each type. My first task is to show that the stable set of the continuum model is generically continuous in the distribution of student-types and college capacities. I leverage this insight to show that for generic matching economies, the set of game-theoretic equilibria of the continuum model and the set of equilibria of the finite model approach one another as the number of agents in the finite model grows. From my continuity result, it follows that the matches realized in the finite and continuum economy will be close in generic matching economies.

The use of a continuum model is absolutely crucial to the convex programming approach. Since each student is a negligible part of the continuum, deviations by individual students do not affect the capacity constraints of the colleges, the stability constraints, or the incentive compatibility constraints of the other students. In the finite model deviations from truthfulness can have complex effects as a long sequence of agents can be affected by a single agent's deviation.[2] The limited impact of an agent's deviation in the continuum model makes it possible to write down the full set of feasibility, stability, and incentive compatibility constraints for the convex program in a computationally tractable form.

Having validated the continuum model as an approximation of the finite market, I use the continuum model as the basis for my convex programming exercises. I explicitly encode various policy goals (e.g., diversity, student welfare) into the objective function of the program, and I enforce school capacity limits, stability, and incentive compatibility using constraints on the program. I apply these programs using data from the high school match run by BPS for the 2011-2012 school year.

First I consider the objective of maximizing welfare, calculated as the average preference rank of the school to which the students are assigned, subject to the capacity, stability, and ordinal dominance incentive compatibility (ODIC) constraints. Although the Gale-Shapley algorithm is known to satisfy these constraints, the algorithm may not maximize welfare when random-tie breakers are employed (Erdil and

[2] The sequence of adjustments to the match following a deviation is one reason that techniques from random graph theory have proven so fruitful in analyses of finite matches with many students. Immorlica and Mahdian [28] is a pioneering example of these techniques in a one-to-one matching model.

Ergin [20], Kesten [30]). I find that the range of welfare values of matches that satisfy ODIC is extremely narrow, which implies Gale-Shapley does an excellent job of maximizing welfare. In contrast, the stability constraints allow a range of welfare values roughly 9 times larger than the range permitted by the ODIC constraints. Finally, I use the shadow prices on the capacity constraints to identify the programs with the highest marginal value for an extra seat, which suggests where resources ought to be directed when given an opportunity to expand school programs.

Second, I use the limited information provided by the BPS to infer demographic traits for each student. I use this data to maximize a linear combination of average welfare and ethnic diversity of the student-body within each school. I find that the choice of which match to implement can have large effects on the diversity of the student bodies at popular schools, but the effect on average welfare is minimal. I conclude that there is at most a modest tension between student welfare and diversity.

Third, I use my convex programs to study the trade-offs between welfare and the percentage of students that each school enrolls from its walk-zone. In a 2012 speech, Boston Mayor Thomas Menino articulated the following externality based logic for encouraging neighborhood schools:[3]

> "Something stands in the way of taking our [public school] system to the next level: a student assignment process that ships our kids to schools across our city. Pick any street. A dozen children probably attend a dozen different schools. Parents might not know each other; children might not play together. They can't carpool, or study for the same tests. [. . .] Boston will have a radically different school assignment process, one that puts priority on children attending schools closer to their homes."

My analysis shows that the percentage of students drawn from the schools' walk-zones can be significantly increased (or decreased) by the choice of which stable and ODIC match to implement with a minimal effect on student welfare.

I provide a literature review in section 2, while the model is introduced in section 3. The primary theoretical contributions are presented in section 4, which starts by proving a match exists and can be implemented as a mixture over pure assignments. I close section 4 by proving that the stable set is continuous in the distribution of student-types and college capacities for generic realizations of these variables. Section

[3]Quote taken from Dur et al. [16].

5 provides my numerical analysis of the convex programs using the data from BPS, and section 6 concludes. All of the proofs are relegated to appendices.

2 Related Literature

This paper touches on a number of strands of the school choice literature. My paper is closely related to the literature analyzing matching when schools have preferences that admit indifferences. In practice these indifferences are resolved using random tie-breakers, and then the Gale-Shapley algorithm (or some other mechanism) is used to find a match. However, the use of random tie-breakers can result in an inefficient outcome. Erdil and Ergin [20] provides an algorithm for computing the student-optimal stable match once an initial match has been computed using the Gale-Shapley algorithm. Unfortunately, this algorithm is not incentive compatible. Kesten [30] provides an algorithm for identifying situations in which an agent's priority can be altered without affecting the agent's outcome, which insures incentive compatibility while allowing for welfare improvements.

Abdulkadiroğlu et al. [4] uses data from BPS and the New York City school system to study the welfare-losses caused by random tie-breakers. The analysis technique used is to run the Gale-Shapley algorithm to find a stable match, and then to apply Gale's top trading cycles (TTC) algorithm to find an efficient (but unstable) match. The authors argue that in the case of the BPS elementary school match,[4] the outcome produced by Gale-Shapley with tie-breakers is almost the same as the efficient match, implying that the welfare costs of stability and strategy-proofness are low.

In contrast, I find using the high school data that incentive compatibility has significant costs, while stability is much less constraining. This may be due to the fact that I use a stricter welfare metric. Executing the Gale-Shapley algorithm followed by TTC results in some point on the Pareto frontier. In contrast, my convex programs use a more stringent metric, the average expected rank of the match, and explores the entire Pareto frontier to find a global optimum with respect to this metric. In other words, Gale-Shapley may result in an assignment near the frontier, but the assignment could still be suboptimal with respect to my metric. An additional possible origin for the difference is that I analyze data from years after the publication of Abdulkadiroğlu

[4]The authors mention in a footnote that the results are robust to using the high-school match data.

et al. [4].

Azevedo and Leshno [8] is the most closely related paper to this study in terms of the theoretical methods employed. Azevedo and Leshno [8] assumes an exogenous distribution of student preferences and priority numbers that do not admit indifferences. The authors then provide a succinct characterization of a stable match in terms of priority number cutoffs and show that the continuum and finite-agent models yield approximately the same equilibrium. It is not clear that one can adapt the techniques of Azevedo and Leshno [8] to allow for students to share the same priority or use their model to develop tractable programs for studying matching markets.

Azevedo and Hatfield [9] show that with a continuum of agents, the existence of a stable matching can be insured if the preferences on one side of the market have an appropriate substitutability property. Che, Kim and Kojima [13] studies a continuum model and prove that continuity of agent preferences implies the existence of a stable match even when the preferences admit complementarities. Both papers conclude that large finite analogs of their models admit matches that satisfy a weakening of stability. However, example 1 in section 2 demonstrates that approximately stable outcomes may be very different from any exactly stable outcome. In contrast, I begin by proving that the stable set is continuous for generic economies, and I use this result to prove that each stable match of my limit model can be made arbitrarily close to an exactly stable match of the finite game if sufficiently many students participate.

A number of recent papers study large matching markets, often with a focus on whether truthfulness by all agents can be supported as an approximate equilibrium. Roth and Peranson [40] points out that the set of stable matches in the National Residency Matching Program is small, which implies approximate strategy-proofness. Immorlica and Mahdian [28] and Kojima and Pathak [32] analyze models of one-to-one and many-to-one matching markets (respectively) using the Gale-Shapley algorithm when the students have strict preference rankings of fixed length independent of the market size. The main result of both of these papers is that the incentive to nontruthfully declare a preference ranking vanishes as the market grows. Lee [35] generates a similar result in an environment without a restriction to short preference lists.

A handful of empirically minded matching papers have also used large game techniques to study issues surrounding identification. Examples include Echenique et al. [17] uses a model with multiple copies of each type of agent to study the identification restrictions imposed by stability. Echenique et al. [17] assumes all agents have strict preferences and do not study the limit as the finite model approaches a continuum game. Menzel [36] uses a large market model where agent types can take values in a continuum and preferences are strict. The primary results of Menzel [36] show that the finite model converges to a limit game wherein the equilibrium objects of interest are identified.

One of the goals of this paper is to provide a method for accommodating diversity goals into stable, incentive compatible matching systems. Recently there has been a number of papers seeking to study the limits on incorporating diversity into school choice schemes. Some of these papers involve adding either quotas or reserves into the mechanism (Budish et al. [12], Halafir et al. [25], Kominers and Sönmez [33]).[5] Echenique and Yenmez [17] axiomatize several notions of diversity that are compatible with stability.

As my methodology relies on treating the matching mechanism as the solution to a programming problem, my work touches on a prior literature that treated matching problems as linear or integer programs. Roth et al. [41] formulates stable matches in a marriage market as the solution to an integer program. Baïou and Balinski [10] extends these ideas to the many-to-one college admissions problem. The thrust of these papers is to use the programming approach to prove properties of the solutions without explicitly solving the programs. Roth et al. [41] and Baïou and Balinski [10] present programs with a size that is a polynomial in the number of agents, while the program I derive has a size that is polynomial in the number of types of agents. This crucial distinction makes my program practical to implement.

Finally, I would like to point out Ashlagi and Shi [6][6] and Featherstone [21] for special mention. To my knowledge, these are the only other works that focus on implementing a large matching market via a programming problem. Featherstone [21] describes a linear program that can be used to find a rank-efficient assignment of over 20,000 applicants to the Teach for America program, but Featherstone [21] does not include incentive compatibility constraints. Ashlagi and Shi [6] focuses on the practical implementation of a match using BPS data, but the model does not include a priority structure. As I discuss below, including a priority structure can complicate the model significantly. Ashlagi and Shi [6] show that the program they solve can be solved in polynomial time if the student utilities are generated by a multinomial

[5]Boston's efforts to provide walk-zone students with priority at their local schools has recently been critiqued by Dur et al. [16] as a de facto quota.

[6]Ashlagi and Shi [6] was developed contemporaneously with this paper.

logit model, which is both an important contribution of the paper as well as a limit to the applicability of the techniques. My work neither requires such a restriction nor undertakes the difficult task of estimating a student utility model. In addition, Ashlagi and Shi [6] cleverly avoids the difficulties caused by incentive constraints by offering the students a simple menu of options.

3 Model

I adopt the language of the college admissions problem and describe one side of the market as students and the other side of the market as colleges. I denote the set of types of students as $\mathcal{S}$ and the set of types of colleges as $\mathcal{C}$. A matching mechanism in the college admissions problem matches each student to one college, but each college may be matched with multiple students.

In traditional matching models, each individual student and college is by definition a unique good in the market, and agent preferences define an ordinal ranking of all possible partners. While I continue to assume that each college is unique, I assume that there exist multiple students of each type. Each student-type is characterized by verifiable traits, nonverifiable traits, and a preference ordering over the colleges. Verifiable traits represent student attributes that are known to the mechanism (e.g., administrative data on truancy or test score) or that cannot be falsely declared except at great cost (e.g., the student's home address). Nonverifiable traits include information that can be nontruthfully revealed to the mechanism (e.g., preferences over colleges, whether a student speaks English at home).

The finite set $\mathcal{V}$ represents the possible verifiable traits of the students, and the nonverifiable traits of each student-type are drawn from the finite set $\mathcal{U}$. Finally, each student's preference ranking is a linear order over the set of colleges, and the full set of these orders is denoted $\mathcal{P}_S$. I do not require that these preferences be strict, and I assume that each student's preferences are nonverifiable. A generic student-type has the form $s = (v, u, \succeq_s) \in \mathcal{S}$ where $v \in \mathcal{V}$, $u \in \mathcal{U}$, and $\succeq_s \in \mathcal{P}_S$.

I assume that each college is defined by a unique trait that identifies the college (the college's name) drawn from the finite set $\mathcal{T}_C$, a verifiable capacity, and a linear order over $\mathcal{S}$ drawn from the set $\mathcal{P}_C$ that defines the priority ordering of each student-type. I do not require the priority orderings be strict. The total capacity of college c is $q_c \in [0, 1]$, where q_c is interpreted as the fraction of the student population that can be enrolled in college c. A generic college-type is then $c = (t_c, \succeq_c, q_c) \in \mathcal{C}$ where $t_c \in \mathcal{T}_C$ and $\succeq_c \in \mathcal{P}_C$. When I take limits as the number of students grows, I am (in effect) assuming that the number of seats at each college grows in proportion.[7]

A final notational detail is required to denote unacceptable matches and cases where agents remain unmatched. When a student is "matched" with partner $\varnothing$, this denotes the student failing to match with any college. Symmetrically, if a college seat is "matched" with student-type $\varnothing$, then that seat is empty. Finally, student preference relations of the form $\varnothing \succeq_s c$ imply that students of type s find college c unacceptable. Let the set of mutually acceptable matches be denoted $\Gamma \subset \mathcal{C} \times \mathcal{S}$.

Since the goal of the paper is to analyze large matching markets using a limit model, I need to define both the sequence of finite economies I am considering as well as the limit economy. In the N-student economy, college c can admit $\lfloor q_c N \rfloor$ students, where $\lfloor q_c N \rfloor$ denotes the largest integer that is weakly smaller than $q_c N$. I assume that the N students have types drawn independently from the probability distribution π^S over $\mathcal{S}$. The empirically realized probability distribution of types in the N-student economy is denoted $\pi_E^{S,N}$, and $\pi_E^{S,N}(s)$ denotes the realized fraction of the N student population that is of type s.

In the limit economy, college c can admit up to a measure q_c of students, and there is a measure 1 continuum of students distributed according to measure π^S exactly. This structure insures that $\pi_E^{S,N} \to \pi^S$ almost surely in the limit as $N \to \infty$. Since the distribution of types is known exactly, the limit is a complete-information economy. When describing conditions (e.g., stability) relevant for both the finite and the continuum economy, I default to the more compact notation of the continuum economy. Since the continuum model will be treated as the limit of a sequence of finite economies, I use the terms *limit model* and *continuum model* interchangeably.

4 Matching in the Finite and the Continuum Case

I now provide a method for describing feasible and stable matches that allows me to use the same structure in matching mechanisms with a finite number or a continuum

[7]Since the colleges are not strategic actors, one is free to interpret each college as a collection of small schools that the students view as perfect substitutes for one another. My preferred interpretation is more realistic for settings that feature a relatively small number of large schools such as Chicago and Boston, whereas the "many perfect substitutes" interpretation may be warranted in settings with many schools such as New York City.

of agents. I define a *match* as a function $x : (\mathcal{C} \cup \varnothing) \times (\mathcal{S} \cup \varnothing) \times \Delta(\mathcal{S}) \to [0,1]$ where $\Delta(X)$ denotes the space of probability measures over X. $x(c, s; \pi^S)$ denotes the fraction of the student population comprised of students of type s matched with college c given a distribution of student-types π^S. Where confusion will not result, I suppress the argument π^S.

To demonstrate the notation, consider a 500 student economy where $C = \{c_1, c_2\}$, $S = \{s_1, s_2\}, \pi_E^{S,N}(s_1) = \frac{3}{5}$, and $x(c_1, s_1) = x(c_2, s_1) = x(\varnothing, s_1) = \frac{1}{5}$. This means that $\frac{1}{5} * 500 = 100$ of the students are of type s_1 and are matched with college c_1, 100 students of type s_1 are matched to college c_2, and 100 students of type s_1 are left unmatched. In the limit game, $x(c_1, s_1) = \frac{1}{5}$ is interpreted as a measure $\frac{1}{5}$ of students of type s_1 (out of the total measure 1 of all students) that are matched with college c_1.

Since my matching does not distinguish between students of the same type, the match (obviously) does not specify the identity of the particular students assigned to each college. My preferred interpretation is that the match represents a stochastic assignment, which I discuss more formally in section 4.2. In the continuum game this yields

$$\Pr\{\text{Student of type } s \text{ matched to college } c\} = \frac{x(c,s)}{\pi^S(s)} \tag{4.1}$$

In the finite game the probability is essentially the same with $\pi_E^{S,N}(s)$ playing the role of $\pi^S(s)$.

I always require that the following feasibility constraints are satisfied by the match

$$\text{For all } s \in \mathcal{S}, \sum_{c \in \mathcal{C}} x(c,s) \le \pi^S(s) \tag{4.2}$$

$$\text{For all } c \in \mathcal{C}, \sum_{s \in \mathcal{S}} x(c,s) \le q_c \tag{4.3}$$

Equation 4.2 insures that the measure of seats occupied by students of type s does not exceed the total measure of these students. Equation 4.3 implies that the measure of students assigned to school c, $\sum_{s \in S} x(c,s)$, does not exceed the school's capacity, q_c.

4.1 Stability

A college-student pair $(c, s) \in \Gamma$ is a *blocking pair* if

$$\sum_{\{s':s' \succeq_c s\}} x(c, s') < q_c \text{ and } \sum_{\{c':c' \succeq_s c\}} x(c', s) < \pi^S(s) \tag{4.4}$$

The first condition implies that college c either has an empty seat or that student-type s is strictly preferred to some student that has a seat at college c. The second condition implies that some student of type s is either unmatched or matched to a university to which c is strictly preferred. A match x is *stable* if it does not admit blocking pairs.[8,9]

Lemma 1 proves that it suffices to consider the least preferred partners assigned to student-type s and college c to determine if (c, s) is a blocking pair with respect to match x. Given a match x and a student-type $s \in \mathcal{S}$, denote the worst possible outcome for students of type s as $\underline{x}_S(s) \in \mathcal{C} \cup \{\varnothing\}$, where $\varnothing$ denotes that the worst possible outcome is that the agent is unmatched. If $\sum_{c \in \mathcal{C}} x(s, c) = \pi^S(s)$, let

$$\underline{x}_S(s) = \max\{c : c \preceq_s c' \text{ for all } c' \text{ such that } x(c', s) > 0\} \tag{4.5}$$

where the maximum is taken with respect to $\succeq_s$.[10] If $\sum_{c \in \mathcal{C}} x(c, s) < \pi^S(s)$, then $\underline{x}_S(s) = \varnothing$. If $\succeq_s$ admits indifferences, then $\underline{x}_S$ can be multivalued. I let $c \succeq_s \underline{x}_S(s)$ denote that c is weakly preferred by s to all members of $\underline{x}_S(s)$.

I use the function $\underline{x}_C(c) \in \mathcal{S} \cup \{\varnothing\}$ to describe the worst partner of college c given a match x. If $\sum_{s \in \mathcal{S}} x(c, s) = q_c$, define $\underline{x}_C$ as

$$\underline{x}_C(c) = \max\{s : s \preceq_c s' \text{ for all } s' \text{ such that } x(c, s') > 0\} \tag{4.6}$$

where the maximum is taken with respect to $\succeq_c$. If $\sum_{s \in \mathcal{S}} x(c, s) < q_c$, then $\underline{x}_C(c) = \varnothing$. If $\succeq_c$ admits indifferences between student-types, then $\underline{x}_C$ can be multivalued.

[8] Stronger notions of stability such as *strong stability* and *super stability* have been proposed for matching models with indifferences. A stable match in our setting (in general) does not exist under these stronger notions of stability.

[9] Our definition of stability is identical to the notion of *ex ante stability* proposed by Kesten and Unver [31]. In the limit model or an economy with multiple agents of the same type, the two notions coincide.

[10] The awkward definition of $\underline{x}_S(s)$ is used so that $\underline{x}_S(s)$ captures all of the colleges c such that the student is indifferent between c and the least preferred college to which student type s is matched.

The following lemma proves that I can characterize stable matches using only the information contained in $\underline{x}_S$ and $\underline{x}_C$. I use this lemma to simplify the stability conditions in my convex programs.

Lemma 1. *$(c,s) \in \Gamma$ is a blocking pair (i.e., equation 4.4 holds) if and only if $c \succ_s \underline{x}_S(s)$ and $s \succ_c \underline{x}_C(c)$ hold.*

Suppose that the "if" condition holds. Then there exists some student of type s matched to a college c' where $c \succ_s c' \succeq_s \underline{x}_S(s)$ and there is a seat at college c assigned to a student s' where $s \succ_c s' \succeq_c \underline{x}_C(c)$. In this case some student of type s and college c would wish to recontract ex post (i.e., equation 4.4 holds). The intuition for the reverse direction is similar.

The set of stable matches in the continuum limit game can be described as the set of $x \in [0,1]^{|\mathcal{C}|*|\mathcal{S}|}$ that satisfies the following constraints given the distribution of students π^S.

$$\text{For all } s \in \mathcal{S}, \sum_{c \in \mathcal{C} \cup \{\varnothing\}} x(c,s) = \pi^S(s) \tag{4.7}$$

$$\text{For all } c \in \mathcal{C}, \sum_{s \in \mathcal{S} \cup \{\varnothing\}} x(c,s) = q_c \tag{4.8}$$

$$\begin{aligned}\text{For all } (c,s) \in \Gamma, & \left(\pi^S(s) - \sum_{\{c' \in \mathcal{C} \cup \{\varnothing\} : c' \succeq_s c\}} x(c',s)\right) * \\ & \left(q_c - \sum_{\{s' \in \mathcal{S} \cup \{\varnothing\} s' \succeq_c s\}} x(c,s')\right) = 0\end{aligned} \tag{4.9}$$

$$\text{For all } (c,s) \notin \Gamma, x(c,s) = 0 \tag{4.10}$$

The first two conditions are feasibility constraints. Equation 4.7 insures that an individual student is matched with at most one college seat. Equation 4.8 implies that college c enrolls at most a fraction q_c of students. Equation 4.9 encapsulates all of the restrictions imposed by stability as per equation 4.4. Equation 4.10 represents the individual rationality constraints. In the N-student game the conditions are the same except that $\pi_E^{S,N}$ plays the role of π^S. Note that including the outcome of being unmatched in the summations allows us to write equations 4.2 and 4.3 as equalities (equations 4.7 and 4.8).

Proposition 1. *Equations 4.7, 4.8, 4.9 and 4.10 are necessary and sufficient for a* *match x to be feasible and stable.*

The following proposition reveals that there is at least one feasible and stable match in the continuum game. The existence result is known for the case of markets with a finite number of agents. When the distribution of student-types and the college capacities of the colleges admit only rational values, I can treat the continuum of students (college seats) of each student-type (college) as if it were composed of a finite number of agents, where each agent represents the same positive measure of students (college seats). Given this analogy between the finite and continuum models, it is straightforward to generate a continuum analog of the deferred acceptance algorithm. When the distribution of student-types or college capacities have irrational values, I use limit analysis techniques to argue feasible and stable matches must exist since matches are known to exist in arbitrarily close economies with rational-valued distributions of students and colleges

Proposition 2. *There is at least one feasible and stable match.*

4.2 Implementation of the Match

In the continuum case, the match outcomes described by equations 4.7 through 4.10 can be implemented exactly and deterministically. I argue in this section that in the finite game the match can be implemented as a randomization over deterministic, stable outcomes. I refer to such a randomization over stable matches as a *stochastic assignment*.

To formalize a stochastic assignment in the N-student game I must delineate between the individual students of the same type. An *assignment* is a $N \times (|\mathcal{C}| + 1)$ matrix $\mathbf{X}$ defining the probability that each of the N students is assigned to each of the $|\mathcal{C}|$ colleges (the first $|\mathcal{C}|$ columns) or remains unmatched (the final column). To define the assignment associated with a match x, all of the cells corresponding to students of type s and college c are assigned the value

$$\frac{x(c,s)}{\pi_E^{S,N}(s)}$$

The probability of remaining unmatched, which is assigned to the final column for all students of type s, is

$$1 - \frac{1}{\pi_E^{S,N}(s)} \sum_{c \in \mathcal{C}} x(c,s)$$

The final allocation of students to colleges must be a *pure assignment* that places each student in a single college. In other words, the final assignment must contain only 0 or 1 values - a student is either matched to a college (1) or not matched with that college (0). A stochastic assignment can be *implemented* if the stochastic assignment is equivalent to randomizing over pure assignments. More formally, the stochastic assignment can be implemented if there exists a set of pure assignments $\{\mathbf{X_i}\}_{i=1}^A$ and positive numbers $\{\lambda_i\}_{i=1}^A$ such that $\sum_{i=1}^A \lambda_i = 1$ and

$$\mathbf{X} = \textstyle\sum_{i=1}^A \lambda_i \mathbf{X_i}$$

where each $\mathbf{X_i}$ satisfies the feasibility and stability requirements.

My proof uses theorem 1 of Budish et al. [12], which proves that the feasible set of a particular class of linear programs can be implemented using a mixture over pure assignments. Since the stability constraints are not linear (much less of the particular form studied in Budish et al. [12]), I use lemma 1 rewrite the constraints in a way that fits within the framework of Budish et al. [12].

Proposition 3. *Any assignment satisfying equations 4.7 through 4.10 can be implemented.*

4.3 Structure and Continuity of the Stable Set

In this section I provide an analysis of the stable set based purely on the description provided by equations 4.7 through 4.10. The first step involves considering the stable set as being composed of a finite collection of subsets, each of which is defined by a combination of $\underline{x}_S(s)$ and $\underline{x}_C(c)$ that satisfy lemma 1's criteria for there to be no blocking pairs. For the duration of this section, assume that any reference to a pair of $\underline{x}_S(s)$ and $\underline{x}_C(c)$ refers to a pair that does not admit blocking pairs. Given a choice of $\underline{x}_S(s)$ and $\underline{x}_C(c)$, I derive properties of each of these subsets from the feasibility equalities given by equations 4.7 and 4.8.

My proof is based on properties of systems of linear equations. If I eliminate the nonnegativity constraints, I show that for any choice of $\underline{x}_S(s)$ and $\underline{x}_C(c)$ the set of matches satisfying the stability and feasibility constraints is continuous in π^S

and $(q_c)_{c\in\mathcal{C}}$. The primary difficulty of the proof is showing that the nonnegativity constraints do not interfere with the continuity of the set of stable matches for generic choices of π^S and $(q_c)_{c\in\mathcal{C}}$. To state it differently, I rule out the possibility that the nonnegativity constraints cause the set of stable matches for a choice of $\underline{x}_S(s)$ and $\underline{x}_C(c)$ to vacillate between being empty and nonempty when I consider perturbations around generic π^S and $(q_c)_{c\in\mathcal{C}}$.[11]

Let $S((q_c)_{c\in\mathcal{C}}, \pi^S; \underline{x}_C, \underline{x}_S)$ denote the set of solutions that satisfy the stability constraints defined by $\underline{x}_C$ and $\underline{x}_S$ as well as the feasibility and nonnegativity constraints. A property is *topologically generic* if it holds over a dense open set within the space of π^S and $(q_c)_{c\in\mathcal{C}}$.[12] The following theorem implies that the desired continuity property is topologically generic.

Theorem 1. *For a topologically generic set of* $(q_c)_{c\in\mathcal{C}}, \pi^S \gg \mathbf{0}$ *one of the following is true:*

1. $S((q_c)_{c\in\mathcal{C}}, \pi^S; \underline{x}_C, \underline{x}_S)$ is empty.
2. $S((q_c)_{c\in\mathcal{C}}, \pi^S; \underline{x}_C, \underline{x}_S)$ is non-empty and continuous in $(q_c)_{c\in\mathcal{C}}, \pi^S$.

The full stable set is

$$S((q_c)_{c\in\mathcal{C}}, \pi^S) = \bigcup\nolimits_{\underline{x}_C, \underline{x}_S} S((q_c)_{c\in\mathcal{C}}, \pi^S; \underline{x}_C, \underline{x}_S)$$

I obtain the following corollary since the set of possible $(\underline{x}_C, \underline{x}_S)$ is finite.

Corollary 1. *For generic choices of* $((q_c)_{c\in\mathcal{C}}, \pi^S) \gg \mathbf{0}$, $S((q_c)_{c\in\mathcal{C}}, \pi^S)$ *is continuous in* $(q_c)_{c\in\mathcal{C}}, \pi^S$.

Consider the generic set of $((q_c)_{c\in\mathcal{C}}, \pi^S)$ such that point 1 or 2 of theorem 1 holds. For there not to exist a locally continuous selection at one of these generic points, it must be that $S((q_c)_{c\in\mathcal{C}}, \pi^S; \underline{x}_C, \underline{x}_S)$ satisfies condition 1 for all choices of $(\underline{x}_C, \underline{x}_S)$. However, this implies that $S((q_c)_{c\in\mathcal{C}}, \pi^S)$ is empty in an open set around $((q_c)_{c\in\mathcal{C}}, \pi^S)$, which contradicts my claim that a stable match exists.

[11] In appendix D I provide an alternative discussion of the continuity issue that is constructive. Appendix D provides a modification of the Gale-Shapley algorithm and describes conditions under which the match is generically Lipschitz continuous. One advantage of the constructive approach is that the underlying mechanics of the continuity are clear. The second advantage is that it links the model more tightly with the prior literature, which often uses combinatoric and algorithmic arguments.

[12] We employ the usual topology over Euclidean spaces.

The following example demonstrates what can happen at a nongeneric point of discontinuity. At this point, $S((q_c)_{c\in\mathcal{C}}, \pi^S; \underline{x}_C, \underline{x}_S)$ becomes empty and $S((q_c)_{c\in\mathcal{C}}, \pi^S; \underline{x}'_C, \underline{x}'_S)$ becomes nonempty for another choice of $(\underline{x}'_C, \underline{x}'_S)$.[13]

Example 1. *Consider an economy where* $\mathcal{S} = \{s_1, s_2\}$ *and* $\mathcal{C} = \{c_1, c_2\}$. *Suppose* $c_1 \succ_{s_1} c_2$ *and* $c_2 \succ_{s_2} c_1$, *while* $s_2 \succ_{c_1} s_1$ *and* $s_1 \succ_{c_2} s_2$. *Assume throughout that* $q_{c_1} = q_{c_2} = \frac{1}{2}$. *There are four cases to consider:*

- *If* $\pi^S(s_1), \pi^S(s_2) < \frac{1}{2}$, *then the unique stable match is the student-optimal match* $x(s_1, c_1) = \pi^S(s_1)$ *and* $x(s_2, c_2) = \pi^S(s_2)$.
- *If* $\pi^S(s_1) \geq \frac{1}{2}$ *and* $\pi^S(s_1) + \pi^S(s_2) < 1$, *then the match is* $x(c_1, s_1) = \frac{1}{2}$, $x(c_2, s_1) = \pi^S(s_1) - \frac{1}{2}$, *and* $x(c_2, s_2) = \pi^S(s_2)$
- *If* $\pi^S(s_1) \geq \frac{1}{2}$ *and* $\pi^S(s_1) + \pi^S(s_2) > 1$, *then then the unique stable match is the college-optimal match* $x(s_1, c_2) = x(s_2, c_1) = \frac{1}{2}$ *with the remaining students unmatched.*
- *If* $\pi^S(s_1) \geq \frac{1}{2}$ *and* $\pi^S(s_1) + \pi^S(s_2) = 1$, *then any match wherein all the students are enrolled at some college is stable.*

In other words, lower hemicontinuity of the stable set fails as $\pi^S(s_1) + \pi^S(s_2)$ *crosses* $\frac{1}{2}$.

As mentioned in the literature review, several prior works show that stable matches of a limit model yield approximately stable matches in a finite economy. One can recast example 1 as a demonstration that an approximately stable match need not resemble any stable match. To be concrete, define a match as $\varepsilon - stable$ if the match admits at most a measure ε of students that can form blocking pairs. Then if $\pi^S(s_1) = \pi^S(s_2) = \frac{1}{2} + \frac{\varepsilon}{2}$, the following match is $\varepsilon - stable$:

$$\begin{aligned} x(s_1, c_1) &= x(s_2, c_2) = \frac{1}{2} \\ x(s_1, \varnothing) &= x(s_2, \varnothing) = \frac{\varepsilon}{2} \end{aligned}$$

Since only a measure ε of unmatched students are interested in recontracting following this outcome, the match is indeed $\varepsilon - stable$. However, the exactly stable match,

[13] I would like to thank an anonymous referee of a related paper for this example.

as described in example 1, yields a different outcome for all but the measure ε of unmatched agents in the economy.

Theorem 1 implies that the correspondence of feasible and stable matches is generically continuous, but it is easy to generate a discontinuous selection from the correspondence if the correspondence is not single-valued. Assumption 1 requires that the designer not choose such a discontinuous selection. This assumption is important since it insures that the limit model with a continuum of agents approximates the finite game with many participants if the students declare their types truthfully. To see this, note that the distribution of types in a large market is with high probability close to the exact distribution from which the types are drawn. Assumption 1 implies that when the empirical and exact distributions are close (which is true in large markets with high probability), the match realized in the finite model is nearly the same as the outcome of the limit model with a continuum of agents.

Assumption 1. *$x(c, s; \pi^S)$ is continuous in $(q_c)_{c\in\mathcal{C}}$ and π^S.*

4.4 Strategic Concerns

Assumption 1 implies that if the agents declare their types truthfully, then the outcomes for the finite model with many participants and the continuum model will be very similar with high probability. However, these results are mute on the potential differences in outcomes between these two classes of models when the students act strategically. My goal in this section is to define the game of incomplete information played by the students, describe the relevant equilibrium concepts, and argue that the sets of equilibria in the finite and continuum games are close. With these results in hand, I can use the continuum model to find incentive compatible matches using programming techniques. Throughout I treat the colleges as non-strategic, but it is straightforward to extend the analysis to models with a continuum of strategic colleges.[14]

I model the game played by the students as a revelation mechanism in which the students must declare their types to the mechanism without knowing the types of the other students. I assume that each student of type $s \in \mathcal{S}$ has a von Neumann-Morgenstern representation of his preferences with a felicity function $v_s : \mathcal{C} \to \mathbb{R}$ that represents the student's utility for matching with each college.

[14]The only difficulty is defining the college utility in terms of groups of students.

Since the game is a revelation mechanism, the action space of the game is $\mathcal{S}$. I require that students truthfully declare the verifiable component of their type ($v \in \mathcal{V}$), but the students are free to declare the nonverifiable elements untruthfully ($u \in \mathcal{U}, \succeq_s \in \mathcal{P}_s$). Recalling equation 4.1, which interprets $x(c, s; \pi_N^S)$ as equivalent to a probability distribution over the college to which a student of type $s \in \mathcal{S}$ is assigned, student utility has the form

$$\sum_{c \in \mathcal{C}} \frac{x(c, s; \pi_N^S)}{\pi_N^S(s)} * v_s(c)$$

I focus on symmetric equilibria with the student strategy denoted $\sigma_S : \mathcal{S} \to \Delta(\mathcal{S})$ where $\sigma_S[s](s')$ denotes the probability that a student of type $s \in \mathcal{S}$ declares he is of type s'. The strategy space is denoted Σ. $\Delta_N(X)$ is the set of empirical measures that can be generated by N draws from the set X. Given a realized distribution of student-types in the N-agent game $\pi_E^{S,N}$, the declared distributions of types given the strategy σ_S is $\widehat{\pi}_E^{S,N}(s; \sigma_S) \in \Delta_N(\mathcal{S})$. If the students use pure strategies, $\widehat{\pi}_E^{S,N}(s; \sigma_S)$ can be written

$$\text{For any } s' \in \mathcal{S},\ \widehat{\pi}_E^{S,N}(s'; \sigma_S) = \sum_{s \in \mathcal{S}} \sigma_S[s](s') * \pi_E^{S,N}(s)$$

If the agents employ mixed strategies in the finite game, then the distribution of declared student-types conditional on the realized distribution of types is a random variable defined analogously.

Interim preferences for a student of type $s \in \mathcal{S}$ declaring a type $s' \in \mathcal{S}$ are

$$\mathbb{E}\left[\sum_{c \in \mathcal{C}} \frac{x(c, s'; \widehat{\pi}_E^{S,N})}{\widehat{\pi}_E^{S,N}(s')} * v_s(c)\right] \tag{4.11}$$

where the expectation reflects uncertainty over the realization of the declared types of the other agents (i.e., $\widehat{\pi}_E^{S,N}$). The definition of a Bayes-Nash equilibrium in the finite game is standard.

In the limit game, the strategy σ_S^∞ induces a declared distribution of types $\widehat{\pi}_\infty^S(s'; \sigma_S^\infty)$ defined by

$$\text{For any } s' \in \mathcal{S},\ \widehat{\pi}_\infty^S(s'; \sigma_S^\infty) = \sum_{s \in \mathcal{S}} \sigma_S^\infty[s](s') * \pi^S(s) \tag{4.12}$$

The definition of a Bayes-Nash equilibrium in the limit game is slightly nonstandard in that a deviation by a single student has no effect on the aggregate distribution of declared types.

Definition 1. *The strategy $\sigma_S^\infty : \mathcal{S} \to \Delta(\mathcal{S})$ is an ε-Nash equilibrium (ε-NE) of the limit game if for all types $s \in \mathcal{S}$, any $\widetilde{s} \in supp[\sigma_S(s)]$, and any $s' \in \mathcal{S}$*

$$\sum_{c\in\mathcal{C}} \frac{x(c, \widetilde{s}; \widehat{\pi}_\infty^S)}{\widehat{\pi}_\infty^S(\widetilde{s})} * v_s(c) + \varepsilon \geq \sum_{c\in\mathcal{C}} \frac{x(c, s'; \widehat{\pi}_\infty^S)}{\widehat{\pi}_\infty^S(s')} * v_s(c)$$

Since individual agents cannot influence the declared distribution of types in the limit game, the match is $x(\circ, \circ; \widehat{\pi}_\infty^S)$ with or without the deviation. To compute the effect of a deviation in the finite model, I would have to recompute the match outcome given the deviation. To insure the the outcome following a deviation was incentive compatible, I would have to recompute the effect of all possible deviations from the deviation. Following this logic, writing the incentive constraints in the finite model would require computing the matches generated by *any* possible type declaration, which is impractical. In the limit model, since a deviation by a single agent does not affect the aggregate match, I do not need to compute any other matches when writing down the incentive compatibility constraints, which is the source of the computational tractability of the limit model.

With the definition of an equilibrium of the limit game in place, I can describe the relationship between the equilibria of the finite and limit game. My proof is based on (1) the continuity of the match and (2) the fact that the empirical distribution of student-types in the N agent game is close to π^S with high probability for large N. Together, these facts show that the incentives facing the students in the limit game and a finite game with many players are similar.

Facts (1) and (2) almost immediately imply that a 0-Nash equilibrium of the limit game is an approximate equilibrium in the finite game (and vice versa). However, this does not immediately imply that the exact equilibrium strategies are close to each other. With a bit more work I show that each exact equilibrium strategy of the limit game is close to an exact equilibrium strategy of the finite game and vice versa. Therefore the behavior of the students in the equilibria of the two models are similar. This stronger result is the bedrock motivation for using the convex programs (based on the limit model) that I describe in the next section to analyze matching in the finite economies of the real-world.

Proposition 4. *Consider a 0-Nash equilibrium of the limit game, σ_S^∞. For any $\delta > 0$,* *I can choose N^* so that for $N > N^*$ there is a Bayes-Nash equilibrium of the N agent game, σ_S^N, such that*

$$\left\|\sigma_S^\infty - \sigma_S^N\right\| < \delta$$

Symmetrically, for any $\delta > 0$, I can choose N^* so that for $N > N^*$ and any Bayes-Nash equilibrium of the N agent game, σ_S^N, there is a 0-Nash equilibrium of the limit game, σ_S^∞, such that

$$\left\|\sigma_S^\infty - \sigma_S^N\right\| < \delta$$

5 The Boston High School Match

The following section uses data provided by BPS on the match between rising 9th grade students and public high schools in Boston for the 2011 - 2012 school year. In total I have data on 3479 students that participated in the Gale-Shapley mechanism used by BPS. The data I have on each student includes:

- Preference ranking of up to 10 high school programs that is submitted to the match
- Priority at each of the high school programs ranked
- Final allocation of students to schools
- Special education evaluation (if any)
- English language proficiency ranked on a five point scale
- Zip code of home residence

Throughout the analysis I exclude students that require either special education or English as a second language programs as these students are effectively matched through a different market. Once I exclude these students, I am left with 2,521 students.

The students rank particular programs within each school (e.g., a Chinese-language immersion program at Snowden International high school). Once I have eliminated special education and English proficiency programs, I am left with 28 high school programs of which 10 are overdemanded. The school district divides the seats at each program into equal numbers of *walk-zone* and *open seats*, and students living within a school's walk-zone have higher priority for walk-zone seats than other students. The Gale-Shapley algorithm used by BPS treats the two types of seats within each program as distinct programs.

Students fit into one of either three (open seat program) or six (walk-zone seat program) priority classes as follows

<table>
<tr><th>Walk-Zone Seats</th><th>Open Seats</th></tr>
<tr><td>Guaranteed</td><td rowspan="2">Guaranteed</td></tr>
<tr><td>Sibling-Walk</td></tr>
<tr><td>Walk</td><td rowspan="2">Sibling-Walk, Sibling, NoWalkZoneInGeo</td></tr>
<tr><td>Sibling</td></tr>
<tr><td>NoWalkZoneInGeo</td><td rowspan="2">Walk, No Priority</td></tr>
<tr><td>No Priority</td></tr>
</table>

A student with a *guaranteed* priority at a program are insured to have a place in that program if the student ranks it.[15] *Sibling* and *sibling-walk* apply if the student has a sibling already enrolled in that program. *NoWalkZoneInGeo* is the priority of students that are not in any school's walk-zone. *Walk* applies if the student is in the school's walk zone.

Students that have a priority of guaranteed, sibling, sibling-walk, or NoWalkZoneInGeo at a program are effectively insured admission to the program if they rank it. This insight implies that the effective priority structure can be described as:

<table>
<tr><th>Walk-Zone Seats</th><th>Open Seats</th></tr>
<tr><td>Insured-Walk</td><td>Insured-Open</td></tr>
<tr><td>Walk</td><td rowspan="2">Walk, No Priority</td></tr>
<tr><td>No Priority</td></tr>
</table>

The new priorities *insured-walk* and *insured-open* mean that the student is automatically assigned to the respective walk or open seat program if he or she is not matched to a more preferred school.

Of the 2,521 students I consider, 48.31% of them rank an underdemanded school with empty seats as their most preferred option. These students can be automatically

[15]Students are typically guaranteed a spot within a school if they are currently enrolled in a middle school affiliated with the high school.

assigned to these schools, which means the students are in effect not participating in the match. Another 4.62% of the students rank a school to which they are insured admission as their top choice, and (again) these students are effectively not participating in the match. The statistics I provide below are based on the 47.08% remaining students.

5.1 Incentive Compatibility as a Programming Constraint

In this section I discuss the various forms of incentive compatibility one could impose on the convex programs I study. Before leaping into a discussion of incentive compatibility, I would like to discuss how and when the priority structure can interfere with incentive compatibility. To fix ideas, consider a first-preference-first school choice mechanism wherein each college gives higher priority to students that rank the college as his or her first choice. The original school choice system used by BPS also assigned priorities based on the declared ranking (Abdulkadiroğlu and Sönmez [5]). This system led strategic students that most want to attend a very overdemanded school to improve their outcomes by top-ranking a less popular school program to obtain a higher priority at that program.

Proposition 5 implies that any mechanism that assigns priorities based on the nonverifiable attributes of a student (e.g., the declared preference ranking of colleges) is susceptible to manipulation.

Proposition 5. *The match can be chosen to be incentive compatible for all* $(q_c)_{c\in\mathcal{C}}$ *and* π^S *only if for all pairs of student-types of the form* $s = (v, u, \succeq)$ *and* $\widetilde{s} = (v, \widetilde{u}, \widetilde{\succeq})$ *and colleges* c *we have* $s \sim_c \widetilde{s}$.

Proof. Suppose by contradiction that there exists students $s = (v, u, \succeq)$ and $\widetilde{s} = (v, \widetilde{u}, \widetilde{\succeq})$ and $s \succ_c \widetilde{s}$. Consider an economy where college c is the only college, $q_c = \frac{1}{2}$, and $\pi^S(s) = \pi^S(\widetilde{s}) = \frac{1}{2}$. If the students are truthful, type s occupies all of the seats at college c. All students of type $\widetilde{s}$ would remain unmatched if truthful, but they can obtain a seat in college c by deviating and declaring type $(v, u, \succeq)$. Therefore $\widetilde{s}$ has a strict incentive to make a non-truthful declaration. □

Now I turn to the three notions of incentive compatibility I consider. The most restrictive notion is based on the concept of ordinal dominance. Ordinal dominance incentive compatibility (ODIC) is based on the ordinal ranking of the students' preferences, which makes this notion of incentive compatibility convenient given my data limitations. ODIC is the incentive compatibility condition I use in my convex programs.

Definition 2. *The match $x(c,s)$ satisfies **ordinal dominance incentive compatibility** if for all types $s=(v,u,\succeq_s)\in\mathcal{S}$, $s'\in\mathcal{S}$ and $c\in\mathcal{C}\cup\{\varnothing\}$ we have*

$$\sum_{\{c':c'\succeq_s c\}} x(c',s) \geq \sum_{\{c':c'\succeq_s c\}} x(c',s') \tag{5.1}$$

Given the stringency of the definition, one might question whether any feasible, stable match satisfies the ODIC condition. Helpfully, since any outcome of the Gale-Shapley algorithm with random tie-breakers satisfies ODIC, there must be a match in the limit game that satisfies ODIC. Since ODIC is the strictest incentive compatibility condition I consider, I know that all of the incentive compatibility conditions can be satisfied.

Imposing ODIC constraints allows the designer to avoid taking a stand on the utility functions of the agents. However, if the designer is willing to assume some structure on the students' cardinal utilities, weaker notions of incentive compatibility can be imposed. These weaker notions of incentive compatibility allow more leeway for the designer to achieve his or her objectives.

Definition 3. *The match $x(c,s)$ satisfies **cardinal incentive compatibility** (CIC) if for all types $s\in\mathcal{S}$, $s'\in\mathcal{S}$, and $c\in\mathcal{C}\cup\{\varnothing\}$ we have*

$$\sum_{c\in\mathcal{C}} v_s(c)\frac{x(c,s)}{\pi(s)} \geq \sum_{c\in\mathcal{C}} v_s(c)\frac{x(c,s')}{\pi(s')} \tag{5.2}$$

The third and final notion of incentive compatibility is the weakest and is again based on the idea of ordinal dominance.

Definition 4. *The match $x(c,s)$ satisfies **ordinally undominated incentive compatibility** (OUIC) if for all types $s=(v,u,\succeq_s)\in\mathcal{S}$ there is no $s'\in\mathcal{S}$ such that for all $c\in\mathcal{C}\cup\{\varnothing\}$*

$$\sum_{\{c':c'\succeq_s c\}} x(c',s) \leq \sum_{\{c':c'\succeq_s c\}} x(c',s') \tag{5.3}$$

where equation 5.3 is strict for some $c\in\mathcal{C}\cup\{\varnothing\}$.

The primary benefit of this notion of incentive compatibility is that, like ODIC, it is based on ordinal information. However, OUIC only insures that there are no deviations that can be identified through a dominance argument.

Two interpretations of the OUIC condition can be provided. First, any match that satisfies OUIC is justifiable in the sense that there exists some $v_s(c)$ for which CIC holds. Second, if truthfulness is ordinally dominated, then an obviously better (i.e., ordinally dominating) action can be found. Therefore one might expect nontruthful behavior in mechanisms that do not satisfy OUIC even if the agents are not sophisticated. Stated differently, if agents are not too sophisticated, then OUIC might be enough to insure that profitable deviations are hard to find and truthful behavior is adopted. As demonstrated in the following example, the failure of incentive compatibility in the Boston mechanism can be framed in terms of truthfulness being ordinally dominated.

Example 2. *Consider two colleges, c_1 and c_2, and three types of students, $s_1 = (v_1, u, c_1 \succ c_2), s_2 = (v_2, u, c_1 \succ c_2)$, and $s_3 = (v_3, u, c_2 \succ c_1)$. Assume there is a measure 1 of seats at each college and that each student-type has measure 1. To discuss the deviation, I need to introduce a new type $s_2' = (v_2, u, c_2 \succ c_1)$ that does not exist in the student population. Finally, given my use of the Boston mechanism, I define college priorities as*

College	*Preferences*
c_1	$s_1 \succ_{c_1} s_2 \succ_{c_1} s_2' \succ_{c_1} s_3$
c_2	$s_2' \succ_{c_2} s_3 \succ_{c_2} s_1 \succ_{c_2} s_2$

Importantly, the preference of college c_2 for s_2' over s_2 is only due to the fact that student-type s_2' declares that c_2 is his or her most preferred college.

If the agents declare truthfully, then college c_1 (c_2) accepts students of type s_1 (s_3) and students of type s_2 are unmatched. If students of type s_2 had instead declared type s_2', then college c_1 (c_2) accepts students of type s_1 (s_2') and students of type s_3 are unmatched. Declaring truthfully is ordinally dominated by declaring s_2' for students of type s_2.

5.2 Welfare Maximization

My first objective is to find the welfare maximizing match, compare the outcome to the result generated by the Gale-Shapley algorithm, and then assess the welfare losses caused by imposing incentive compatibility constraints on top of stability and feasibility constraints. For details on the convex programs used to compute these numbers, please see appendix C.

I describe my welfare objective in three steps. First, $|\{c' : c' \succeq_s c\}|$ is the number of schools that student s strictly prefers to college c, which is equivalent to the rank of c in the preference list of s. The expected rank of the college to which students of type s are assigned is

$$\sum_{c \in \mathcal{C}} x(c, s) * |\{c' : c' \succeq_s c\}|$$

My welfare objective is the average expected rank produced by the match

$$R(x) = \sum_{s \in \mathcal{S}} \pi^S(s) \sum_{c \in \mathcal{C}} x(c, s) * |\{c' : c' \succeq_s c\}|$$

Maximizing welfare involves minimizing the average expected rank. Any match that minimizes $R(x)$ will be Pareto optimal, but the reverse need not be true. Throughout I refer to $R(x)$ as the average welfare generated by match x.

The following table describes the average welfare of the best and worse stable and feasible matches; the best and worst stable, ODIC, and feasible matches; and the match produced by the Gale-Shapley algorithm.[16] Recall that the welfare measure employed is the average preference rank of the school to which the students are assigned for the 47.08% of the students that effectively participate in the market.[17]

Match Welfare

Best Stable Match	Best Stable and ODIC match	Gale-Shapley	Worst Stable and ODIC match	Worst Stable Match
2.2492	2.5960	2.6151	2.6852	3.0826

As the table shows, Gale-Shapley recovers 78.6% percent of the welfare gap between the best and worst stable and ODIC matches. However, the difference between the best and worst stable (and non-ODIC) matches is over 9 times larger than the difference between the best and worst stable and ODIC matches. In conclusion, the

[16] I compute the Gale-Shapley outcome in a limit economy using the techniques developed by Azevedo and Leshno [8].

[17] Although not reported, I have solved the welfare maximization problem applying the ODIC constraints without the stability constraints. The maximum and minimum welfare achievable are within 1% of the welfare achieved by applying all of the constraints.

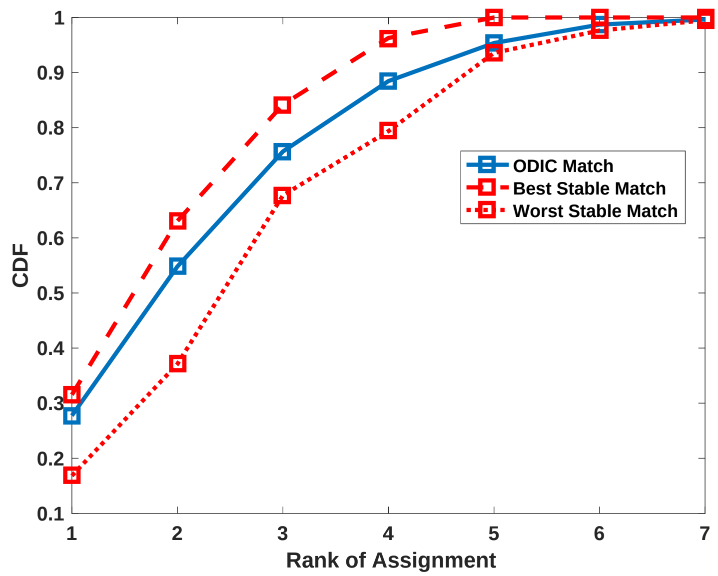

Figure 1: Rank Order Anaysis

ODIC constraints severely limit my ability to improve the welfare of the students, and it does not matter much which stable and ODIC match one implements.[18]

For completeness, figure 1 provides the distribution of the rank assignments (i.e., the percentage of students assigned to their top ranked program, the percentage of students assigned to their second ranked program or better, etc.). To illustrate the effects of the ODIC constraints, I have included the best and worst stable (but not ODIC) matches as well.

Finally, as mentioned in the introduction, one can use the shadow prices of the constraints of the program to evaluate how to allocate future resources to maximize welfare. As a concrete example, the shadow prices on the capacity constraints represent the marginal value (in terms of student welfare) of seats in each school. The highest shadow prices are associated with the four foreign language programs offered by Snowden International. This means that if there is money for an extra seat available, welfare would be maximized by adding that seat to one of Snowden International's programs. Moreover, the difference is significant - the marginal value of an extra seat at Snowden International is more than 84% larger than at other overdemanded schools.

[18]In section 5.3 I condition the ODIC constraints on the students' zip codes, which slackens the constraints and allows us to achieve an average expected rank of under 2.5.

5.3 Diversity

The data provided by BPS does not include any demographic information immediately relevant to ethnicity other than the students' English language proficiency. However, I can compare the zip code of the student's home residence to the demographics collected for the 10-17 year old age group in the 2010 Census of permanent residents of these zip codes. I find that on average this age group is 23.8% Caucasian with zip code rates running from as low as 1% in the 02120 zip code in the Mission Hill neighborhood[19] to as high as 91% Caucasian in the Beacon Hill neighborhood.[20]

Throughout this section I treat the zip code of a student's home residence as a verifiable component of the student's type, which means that the objective function and constraints of the program (as well as the solution) will be conditioned on the students' zip codes. This has the effect of slackening the ODIC constraints. To define my objective function, first I need to compute the fraction of Caucasian students enrolled at school c for a given match x

$$\sum_s p_s \frac{x(c,s)}{q_c}$$

where p_S is the percentage of Caucasian school-age children in the zip code of student-type s. Letting p_{Boston} denote the percentage of school-age children that are Caucasian in Boston as a whole, I use the following measure of diversity within a given school:

$$\left[\sum_s p_s \frac{x(c,s)}{q_c} - p_{Boston}\right]^2$$

In other words, I compute a proxy for the square distance between the percentage of Caucasian children in a school and the average over Boston as a whole. The average

[19] This neighborhood, which is very popular with students in their 20s, has a relatively low white population in the pre-college population.

[20] Demographic data from the 2015-2016 school year indicate that just 14.2% of BPS students are Caucasian, which is understandable if wealthier Caucasian students are prone to enrolling in private schools. However, we have no reason to believe that the qualitative message of this section is affected by this discrepancy (Demographic data drawn from http://profiles.doe.mass.edu/state_report/enrollmentbyracegender.aspx on 17 August 2016).

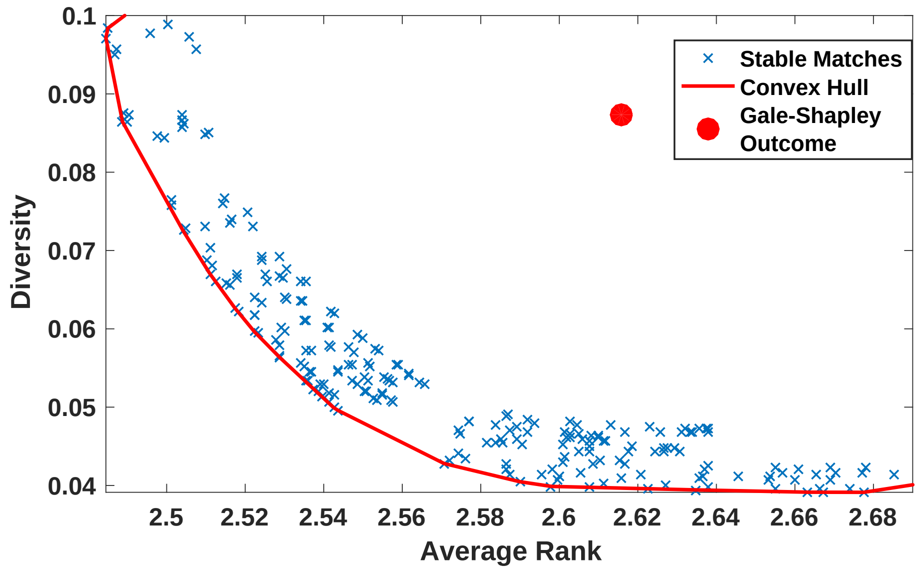

Figure 2: Welfare - Diversity Frontier

diversity across all of the BPS high schools is

$$\frac{1}{\sum_c q_c} \sum_c q_c \left[\sum_s p_s \frac{x(c,s)}{q_c} - p_{Boston} \right]^2 \tag{5.4}$$

Lower values of equation 5.4 imply that the population at each school better mirrors the demographic distribution of Boston as a whole.

To understand the trade-offs between average welfare and diversity, I seek to minimize the following objective function:

$$\alpha * R(x) + (1-\alpha) * \frac{1}{\sum_c q_c} \sum_c q_c \left[\sum_s p_s \frac{x(c,s)}{q_c} - p_{Boston} \right]^2$$

where $R(x)$ is the average expected rank of the school assignment x. In other words, the objective function is a weighted average of the welfare and diversity criteria. I respect the capacity and stability constraints throughout my analysis. Because of this, I can only influence the diversity at overdemanded schools since moving students between underdemanded schools will immediately create a blocking pair. Since my match needs to condition on the students' zip codes, my ODIC constraints must also condition on the zip codes.

Figure 2 presents the welfare-diversity possibility frontier generated by varying

$\alpha \in [0, 1]$. I have, for expositional clarity, used the following absolute difference criterion as my metric of diversity in the figure.[21]

$$\frac{1}{\sum_c q_c} \sum_c q_c \left\| \sum_s p_s \frac{x(c,s)}{q_c} - p_{Boston} \right\|$$

Each "x" mark represents a stable, feasible and ODIC match that is optimal with respect to my objective for some value of α. The dot represents the Gale-Shapley outcome, which is within the frontier and lacking in diversity relative to many of the matches I compute. The frontier, the thick line, is the convex hull of the welfare-diversity possibility frontier. Finally, since my incentive compatibility conditions are slackened by conditioning on the students' zip codes, I am able to achieve a slightly lower average rank than the welfare results presented in section 5.2.

As suggested by the welfare section, the choice of which match to implement has a small effect on the welfare of the students. One might have assumed that other objectives might be similarly constrained by the ODIC constraints. My analysis here reveals that, in fact, the choice of which match to implement has a large effect on the average diversity of the schools, while at the same time having little effect on the student welfare.

5.4 Neighborhood Schools

The origin of the walk-zone priority system was a desire to strengthen neighborhood schools. In this section I assess the trade-off between student welfare and having schools that disproportionately serve the student population in the school's walk-zone. In lieu of a priority scheme, I directly encode the neighborhood school requirement as a constraint in my program. I insist that a fraction α of seats within each school be assigned to students within the walk-zone. My results indicate that there is essentially no welfare cost to requiring that up to $\alpha = 50\%$ of the student body be drawn from the walk-zone.

I let $\mathcal{W}(c)$ denote the set of student-types within school c's walk-zone, which means

[21] While I would have liked to use the absolute difference criterion in the program, it is computationally easier to solve a quadratic program.

I can write my walk-zone constraint as follows:

$$\text{For all schools } c, \quad \sum_{s \in \mathcal{W}(c)} x(c, s) \geq \alpha q_c \tag{5.5}$$

To compute the trade-off between neighborhood schools and student welfare, I solved a program where I maximize welfare subject to:

- ODIC constraints
- Feasibility constraints
- Stability constraints
- At least a fraction α of the assigned students be drawn from each schools' walk-zone

Since I do not impose a priority structure on the market, the stability constraints only require that I cannot match a student of type s to a college c when the student would prefer to be matched to an underdemanded school. However, this limited stability requirement can cause infeasibilities. Some schools may be overdemanded, but not sufficiently popular within the set of walk-zone students to attract at least a fraction α of their student body from the walk-zone. My analysis reveals that two schools cannot attract more than 9% of their students from their walk-zones, while a third school cannot attract more than 20% of its students from its walk-zone.

To handle this issue I solve the problem starting with values of α ranging from 0 to 0.5. When I hit a feasibility problem for college c, I cap the neighborhood school requirement at that level and continue to raise α beyond that point for the remaining schools. For example, when I compute $\alpha = 0.15$, I require that the schools that can attract at most 9% of their students from their walk-zone attract 9%, while the other schools are required to attract at least 15%.[22]

The trade-off between α, the percentage of students drawn from the walk-zone, and student welfare is plotted in figure 3. Since the ODIC constraints are the primary barriers to welfare maximization, the welfare generated by the match is not very

[22] Perhaps unsurprisingly, the three school programs that cannot attract 50% of their students from their walk-zones are the same three school programs where non-walk-zone students occupy walk-zone slots in the status quo match.

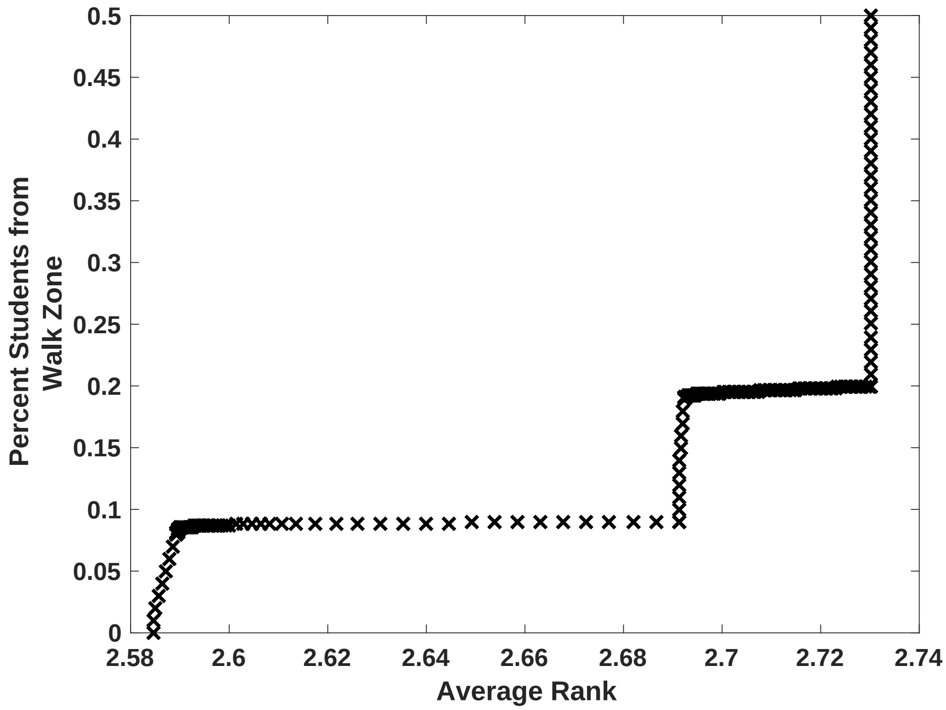

Figure 3: Welfare - Walk Zone Frontier

sensitive to α, the fraction of walk-zone students.[23] In summary, I do not see a strong trade-off between welfare and the encouragement of neighborhood schools by requiring a particular level of walk-zone enrollment.

6 Conclusion

In this paper I introduce a novel analysis tool for school choice problems based on convex programs that assign a continuum of students to a continuum of college seats. By assuming that I have a continuum of agents on each side of the market, I am able to describe the feasibility, stability and incentive compatibility constraints in a computationally tractable fashion. I provide a characterization of the stable set of the model and prove that the match can be chosen to be continuous in the distribution of preferences declared to the mechanism and the school capacities. I use my continuity result to prove that the game-theoretic equilibria of the continuum and finite models

[23]The points at which the average expected rank appears to rapidly increase are due to equation 5.5 causing a near infeasibility in the program.

approach each other as the number of students in the finite game grows. This final result implies that tractable convex programs based on the continuum model can be used to describe matches for finite, real-world school choice settings.

The main advantage of the programming approach is that market designers can explicitly encode their goals into an objective function, and the program's solution is ensured to achieve the global optimum subject to the feasibility, stability, and incentive compatibility constraints of the problem. Current matching mechanisms often use priority structures to nudge the mechanism towards desired outcomes, in effect trying to use constraints to achieve the designer's objectives. In contrast, the programming approach transparently distinguishes between objectives and constraints. An additional advantage of transparently delineating objectives and constraints is that I am able to analyze the trade-offs between the designer's goals and use the tools of convex programming to study the benefits of relaxing the constraints by (for example) increasing a school's capacity.

In the context of the BPS high school match, numerical solutions to my convex programs reveal that the ODIC constraints severely constrain the welfare that can be achieved. This implies that the Gale-Shapley algorithm attains essentially the highest welfare possible. On the other hand, the set of stable and ODIC matches admit significant variation in the diversity of the student body and the fraction of the student body drawn from each school's walk-zone. This implies that while there may be little that can be done to improve the average expected rank of the assignment, the choice of which stable and incentive compatible match to implement can have significant effects on goals related to the composition of the student body at each school without significantly affecting student welfare.

References

[1] Abdulkadiroğlu, A.; Y-K Che; and Y. Yasuda (2011) "Resolving Conflicting Preferences in School Choice: The Boston Mechanism Reconsidered," *American Economic Review*, 101, pp. 399-410.

[2] Abdulkadiroğlu, A.; P. Pathak; and A. Roth (2005) "The NYC High School Match," *American Economic Review*, 95 (2), pp. 364 - 367.

[3] Abdulkadiroğlu, A.; P. Pathak; A. Roth; and T. Sönmez (2005) "The Boston Public School Match," *American Economic Review*, 95 (2), pp. 368 - 371

[4] Abdulkadiroğlu, A.; P. Pathak; and A. Roth (2009) "Strategy-proofness versus Efficiency in Matching with Indifferences: Redesigning the NYC High School Match," *American Economic Review*, 99, pp. 1954 - 1978.

[5] Abdulkadiroğlu, A. and T. Sönmez (2003) "School Choice: A mechanism design approach, "*American Economic Review*, 93 (3), pp. 1954 - 1978.

[6] Ashlagi, I. and P. Shi (2014) "Optimal Allocation without Money: an Engineering Approach," *Management Science*, 62 (4), pp. 1078 - 1097.

[7] Azevedo, E. and J. Hatfield (2012) "Complementarity and Multidimensional Heterogeneity in Matching Markets," *mimeo.*

[8] Azevedo, E. and J. Leshno (Forthcoming) "A Supply and Demand Framework for Two-Sided Matching Markets ," *Journal of Political Economy.*

[9] Azevedo, E. and J. Leshno (2012) "Complementarity and Multidimensional Heterogeneity in Matching Markets," *mimeo.*

[10] Baïou, M. and M. Balinski (2000) "The stable admissions polytope," *Mathematical Programming*, 87, pp. 427 - 439.

[11] Bodoh-Creed, A. (2013) "Efficiency and information aggregation in large uniform-price auctions," *Journal of Economic Theory*, 148 (6), pp. 2436 - 2466.

[12] Budish, E.; Y-K Che; F Kojima; and P. Milgrom (2013) "Designing Random Allocation Mechanisms: Theory and Applications," *American Economic Review*, 103, pp. 585-623.

[13] Che, Y-K; J. Kim; and F. Kojima (2014) "Stable Matching in Large Economies," *mimeo.*

[14] Chen, Y. and O. Kesten (Forthcoming) "Chinese College Admissions and School Choice Reforms: A Theoretical Analysis," *Journal of Political Economy.*

[15] Curtis, C. (1984) *Linear Algebra*, Springer-Verlag: New York.

[16] Dur, U.; S. Kominers; P. Pathak; and T. Sönmez (2014) “The Demise of Walk Zones in Boston: Priorities vs. Precedence in School Choice,” *mimeo.*

[17] Echenique, F.; S. Lee; M. Shum; and M.B. Yenmez (2013) “The Revealed Preference Theory of Stable And Extremal Stable Matchings,” *Econometrica*, 83, pp. 153 - 171.

[18] Echenique, F and M.B. Yenmez (2014) “How to Control Controlled School Choice,” *The American Economic Review*, 105 (8), pp. 2679 - 2694.

[19] Embrechts, P.; C. Klüppelberg; and T. Mikosch (1997) *Modelling Extremal Events for Insurance and Finance*, Springer: New York.

[20] Erdil, A. and H. Ergin (2006) “What’s the Matter with Tie-Breaking? Improving Efficiency in School Choice,” *American Economic Review*, 98, pp. 669-689.

[21] Featherstone, C. (2006) “What’s the Matter with Tie-Breaking? Improving Efficiency in School Choice,” *American Economic Review*, 98, pp. 669-689.

[22] Featherstone, C. and M. Niederle (2013) “Improving on Strategy-proof School Choice Mechanisms: An Experimental Investigation,” *mimeo.*

[23] Gale, D. and L.S. Shapley (1962) ”College Admissions and the Stability of Marriage,” *The American Mathematical Monthly*, 69, pp. 9-15

[24] Haeringer, G. and F. Klijn (2009) ”Constrained School Choice,” *Journal of Economics Theory*, 144, pp. 1921-1947.

[25] Halafir, I.; M. Yenmez; and M. Yildirim (2013) “Effective affirmative action in school choice,” *Theoretical Economics*, 8, pp. 325-363.

[26] Hatfield, J. and Milgrom, P. (2005) ”Matching with Contracts,” *The American Economic Review*, 95 (4), pp. 913-935.

[27] Hitsch, G.; A. Hortaçsu; and D. Ariely (2010) “Matching and Sorting in Online Dating,” *The American Economic Review*, 100 (1), pp. 130-163.

[28] Immorlica, N. and M. Mahdian (2005) “Marriage, Honesty, and Stability,” *SODA 2005*, pp. 53-62.

[29] Kelso, A. and V. Crawford (1982) "Job Matching, Coalition Formation, and Gross Substitutes," *Econometrica*, 50 (6), pp. 1483-1504.

[30] Kesten, O. (2010), "School Choice with Consent," *The Quarterly Journal of Economics*, 125 (3), pp. 1297 - 1348.

[31] Kesten, O. and U. Unver (2015) "A theory of school-choice lotteries," *Theoretical Economics*, 10, pp. 543 - 595.

[32] Kojima, F. and P. Pathak (2009) "Incentives and Stability in Large Two Sided Matching Markets," *American Economic Review*, 99 (3) pp. 608 - 627.

[33] Kominers, S. and T. Sönmez (2016) "Matching with Slot-Specific Priorities: Theory," *Theoretical Economics*, 11 (2), pp. 683-710

[34] Kozlov, M.; S. Tarasov; and L. Khachiyan (1980) "The polynomial solvability of convex quadratic programming," *USSR Computational Mathematics and Mathematical Physics,* 20 (5), pp. 223 - 228.

[35] Lee, S. (2012) "Incentive Compatibility of Large Centralized Matching Markets," *mimeo.*

[36] Menzel, Konrad (2015) "Large Matching Markets as Two-Sided Demand Systems," *Econometrica*, 83 (3) pp. 897 - 941.

[37] Niederle, M.; A. Roth; U. Unver (2013) "Unraveling Results From Comparable Demand and Supply: An Experimental Investigation," *Games and Economic Behavior*, 4 (2), pp. 243 - 282.

[38] P. Pathak and T. Sönmez (2013) "School Admissions Reform in Chicago and England: Comparing Mechanisms by Their Vulnerability to Manipulation," *American Economic Review*, 103 (1), pp. 80-106.

[39] Roth, A. (1984) "The Evolution of the Labor Market for Medical Interns and Residents: A Case Study in Game Theory, "*Journal of Political Economy*, 92 (6), pp. 991 - 1016.

[40] Roth, A. and E. Peranson (1999) "The Redesign of the Matching Market for American Physicians: Some Engineering Aspects of Market Design," *The American Economic Review,* 89 (4), pp. 748 - 780.

[41] Roth, A.; U. Rothblum; and J. Vande Vate (1993) "Stable Matchings, Optimal Assignments, and Linear Programming," *Mathematics of Operations Research*, 18, pp. 803-828.

[42] Roth, A.; T. Sönmez; and M.U. Unver (2007) "Efficient Kidney Exchange: Coincidence of Wants in Markets with Compatability-Based Preferences," *The American Economic Review,* 97 (3), pp. 828-851.

[43] Roth, A. and M. Sotomayor (1990) *Two Sided Matching: a Study in Game-Theoretic Modeling and Analysis*, Cambridge University Press: Cambridge, England.

[44] Roth, A. and J. Vande Vate (1990) "Random Paths to Stability in Two-Sided Matching," *Econometrica,* 58 (6), pp. 1475 - 1480.

[45] Roth, A. and J. Vande Vate (1991) "Incentives in Two-Sided Matching with Random Stable Matches," *Economic Theory,* 1, pp. 31 - 44.

[46] Sönmez, T. and T. Switzer (2013) "Matching with (Branch-of-Choice) Contracts at the United States Military Academy" *Econometrica,* 81, pp. 451-488.

[47] van der Vaart, A. and J. Wellner (1996) *Weak Convergence and Empirical Processes*, Springer-Verlag: New York.

A Proofs

Lemma 1. *Equation 4.4 holds for a pair* (c, s) *if and only if* $c \succ_s \underline{x}_S(s)$ *and* $s \succ_c \underline{x}_C(c)$

Proof. Note that if $c \succ_s \underline{x}_S(s)$, then there must be a positive mass of students either unmatched or matched to colleges to which c is strictly preferred by s. Therefore

$$\sum_{\{c' : c' \succeq_s c\}} x(c', s) < 1.$$

Similarly, if $s \succ_c \underline{x}_C(c)$, then there must be a positive mass of college seats that are either unoccupied or occupied by students that have lower priority at college c than s. Therefore $\sum_{\{s' : s' \succeq_c s\}} x(c, s') < q_c$. The logic of the reverse direction is similar (and omitted). $\square$

Proposition 1. *Equations 4.7, 4.8, 4.9, and 4.10 are necessary and sufficient for a match x to be feasible, stable, and individually rational.*

Proof. Equations 4.7 and 4.8 are required feasibility constraints and the if-and-only-if argument for feasibility is immediate. Similarly individual rationality is (by definition) satisfied if-and-only-if equation 4.10 holds. Equation 4.9 is a compressed formulation of my definition of stability. □

Proposition 2. *There is at least one feasible and stable match.*

Proof. This result is known in the case with a finite set of agents through a constructive argument based on the Gale-Shapley deferred acceptance algorithm. I extend this result to the continuum case in two steps. First I prove the theorem when the measure of the seats at each college is a rational number and the measure of each type of agent is a rational number, which is the natural extension of the finite agent case. Second, I use the compactness of the space of matches to extend this result to the case where distributions of student-types or college capacities take irrational values.

Lemma 2. *When the college capacities and the measure of each student-type take on rational values, a feasible and stable match exists.*

Proof. Let $m > 0$ be the smallest real number such that for any $s \in \mathcal{S}$ ($c \in \mathcal{C}$) I have $m\pi^S(s)$ (mq_c) is an integer. Such a finite m must exist since I have assumed that $\pi^S(s)$ and q_c are rational numbers. Treating each m^{-1} measure of students and college seats as atomistic agents and running the usual Gale-Shapley algorithm on this finite economy yields a stable match. If I define $x(c,s) = m * \#\{(c,s)$ matches in the finite economy$\}$, I obtain a feasible and stable match in the continuum model. □

Now consider a continuum economy where $\pi^S(s)$ or q_C takes on an irrational value for some s or c. Since the rational numbers are dense in the real numbers, I can construct a sequence $\{\pi_n^S, (q_C^n)_{c\in\mathcal{C}}\}_{n=1}^{\infty}$ such that $\pi_n^S(s)$ and $(q_C^n)_{c\in\mathcal{C}}$ assume rational values and $\pi_n^S \to \pi^S$ and $q_C^n \to q_C$. Consider an associated sequence $\{x_n\}_{n=1}^{\infty}$, where x_n is a feasible and stable match in the $(\pi_n^S, (q_C^n)_{c\in\mathcal{C}})$ economy. Recall that $x \in [0,1]^{(|\mathcal{C}|+1)(|\mathcal{S}|+1)}$ and note that this space is compact. From the compactness of the space, there exists a convergent subsequence of $\{x_n\}_{n=N}^{\infty}$ with a limit x_∞. Since the feasibility and stability equations are equalities, x_∞ must satisfy these equations for the $(\pi^S, (q_C)_{c\in\mathcal{C}})$ economy. Therefore x_∞ is a feasible and stable match for $(\pi^S, (q_C)_{c\in\mathcal{C}})$. □

Proposition 3. *Any assignment satisfying equations 4.7 through 4.10 can be implemented.*

Proof. I rely on theorem 1 of Budish et al. [12]. From this result, it suffices to argue that the constraints defining my match are a bihierarchy. The capacity constraints on students and colleges and the individual rationality constraints form a bihierarchy. What remains are the stability constraints, which are difficult to interpret in the context of Budish et al. [12] since they do not fit into the linear constraint structure studied therein. However, for any choice of $\underline{x}_S$ and $\underline{x}_C$ the stability constraints can be written as a family of constraints as follows:

$$\text{For all } s, \quad \sum_{\{c' \in \mathcal{C}: c' \succeq_s \underline{x}_C(s)\}} x(c', s) = \pi^S(s) \tag{A.1}$$

$$\text{For all } c, \quad \sum_{\{s' \in \mathcal{S}: s' \succeq_c \underline{x}_S(c)\}} x(c, s') = q_c \tag{A.2}$$

where I have implicitly used $x(c, \varnothing)$ and $x(\varnothing, s)$ to denote unmatched agents. When written in this way, it is clear that each equation involves a sum over a subset of (c, s) pairs that appear in the capacity constraint for some student or college. So equations 4.7 and A.1 form a hierarchy, while equations 4.8 and A.2 form a second hierarchy. The individual rationality constraints mandate that certain cells of $\mathbf{X}$ be 0, which implies that the corresponding cells of each of $\{\mathbf{X_i}\}_{i=1}^{A}$ also be zero. Since each $\{\mathbf{X_i}\}_{i=1}^{A}$ is weakly positive, these requirements do not affect my ability to implement $\{\mathbf{X}\}$. With these difficulties resolved, theorem 1 of Budish et al. [12] implies $\mathbf{X}$ can be implemented. □

Proposition 4. *Consider a 0-Nash equilibrium of the continuum game, σ_S^∞ . For any $\delta > 0$, I can choose N^* so that for $N > N^*$ there exists a Bayes-Nash equilibrium of the N agent game, σ_S^N, such that*

$$\left\|\sigma_S^\infty - \sigma_S^N\right\| < \delta$$

Symmetrically, for any $\delta > 0$, I can choose N^* so that for $N > N^*$ and any Bayes-Nash equilibrium of the N agent game, σ_S^N, there exists a 0-Nash equilibrium of the continuum game, σ_S^∞, such that

$$\left\|\sigma_S^\infty - \sigma_S^N\right\| < \delta$$

Proof. From the continuity of the mechanism (Assumption 1) and the linearity of the student utilities in probability (equation 4.11), I know that agent utility is continuous in π^S and $(q_c)_{c\in\mathcal{C}}$. From theorems 8 and 9 of Bodoh-Creed [11] I know that for any $\varepsilon > 0$ I can choose N^* so that for $N > N^*$ I have:

- Any Bayes-Nash equilibrium of the N-agent game is an $\varepsilon-$Nash equilibrium of the continuum game
- Any 0-Nash equilibrium of the continuum game is an $\varepsilon-$Bayes-Nash equilibrium of the N-agent game

I now prove that for any $\delta > 0$ I can find N^* so that for $N > N^*$ and any Bayes-Nash equilibrium of the $N-$agent game I can find a 0-Nash equilibrium of the continuum game. The argument for the reverse direction is very similar (and hence omitted).

First I define the following equation for payoffs in the limit game

$$U(a, \sigma, s) = \sum_{c\in\mathcal{C}} \frac{x(c, a; \widehat{\pi}^S_\infty)}{\widehat{\pi}^S_\infty(a)} v_s(s)$$

where $a \in \mathcal{S}$ is a type-declaration allowed by a student of type s and $\widehat{\pi}^S_\infty$ is generated by equation 4.12. $U(a, \sigma, s)$ is continuous in (a, σ, s) since a and s are drawn from discrete sets, x is continuous with respect to $\widehat{\pi}^S$, and $\widehat{\pi}^S$ is continuous with respect to σ.

Given the continuity, define the benefit to a student of deviating from a candidate strategy σ as

$$r(\sigma) = \max_{a,s\in\mathcal{S}} [U(a, \sigma, s) - U(\sigma(s), \sigma, s)]$$

Given this definition, I also define a set-valued inverse of r as

$$r^{-1}(\varepsilon) = \{\sigma : r(\sigma) \leq \varepsilon\}$$

where $R = r^{-1}(0)$ is the set of exact equilibria of the continuum game.

I can describe the set of strategies that allow more than a γ-deviation as $B_\gamma = r^{-1}((\gamma, \infty))$. Since r is continuous, B_γ is open and I can write

$$R = \bigcap_{k=1}^{\infty} B_{\frac{1}{k}}$$

The set of strategies that are close to an equilibrium is

$$C_\delta = \{\sigma \in \Sigma : \text{ for some } \sigma^* \in R \text{ I have } \|\sigma^* - \sigma\| < \delta\}$$

which is open, so $\Sigma - C_\delta$ is compact. Since $\{B_{\frac{1}{k}} : k \in \mathbb{N}\}$ is an open cover of $\Sigma - C_\delta$, there is a finite cover of $\Sigma - C_\delta$. Let the largest element of this cover (in which all other elements of the cover are nested) be denoted B_ε, and I have $\Sigma - C_\delta \subset B_\varepsilon$. Consider $\sigma \notin B_\varepsilon$, in which case σ is an $\varepsilon-$Nash equilibrium of the continuum game. Since $\sigma \notin B_\varepsilon$ implies $\sigma \in C_\delta$, there is some $\sigma^* \in R$ so that $\|\sigma^* - \sigma\| < \delta$. From theorem 9 of Bodoh-Creed [11] I know that for any $\varepsilon > 0$ I can choose N^* so that for $N > N^*$ any Bayes-Nash equilibrium σ^N of the N-agent game is an $\frac{\varepsilon}{2}-$Nash equilibrium of the continuum game (i.e., $\sigma^N \notin B_\varepsilon$). Together, this yields the desired conclusion. □

B Appendix: Continuity Proofs

The first step is to consider an arbitrary $\underline{x}_S(s)$ and $\underline{x}_C(c)$ that satisfy the stability constraints of lemma 1. Given such a choice, I can write the stability constraints as

$$\begin{aligned}
\text{For all } s, \sum_{c \in \{c' : c \succeq_s \underline{x}_S(s)\}} x(c,s) &= \pi^S(s) \\
\text{For all } c, \sum_{s \in \{s' : s \succeq_c \underline{x}_C(c)\}} x(c,s) &= q_c
\end{aligned}$$

I now define a matrix A that represents the capacity and stability constraints of my problem. There must be one row for each college and student-type, for total of $|\mathcal{C}| + |\mathcal{S}|$ rows. The matrix must have one column for each (college,student-type) pair and a column representing "matching" the agent represented in the row with $\varnothing$ (i.e., the agent is left unmatched). I define A as follows, where I use s denotes a generic student-type and c denotes a generic college. Unless stated otherwise, the entry of A is 0.

$$\begin{aligned}
A(c,(c,s)) &= A(s,(c,s)) = 1 \text{ if and only if } s \succeq_c \underline{x}_C(c) \text{ and } c \succeq_s \underline{x}_S(s) \\
A(c,(c,\varnothing)) &= 1 \text{ if and only if } \varnothing \in \underline{x}_C(c) \\
A(s,(\varnothing,s)) &= 1 \text{ if and only if } \varnothing \in \underline{x}_S(s)
\end{aligned}$$

I treat $x(c,s)$ as a column vector, denoted $\mathbf{x}$, ordered in the same fashion as the columns of A. I can then write my set of feasible and stable matches as

$$A\mathbf{y} = \mathbf{q} \tag{B.1}$$

where $\mathbf{q}$ is a $|\mathcal{C}| + |\mathcal{S}|$ element vector ordered as the rows in A with the values

$$\begin{aligned} \mathbf{q}(c) &= q_C \\ \mathbf{q}(s) &= \pi^S(s) \end{aligned} \tag{B.2}$$

Note that there is a bijection between $\mathbf{q}$ and $(q_c)_{c\in\mathcal{C}}$ and π^S, so statements about the genericity of $\mathbf{q}$ and $((q_c)_{c\in\mathcal{C}}, \pi^S)$ are equivalent.

I can use the formulation of the set of feasible and stable matches combined with elementary results from the theory of solutions to systems of equations to obtain a number of results. Denote the set of solutions to equation B.1 as $\mathbf{Y}(\mathbf{q}; \underline{x}_C, \underline{x}_S)$ to reflect the fact that equation B.1 is defined by the choice of $\underline{x}_C, \underline{x}_S$. Since $\mathbf{q}$ is a function of $(q_c)_{c\in\mathcal{C}}$ and π^S, $\mathbf{Y}$ is implicitly a function of the underlying distributions of student-types and college capacities.

Lemma 3. *The following are true*

1. *If A does not have full rank , then for a topologically generic set of $\mathbf{q}$ there is no match that satisfies the feasibility and stability constraints.*

2. *If A has full rank, then $\mathbf{Y}(\mathbf{q}; \underline{x}_C, \underline{x}_S)$ is a convex valued, continuous correspondence with respect to $\mathbf{q}$.*

Proof. Claim 1 follows from theorem 8.8 of Curtis [15], which states that equation B.1 has a solution if and only if the rank of A equal the rank of $[A\ \mathbf{q}]$. This condition holds if A has full rank, and theorem 8.8 of Curtis [15] implies the existence of a solution. Since the system is linear, the set of solutions must be convex. If A does not have full rank, then it the rank of A is strictly less than the rank of $[A\ \mathbf{q}]$ (i.e., no solution exists) for a topologically generic set of $\mathbf{q}$

Now I prove continuity of the correspondence of solutions. First note that upper hemicontinuity follows from my definition of equation B.1 as a system of linear equalities. What remains is to show lower hemicontinuity, which requires that if $A\mathbf{y} = \mathbf{q}$, then for any sequence $(\mathbf{q}_i)_{i=1}^{\infty}$ where $\mathbf{q}_i \to \mathbf{q}$ I can choose $\mathbf{y}_i$ such that $A\mathbf{y}_i = \mathbf{q}_i$ and

$\mathbf{y}_i \to \mathbf{y}$. This amounts to showing that for any $\gamma > 0$ I can choose $\eta > 0$ so that if $\|\mathbf{q}_i - \mathbf{q}\| < \eta$ I can find $\|\mathbf{y}_i - \mathbf{y}\| < \gamma$ where $A\mathbf{y}_i = \mathbf{q}_i$.

I use the basic tools of sensitivity analysis with a slight modification since I am dealing with a non-square matrix A. Since A has full rank, I can identify $|\mathcal{C}| + |\mathcal{S}|$ columns that form a basis for the column space of A. Therefore consider the square matrix consisting of only these $|\mathcal{C}| + |\mathcal{S}|$ spanning columns, and denote this matrix B. Let ρ be a $|\mathcal{C}| + |\mathcal{S}|$ vector describing the column indices of the retained basis columns and assume ρ is increasing.

Suppose that I have identified an $\mathbf{x}$ that solves equation B.1. Now I am faced with a perturbation $\boldsymbol{\varepsilon} \in \mathbb{R}^{|\mathcal{C}|+|\mathcal{S}|}$ of $\mathbf{q}$ and wish to solve

$$A(\mathbf{y} + \boldsymbol{\delta}) = \mathbf{q} + \boldsymbol{\varepsilon}$$

which is equivalent to solving

$$A\boldsymbol{\delta} = \boldsymbol{\varepsilon} \tag{B.3}$$

I solve the equivalent problem

$$B\mathbf{z} = \boldsymbol{\varepsilon}$$

which admits a solution since B has full rank. The vector $\mathbf{z}$ can be converted into a vector $\boldsymbol{\delta}$ that solves equation B.3 by setting $\delta_j = z_{\rho(i)}$ when $\rho(i) = j$ and $\delta_j = 0$ otherwise. From the standard theory of the perturbation of square systems of equations

$$\|\boldsymbol{\delta}\| \leq \left\|B^{-1}\right\| \|\boldsymbol{\varepsilon}\|$$

In the language of lower hemicontinuity, I have shown that I can choose $\mathbf{y}_i$ to satisfy

$$\|\mathbf{y}_i - \mathbf{y}\| \leq \left\|B^{-1}\right\| \|\mathbf{q}_i - \mathbf{q}\|$$

which implies lower hemicontinuity (and hence continuity) of the correspondence of solutions to equation B.1. $\square$

I am still one step removed from making claims regarding sets of stable, feasible matches because $\mathbf{Y}(\mathbf{q}; \underline{x}_C, \underline{x}_S)$ may include solutions that do not satisfy the non-negativity constraints of the match. Let $\boldsymbol{S}(\mathbf{q}; \underline{x}_C, \underline{x}_S)$ denote the set of matches that satisfy the stability constraints defined by $\underline{x}_C$ and $\underline{x}_S$, the capacity constraints defined by $\mathbf{q}$, and the non-negativity constraints. I have to be careful at this juncture, because equation B.1 does not constrain $x(c,s)$ when (c,s) is a blocking pair with respect to $\underline{x}_C$ and $\underline{x}_S$. To this end let $\mathbb{R}_{SC} \subset \mathbb{R}^{(|\mathcal{C}|+1)(|\mathcal{S}|+1)}$ be a product set whose i^{th} dimension is equal to $\mathbb{R}_+$ if the i^{th} column of A denotes a (c,s) that is not a blocking pair given my choice of $\underline{x}_C$ and $\underline{x}_S$. The i^{th} dimension of $\mathbb{R}_{SC}$ is equal to $\{0\}$ otherwise. Given this definition

$$\boldsymbol{S}(\mathbf{q};\underline{x}_C,\underline{x}_S) = \cup_{\mathbf{y}\in\mathbf{Y}(\mathbf{q};\underline{x}_C,\underline{x}_S)}\mathbf{y} \cap \mathbb{R}_{SC} \tag{B.4}$$

which insures that $\boldsymbol{S}(\mathbf{q};\underline{x}_C,\underline{x}_S)$ contains only the solutions of equation B.1 that satisfy the non-negativity constraints.

The following theorem implies that the desired continuity property holds for generic values of **q.**

Theorem 1. *For a topologically generic set of* $\mathbf{q} \gg \mathbf{0}$ *one of the following is true:*

1. $\mathbf{S}(\mathbf{q};\underline{x}_C,\underline{x}_S)$ *is empty.*

2. $\mathbf{S}(\mathbf{q};\underline{x}_C,\underline{x}_S)$ *is non-empty and continuous at* $\mathbf{q}$.

Proof. First I must show that for any $\mathbf{q}$ where one of the two statements is true, that it must remain true for an open neighborhood of $\mathbf{q}$. Second, I must show that $\mathbf{q}$ where the claim holds are dense in that I can find such a $\mathbf{q}'$ arbitrarily close to any $\mathbf{q} \in \mathbb{R}_+^{|\mathcal{C}|+|\mathcal{S}|}$ such that either claim 1 or 2 holds.

As an initial step, note that if A is not of full rank, theorem 3 implies that statement (1) is true for a topologically generic set and I am done. For the remainder, assume that A has full rank. Furthermore, I consider only $\mathbf{q}$ such that $\mathbf{Y}(\mathbf{q};\underline{x}_C,\underline{x}_S)$ is nonempty and continuous since theorem 3 implies these $\mathbf{q}$ are topologically generic.

Note that if claim (2) holds at $\mathbf{q}$ then it must, by definition, hold for all $\mathbf{q}'$ within an open neighborhood of $\mathbf{q}$. Now I show that if claim (1) holds, it must hold in an open neighborhood of $\mathbf{q}$. If $\mathbf{S}(\mathbf{q};\underline{x}_C,\underline{x}_S)$ is empty, then $\mathbf{Y}(\mathbf{q};\underline{x}_C,\underline{x}_S)$ lies in the open set $\mathbb{R} - \mathbb{R}_{SC}$. Since I know that $\mathbf{Y}(\mathbf{q};\underline{x}_C,\underline{x}_S)$ is continuous in $\mathbf{q}$, for any $\mathbf{q}'$ sufficiently close to $\mathbf{q}$ I must have that $\mathbf{Y}(\mathbf{q}';\underline{x}_C,\underline{x}_S) \subset \mathbb{R} - \mathbb{R}_{SC}$. This last fact implies $\mathbf{S}(\mathbf{q}';\underline{x}_C,\underline{x}_S)$ is empty. This concludes the first step of the proof.

For the second step, suppose $\mathbf{S}(\mathbf{q};\underline{x}_C,\underline{x}_S)$ is not empty and not continuous in $\mathbf{q}$. This can only be the case if all $y \in \mathbf{Y}(\mathbf{q};\underline{x}_C,\underline{x}_S)$ are on the "non-trivial" exterior of $\mathbb{R}_{SC}$ in the sense that for each such y there exists $y' \in \mathbb{R}_{SC}$ and a dimension j such that $y_j = 0 < y'_j$. In this case, consider $\boldsymbol{\delta} \in \mathbb{R}_{SC}$ small and strictly positive in all dimensions where that is possible.[24] Let $\boldsymbol{\varepsilon} = -A\boldsymbol{\delta} < \mathbf{0}$, and note that as $\boldsymbol{\delta} \to 0$ I have $\boldsymbol{\varepsilon} \to 0$.

Consider the following system:

$$A\mathbf{z} = \mathbf{q} + \boldsymbol{\varepsilon} \tag{B.5}$$

For $\boldsymbol{\varepsilon}$ sufficiently small, $\mathbf{q} + \boldsymbol{\varepsilon}$ is a valid capacity vector, so the set of solutions to this system that lie in $\mathbb{R}_{SC}$ define $S(\mathbf{q} + \boldsymbol{\varepsilon}; \underline{x}_C, \underline{x}_S)$. But also note that $\mathbf{z}$ solves equation B.5 if and only if there is some $\mathbf{y} = \mathbf{z} + \boldsymbol{\delta} \in \mathbf{Y}(\mathbf{q}; \underline{x}_C, \underline{x}_S)$ such that $\mathbf{y}$ solves equation B.1. Since $\mathbf{y}$ is on the non-trivial exterior of $\mathbb{R}_{SC}$ and I chose $\boldsymbol{\delta} \in \mathbb{R}_{SC}$ to be strictly positive in all dimensions where possible, it must be the case that $\mathbf{z} = \mathbf{y} - \boldsymbol{\delta} \notin \mathbb{R}_{SC}$. Since this holds for all solutions to equation B.5, it must be the case that

$$S(\mathbf{q} + \boldsymbol{\varepsilon}; \underline{x}_C, \underline{x}_S) = \mathbf{Y}(\mathbf{q} + \boldsymbol{\varepsilon}; \underline{x}_C, \underline{x}_S) \cap \mathbb{R}_{SC} = \varnothing$$

In other words, I can choose $\mathbf{q}'$ arbitrarily close to $\mathbf{q}$ such that claim 1 holds. This concludes the proof. □

C For Online Publication - Implementing the Programming Approach

The goal of this section is to describe the computational methods used to solve the programs used in my analyses of welfare, diversity, and neighborhood schools. A recurring theme is a series of techniques that I have used to render the program more tractable. The key difficulty I need to address is the potentially large type space and the correspondingly large number of incentive compatibility constraints.

C.1 Simplifying the Type Space

The first (and most important) simplification I use is the fact that only 10 of the school programs are overdemanded. This means that I can code a student's rank list of school programs as an ordering of the overdemanded school programs and a generic underdemanded school program to which the student will be admitted if he or she

[24] Recall that $\mathbb{R}_{SC}$ has some dimensions without any nonzero elements.

does not receive a seat at an overdemanded school program. Unless otherwise stated, c refers to an overdemanded school program. In addition, I treat the walk-zone slots and the open slots at each school program as different programs.

The second simplification I make is with regard to the priority structure. Recalling that the priority structure is a component of the students' types, this again amounts to a simplification of the type space. BPS defines six levels of priority, but for all practical purposes there are three levels of priority a student can have at a school. The highest priority students are insured a seat, meaning he or she is formally guaranteed a seat, has a sibling enrolled in the school program, or has the NoWalkZoneInGeo priority. Students of intermediate priority live within the walk-zone of the school, and have a higher priority at the walk-zone seats. Third, students may have no priority for any seats at a school program.

C.2 Simplifying the Stability Constraints

Another source of significant simplification is the special structure of the stability constraints. First, all students that are insured a seat must be assigned to a school that is at least as preferred as the insured outcome. The second kind of stability requirement relates to whether walk-zone slots are filled by students with walk-zone priority or not. If the walk-zone slots at a school are filled by students with walk-zone priority, then the constraint is that no student-types without walk-zone priority can be matched with the walk-zone seats. If the walk-zone seats in a school program are not all filled by students with walk-zone priority, then any student-type with walk-zone priority at that school program must be assigned to school programs at least as preferred as the school program to which the student-type has walk-zone priority.

Now I state the stability constraints formally. Denote the set of student-types with insured seats at school program c as $\mathcal{G}(c)$. Let $\mathcal{W}(c)$ refers to student-types that have at least walk-zone priority at school program c. If the walk-zone slots at a school program admit only students with walk-zone priority or higher, the stability constraints are satisfied if and only if

$$\begin{aligned}
\text{For all } s &\in \mathcal{G}(c)\ ,\ \pi^S(s) = \sum_{\{c' \in \mathcal{C} \cup \{\varnothing\} : c' \succeq_s c\}} x(c', s) && \text{(C.1)}\\
\text{For all } s &\notin \mathcal{W}(c)\ , x(c,s) = 0 \text{ where c is a walk-zone school program} && \text{(C.2)}\\
\sum_{s \in \mathcal{S}} x(c,s) &= q_c && \text{(C.3)}
\end{aligned}$$

Equation C.1 requires that all student-types insured a seat at school program c must be matched with a school program at least as preferred as college c. Equation C.2 requires that admitted students must have walk-zone priority to be matched with the walk-zone slots of the school program. Equation C.3 requires that there be no empty seats at an overdemanded school program. If a school program admits a student in the "No Priority" group, then it must be the case that the stability requirements have the form

$$\begin{aligned}
\text{For all } s &\in \mathcal{G}(c)\ ,\ \pi^S(s) = \sum_{\{c' \in \mathcal{C} \cup \{\varnothing\} : c' \succeq_s c\}} x(c', s) \\
\text{For all } s &\in \mathcal{W}(c)\ ,\ \pi^S(s) = \sum_{\{c' \in \mathcal{C} \cup \{\varnothing\} : c' \succeq_s c\}} x(c', s) && \text{(C.4)}\\
\sum_{s \in \mathcal{S}} x(c,s) &= q_c
\end{aligned}$$

where equation C.4 requires that all students with at least walk-zone priority at school program c must be admitted to a school program at least as preferred as c.

Given this structure, each school program has one of the two combinations of stability constraints: either the walk-zone slots are filled with walk-zone students or there are walk-zone slots available for students without walk-zone priority. Given 10 overdemanded school programs, this means that there are 1,024 possible combinations of stability conditions I need to consider. In my computational study I solve each of these 1,024 problems separately. Note that not all of these problems are feasible. One of these problems must be feasible and, in practice, many are. For example, 24 of the 1,024 problems solved for my welfare analysis are feasible. The global optimum for my problem is then the optimum over these 24 subproblems.

C.3 Simplifying the Incentive Constraints

The incentive constraints are the source of most of my computational difficulties. The ODIC standard requires the imposition of $|\mathcal{S}| * (|\mathcal{S}| - 1) * R$ constraints where $|\mathcal{S}|$ is the size of the student-type space and R is the maximum length of the rank-order lists that can be submitted to the mechanism. I can further reduce the number of constraints by noting that for a student of type s, I need only $(|\mathcal{S}| - 1) * R(s)$ constraints where $R(s)$ denotes the number of overdemanded schools in the preference ranking of students with type s. In the Boston context this simplification is particularly useful since many students only rank a few overdemanded school programs.

My final simplification is to relax the incentive constraints, although I argue below that this relaxation is without loss of generality. Let $\mathcal{S}_{\exists}$ denote the types of student that exist in the data, and let $\mathcal{S}_{\nexists} = \mathcal{S} - \mathcal{S}_{\exists}$ denote "hypothetical" types that do not exist in the data. For all types in $s \in \mathcal{S}_{\exists}$, I impose constraints to insure that it is ODIC to not declare another type $s' \in \mathcal{S}_{\exists}$ where s and s' have identical verifiable traits. My constraints do not (formally) require that it be ODIC for s to not declare some type $s' \in \mathcal{S}_{\nexists}$.

Before formally justifying this simplification, let us first consider what would happen if I solved the model with the full $|\mathcal{S}| * (|\mathcal{S}| - 1) * R$ set of incentive constraints. Recall that my objective functions are based only on the assignment of student-types that participate in the market (i.e., the objective function is not affected by the outcome of types that no student declares). Consider a type of student s' that does not exist in the market. The solution to my convex program will assign an outcome to type s' that slackens the incentive constraints for the student-types that do exist in the market as much as possible. One might conjecture that the easiest way to slacken the constraints would be to simply assign hypothetical types s' to an underdemanded school. However, this solution may not be feasible if doing so violates the stability constraints (e.g., student-type s' might be insured admission to some school) or the incentive constraints of a hypothetical type s'.

Now, let us turn to what I am implicitly assuming when I require only that types in $\mathcal{S}_{\exists}$ be disincentivized from mimicking another type in $\mathcal{S}_{\exists}$. In other words, what outcome am I implicitly assigning to hypothetical types like $s' \in \mathcal{S}_{\nexists}$? First, suppose that s' is insured a seat in some overdemanded school, in which case my match implicitly assigns s' to that school. Any student that could declare him or herself to have type s' must also have a (verifiable) insured seat at this school, so he or she can do not better by mimicking type s'. Second, suppose that type s' has walk-zone access at a school c that allows students without walk-zone priority to have access to walk-zone slots. Then, as noted in my discussion of the stability constraints, type s' has in effect been insured a seat at c. In line with my first point, I implicitly assign students of type s' to school c. Again, the only students that could mimic type s' must also have walk-zone access to school c and are assigned to a slot at least as preferred as school c by my relaxed program.

Finally, consider types $s' \in \mathcal{S}_{\nexists}$ that do not satisfy either of these earlier points. My program implicitly assigns type s' his or her most preferred assignment out of all of those given to types $s \in \mathcal{S}_{\exists}$ that have the same priorities as type s'. Written more formally:

$$x(c, s') = \max\{x(c, s) : s \in \mathcal{S}_{\exists} \text{ and } s \text{ has the same priorities as } s'\}$$

where the maximum is taken with respect to the true preferences of type s'. In the event that there are no students $s \in \mathcal{S}_{\exists}$ that have the same priorities as s', I simply assign s' to an underdemanded school. The key insight here is that the range of $x(c, \circ)$ over $\mathcal{S}_{\exists}$ and $\mathcal{S}$ are the same. In other words, any outcome that a type $s \in \mathcal{S}_{\exists}$ can obtain by declaring $s' \in \mathcal{S}_{\nexists}$, again where s and s' have identical priorities, can be obtained by declaring some type $s'' \in \mathcal{S}_{\exists}$.

The astute reader might wonder what happens if $x(c, s')$, $s' \in \mathcal{S}_{\nexists}$, yields the possibility of an assignment to an unacceptable school. In some practical matching contexts, such as the medical resident match, there is a real possibility that a would-be resident would prefer to forgo a residency than be placed in a extremely disliked hospital. In practice no school programs are unacceptable since any student that fails to match with a school is administratively assigned to an under-demanded school program.

D For Online Publication - Constructive Description of the Continuity of the Match

My purpose in this section is two fold. First one might have conjectured that if there is a small perturbation of the population of agents, running the Gale-Shapley algorithm to find new partners for agents affected by the perturbation ought to, in an informal sense, generically result in a nearby stable match. This section describes conditions under which this intuition holds. Second, the result in section 4.3 leveraged the analytic structure of the equations describing the stable set. This section provides arguments in the form of constructive algorithms and provides a link between the combinatoric arguments often used in the matching literature and the analytic techniques used in section 4.3.

Now I provide conditions under which one can adjust the match using an algorithm similar to the Gale-Shapley algorithm to account for perturbations of the distribution of student-types, π^S. In the tradition of the matching literature, the result is a constructive proof. My algorithm is based on the resolving blocking pairs that result from perturbing π^S, and bears some similarity to the decentralized paths to stability described by Roth and Vande Vate [44].

To describe how my algorithm accounts for a small perturbation of the distribution of student-types, I define *displacement paths* and *vacancy paths*. A displacement path is an alternating sequence of students and colleges of the form $(s_1, c_2, s_3, c_4, ...)$. The displacement path describes how the matching adjusts to a small perturbation of π^S that increases the population of a particular student-type. The s_1 at the beginning of the sequence denotes a small measure of students of type s_1 that have been added to the economy. The second element c_2 denotes that the additional students of type s_1 are accepted at c_2, and the third element s_3 indicates that an equal measure of previously admitted students of type s_3 that must be rejected by c_2. The newly rejected students of type s_3 are in turn accepted by college c_4.

Displacement paths need not be finite, but if they are finite they must terminate with either $(.., s_5, c_6)$ or $(..., s_5, \varnothing)$. If the final element denotes a college-type, then a small measure of students of type s_5 can be accepted at college c_6 by filling empty seats at the college. If the final element is $\varnothing$, then the small measure of students of type s_5 are left unmatched at the end of the displacement path.

A vacancy path is a sequence of the form $(c_1, s_2, c_3, s_4, ...)$ that describes how the matching adjusts to a small perturbation of π^S that decreases the population of a particular student-type, which I denote s'. If some students of type s' are unmatched before the perturbation, then the vacancy path $(\varnothing)$ is admissible, which denotes that the decrease in the population of student-type s' only has the effect of reducing the number of unmatched agents of this type. Otherwise, c_1 denotes a college to which students of type s' were admitted that now has empty seats as a result of students of type s' leaving the economy. These seats can be filled by students of type s_2, who leave college c_3 to obtain a newly vacant seat at c_1. The newly empty seats at college c_3 are in turn filled by students of type s_4.

As with displacement paths, vacancy paths need not be finite. Finite vacancy paths terminate with either $(.., c_5, s_6)$ or $(..., c_5, \varnothing)$. The former indicates that the newly vacant seats at c_5 are taken by otherwise unmatched students of type s_6. The second case indicates that the newly vacant seats at c_5 remain empty.

Below I provide an example of a displacement path with two colleges and three student-types.

Example 3. *Consider an economy where $\mathcal{S} = \{s_1, s_2, s_3\}$, $\mathcal{C} = \{c_1, c_2\}, \pi^S(s_i) = \frac{1}{3}$ and $q_c = \frac{1}{2}$. Suppose all students strictly prefer c_1 to c_2 and both colleges have priorities $s_1 \succ_c s_2 \succ_c s_3$. The only stable match is assortative*

$$\begin{aligned} x(c_1, s_1; \pi^S) &= \frac{1}{3} \\ x(c_1, s_2; \pi^S) &= \frac{1}{6} \\ x(c_2, s_2; \pi^S) &= \frac{1}{6} \\ x(c_1, s_3; \pi^S) &= \frac{2}{3} \end{aligned}$$

Suppose an additional mass $\varepsilon > 0$ of students of type s_1 are present. This ε mass of students must be accepted at college c_1, which displaces an ε mass of students s_2 from c_1. This ε mass of students of type s_2 are in turn accepted at college c_2, which results in an ε mass of students of type s_3 being displaced from c_2. Finally, this ε mass of students of type s_3 must remain unmatched. This displacement path is described as $(s_1, c_1, s_2, c_2, s_3, \varnothing)$.

To insure the continuity of feasible, stable matches, it suffices that the match admit a sufficiently rich set of displacement and vacancy paths that (1) do not admit cycles, (2) respect the stability constraints of the match, and (3) are feasible given $x(c, s; \pi^S)$. The acyclicity requirement is violated in example 1, which follows a displacement path of the form $(s_1, c_2, s_2, c_1, s_1, c_2, ...)$. The stability condition requires that each student-college pair in the displacement path respect the stability conditions of the original match, which insures that following the displacement path does not generate blocking pairs. The feasibility condition requires that when I break matches to displace students or fill vacant seats, these broken matches must be described by $x(c, s; \pi^S)$. In example 3, the displacement path $(s_1, c_1, s_3, c_2, ...)$ would not be feasible since it requires students of type s_1 to apply to c_1 and displace students of type s_3 from college c_1. This path is not feasible since $x(c_1, s_3; \pi^S) = 0$, which means that no students of type s_3 were assigned to c_1 and so none can be displaced.

If I wanted to simply insure that the match is continuous with respect to π^S, it would suffice that π^S admit a match with a sufficiently rich set of acyclic, stable, feasible displacement and vacancy paths. However, I would like to insure a stronger notion of continuity, which I denote as *conservative continuity.*

Definition 5. *A match $x(c, s; \pi^S)$ satisfies* ***conservative continuity*** *if*

1. *$x(c, s; \circ)$ is continuous for all c and s*
2. *For all $\widetilde{\pi}^S$ sufficiently close to π^S, I have that $x(c, s; \pi^S) = 0$ implies $x(c, s; \widetilde{\pi}^S) = 0$*

Conservative continuity has the advantage that small changes in π^S can be accommodated through small changes in the measure of existing matches college and student-types. If condition (2) of the definition fails, then a small change in π^S might necessitate that a student-type s suddenly be granted a seat in college c where no other student of type s has been granted a seat. I argue below that conservative continuity can be insured for generic economies, which gives me little reason to settle for weaker properties.

To prove conservative continuity holds, I require that the displacement and vacancy paths *respect existing matches.* A displacement path, which describes an algorithm for accommodating excess students, respects existing matches if the excess students are accommodated by increasing the measure of that student-type at some colleges at which the type is already enrolled. This same logic must apply to all of the students displaced along the path. Symmetrically, vacancy paths respect existing matches when empty seats are filled with student-types that have already been enrolled in the school.

For the remainder of this appendix, I use the terms displacement and vacancy paths to refer to acyclic, feasible paths that respect existing matches. These conditions are formalized as follows:

Definition 6. *An **acyclic, feasible displacement/vacancy path that respects existing matches** satisfies*

1. *(Acyclicity) Each college and student-type appears at most once in the sequence.*
2. *(Feasibility) For each pair (c_i, s_{i+1}) in the path I have $x(c_i, s_{i+1}; \pi^S) > 0$.*
3. *(Respects Existing Matches) For each pair (s_i, c_{i+1}) in the path I have $x(c_{i+1}, s_i; \pi^S) > 0$.*

If $x(\varnothing, s) > 0$, then the displacement path $(s, \varnothing)$ is an acyclic, feasible displacement path that respects existing matches by convention. If $x(c, \varnothing) > 0$, then the vacancy path $(c, \varnothing)$ is an acyclic, feasible vacancy path that respects existing matches by convention.

Note that condition (3) implies that following the path will not generate new blocking pairs since the path merely increases or decreases the measure of existing student-college matches. Theorem 2 proves that if a match admits a sufficiently rich set of acyclic, feasible displacement path that respects existing matches, then the designer can implement a match that satisfies conservative continuity. The richness requirement requires that there exists a path that allows us to adjust the match for any perturbation of π^S.

Definition 7. *A match $x(c, s)$ admits a **rich set of paths** if for each $s \in \mathcal{S}$ and $c \in \mathcal{C}$ there exists an acyclic, feasible displacement or vacancy path that respects existing matches and commences with (respectively) s or c.*

First I prove that the richness criterion is sufficient for continuity. My proof is constructive and involves arguing that my algorithm can follow select paths to make small changes to $x(c, s; \pi^S)$ to account for small changes in π^S without violating the feasibility or stability conditions.

Theorem 2. *Suppose that $x(c, s; \pi^S)$ admits a rich set of paths. Then for some open neighborhood U of π^S one can choose the match so that*

1. *$x(c, s; \pi^S)$ is locally Lipschitz continuous in $\pi^S \in U$*
2. *$x(c, s; \pi^S)$ admits a rich sets of paths for all $\pi^S \in U$*

Proof. Consider any $\widetilde{\pi}^S, \pi^S \in U$. Since the set of student-types is finite, there is a finite sequence $\left(\pi_1^S, \pi_2^S, .., \pi_N^S\right), N \leq |\mathcal{S}| + 1$, such that $\pi_1^S = \pi^S$, $\pi_N^S = \widetilde{\pi}^S$, and for each $i \in \{1, .., N\}$ there exists s_i such that $\pi_i^S(s_i) \neq \pi_{i+1}^S(s_i)$ and for all $s' \neq s_i$ I have $\pi_i^S(s') = \pi_{i+1}^S(s')$. In other words, moving from π_i^S to π_{i+1}^S involves changing the population of one student-type. Since I am making a continuity argument, U can be thought of as a small open set, so the sequence can be chosen to involve only small differences between each element. I assume throughout that $\left(\pi_1^S, \pi_2^S, .., \pi_N^S\right)$ is of the shortest length possible, so that each s_i is used at most once. I prove at the close of this proof that $(\widetilde{\pi}^S, \widetilde{x}) \in U$ admit rich sets of paths,[25] so I assume that the measures π_i^S and the associated matches I construct admit rich sets of paths.

I prove sufficiency by providing an argument for continuity across pairs of any such sequence, which (since the sequences are finite) implies continuity of the match over U. The proof is constructive, so I build a sequence of feasible and stable matches $(x_1, .., x_N)$ where x_i is feasible and stable under π_i^S, and successive matches in the sequence are close to each other. Note that $x_1 = x(c, s; \pi^S)$ and $x_N = x(c, s; \widetilde{\pi}^S)$. Unless otherwise noted, for the duration of this proof I use the term *match* to refer to feasible, stable matches and the term *displacement (vacancy) path* to refer to acyclic, feasible displacement (vacancy) paths that respect existing matches.

Consider x_i and the pair of distributions π_i^S, π_{i+1}^S. Suppose that $\pi_i^S(s_i) < \pi_{i+1}^S(s_i)$, so that I am adding an additional $\varepsilon = \pi_{i+1}^S(s_i) - \pi_i^S(s_i)$ mass of students of type s_i to the economy. I have assumed the existence of a displacement path for s_i given the match x_i, and I denote the path $(s_1, c_2, s_3, ...)$ where the path has K elements and $s_1 = s_i$.

I now explicitly provide an algorithm to construct x_{i+1}. As a first step, initialize $x_{i+1} = x_i$. For $j \in \{2, 4, .., K-2\}$ set

$$\begin{aligned} x_{i+1}(c_j, s_{j-1}) &= x_i(c_j, s_{j-1}) + \varepsilon \\ x_{i+1}(c_j, s_{j+1}) &= x_i(c_j, s_{j+1}) - \varepsilon \end{aligned}$$

The displacement path ends with either $(.., s_{K-1}, c_K)$ or $(..., s_{K-1}, \varnothing)$. In the former

[25] The argument for the richness is discussed at the close of the proof since the details of our construction need to be described first.

case, let

$$
\begin{aligned}
x_{i+1}(c_K, s_{K-1}) &= x_i(c_K, s_{K-1}) + \varepsilon \\
x_{i+1}(c_K, \varnothing) &= x_i(c_j, \varnothing) - \varepsilon
\end{aligned}
$$

In the latter case, let $x_{i+1}(\varnothing, s_{K-1}) = x_i(\varnothing, s_{K-1}) + \varepsilon$. Since the displacement path is feasible, x_{i+1} is feasible given π^S_{i+1} for ε sufficiently small. The displacement respects existing matches, so $\underline{x}_C$ and $\underline{x}_S$ are identical for x_i and x_{i+1}. This implies that if x_i is stable, then x_{i+1} is also stable. Since the displacement path is acyclic, it involves at most $|\mathcal{S}| + |\mathcal{C}|$ displacements of size at most ε, so I know that $\|x_{i+1} - x_i\| \leq (|\mathcal{S}| + |\mathcal{C}|)\,\varepsilon$.

Now consider the case where $\pi^S_i(s_i) > \pi^S_{i+1}(s_i)$, so that I have removed an $\varepsilon = \pi^S_i(s_i) - \pi^S_{i+1}(s_i)$ mass of students of type s_i from the economy. Suppose $x_i(s_i, \varnothing) > 0$, meaning some students of type s_i are unmatched under x_i. Then for ε I can let $x_{i+1}(s_i, \varnothing) = x_i(s_i, \varnothing) - \varepsilon$, $x_{i+1}(s_i, c) = x_i(s_i, c)$ for all c, and $x_{i+1}(s, \circ) = x_i(s, \circ)$ for all $s \neq s_i$. In this case $\|x_{i+1} - x_i\| = \varepsilon$. If $x_i(s_i, \varnothing) = 0$, then there exists some c such that $x_i(s_i, c) > 0$ as well as an acyclic, stable, feasible vacancy chain starting at c. Let this vacancy chain be denoted $(c_1, s_2, c_3, ...)$ where $c_1 = c$. Assume the vacancy chain has K elements. I begin to build x_{i+1} by first initializing $x_{i+1} = x_i$. I then let $x_{i+1}(c_1, s_i) = x_i(c_1, s_i) - \varepsilon$ and $x_{i+1}(c_1, s_2) = x_i(c_1, s_2) + \varepsilon$. For each $j \in \{3, 5, .., K-2\}$ set

$$
\begin{aligned}
x_{i+1}(c_j, s_{j-1}) &= x_i(c_j, s_{j-1}) - \varepsilon \\
x_{i+1}(c_j, s_{j+1}) &= x_i(c_j, s_{j+1}) + \varepsilon
\end{aligned}
$$

If the vacancy chain ends with $(, ..., c_{K-1}, s_K)$ let

$$
\begin{aligned}
x_{i+1}(c_{K-1}, s_{K-2}) &= x_i(c_{K-1}, s_{K-2}) - \varepsilon \\
x_{i+1}(c_{K-1}, s_K) &= x_i(c_{K-1}, s_K) + \varepsilon
\end{aligned}
$$

If the vacancy chain ends with $(, ..., c_{K-1}, \varnothing)$ let

$$
\begin{aligned}
x_{i+1}(c_{K-1}, s_{K-2}) &= x_i(c_{K-1}, s_{K-2}) - \varepsilon \\
x_{i+1}(c_{K-1}, \varnothing) &= x_i(c_{K-1}, \varnothing) + \varepsilon
\end{aligned}
$$

As with the adjustment via a displacement path, since the vacancy path is feasible and respects existing matches, x_{i+1} is a feasible and stable match. Since the vacancy path is acyclic, I again have $\|x_{i+1} - x_i\| \leq (|\mathcal{S}| + |\mathcal{C}|)\,\varepsilon$.

Since the construction holds for each (π_i^S, π_{i+1}^S) pair, by applying the techniques to each successive (π_i^S, π_{i+1}^S) pair I construct a feasible stable match given $\widetilde{\pi}^S$, which I denote $x(c, s; \widetilde{\pi}^S)$, such that

$$\left\|x(c, s; \widetilde{\pi}^S) - x(c, s; \pi^S)\right\| \leq |\mathcal{S}|\,(|\mathcal{S}| + |\mathcal{C}|)\,\left\|\pi_i^S - \pi_{i+1}^S\right\|$$

Therefore $x(c, s; \circ)$ can be chosen to be locally Lipschitz continuous.

Note that for $\varepsilon = \left\|\widetilde{\pi}^S - \pi^S\right\|$ sufficiently small $x(c, s; \pi^S) > 0$ implies $x(c, s; \widetilde{\pi}^S) > 0$. This implies that if the paths are feasible and respect existing matches under $x(c, s; \pi^S)$, then the same will be true under $x(c, s; \widetilde{\pi}^S)$. Obviously, the paths remain acyclic under both $x(c, s; \pi^S)$ and $x(c, s; \widetilde{\pi}^S)$ since this condition has nothing to do with the match, per se. Therefore, the paths that are valid under $x(c, s; \pi^S)$ remain valid under $x(c, s; \widetilde{\pi}^S)$. Since x admits a rich set of paths, the same is true of $x(c, s; \widetilde{\pi}^S)$. $\square$

The existence of a rich set of paths can be difficult to verify in practice. However, I can insure the existence of a rich set of paths if π^S places support over a subset of the student-types that I call *grounders*. In analogy to electrical circuits, continuity fails in example 1 because arbitrarily small fluctuations in the measure of the student-types (a small power surge in my analogy) must generate long cycles of rematching (an overloaded circuit). The role of grounders in the match is to provide an outlet for small fluctuations in the measure of students that insures that large cycles can be avoided.

Definition 8. *Student-type s is a* ***grounder*** *if only a single overdemanded college is an acceptable match.*

Grounders serve two roles. First, a grounder that is matched with his favorite (and only acceptable) college can be rejected by the college without fear that the student will affect the match by applying to another college. Therefore, if a student-type is in surplus and needs to get matched to college c, one can accommodate the surplus students by rejecting admitted grounders without worrying about secondary effects on other colleges. Second, grounders that have not been admitted to their favorite college can be used to fill any deficits if a student-type is in short supply. In the vernacular I use, the first role allows for the construction of short displacement paths, while the second role insures the existence of short vacancy paths.

Let the set of all grounders be $G = \{(v, u, \succeq_s) : c \succeq_s \varnothing$ for a single overdemanded $c \in \mathcal{C}\}$. I require two conditions to establish the existence of a rich set of paths. First, I require that the grounders be sufficiently richly distributed throughout the colleges' priorities.

Definition 9. *The grounders are* ***richly distributed*** *if for each $s \in \mathcal{S}$ and $c \in \mathcal{C}$ there exists $g \in G$ such that $g \sim_c s$.*

If π^S has full support over G and the grounders are richly distributed in the priority structure, then I can construct a match with a rich set of paths by insuring that some, but not all, of each type of grounder are admitted.

Theorem 3. *For a generic π^S with full support in an economy where grounders are richly distributed, there exists a match $x(c, s; \pi^S)$ that admits a rich set of paths.*

Proof. The proof proceeds in two parts. Part one: I describe how to use Grounders to construct a rich set of paths, yielding a match x. The key to my construction is to insure that the grounders fulfill the dual roles described above. In part two I argue that the construction holds for generic π^S.

Part One: The Construction

The key to insuring that the grounders can fulfill both roles is to assign some, but not all, of the grounders of each type to their respective college. Let $\epsilon \in (0, \frac{1}{2}\min\{\pi^S(g) : g \in G\})$, which exists since $\pi^S(g) > 0$ for all $g \in G$. For each grounder $g \in G$ I generate a pair of synthetic types g^H, g^L, which denote high and low priority "versions" of that grounder. For each such pair, I extend the college preferences over the set of student-types as follows to insure the desired match is chosen

$$g \succ_c s \Rightarrow g^H \succ_c g \succ_c g^L \succ_c s$$
$$s \succ_c g \Rightarrow s \succ_c g^H \succ_c g \succ_c g^L$$
$$g \sim_c s \Rightarrow g^H \succ_c g \sim s \succ_c g^L$$

In other words, the high and low priority synthetic types are only distinguished within their priority class and do not affect the ranking between other types in different priority classes. Denote the new set of student-types as $\widetilde{\mathcal{S}}$.

Now that I have defined the synthetic types and extended the colleges' preference over these traits, I endow an ϵ measure of each grounder g with the synthetic type g^H and an ϵ measure with trait g^L. Grounders with trait g^H (g^L) will have higher (lower) priority at each college than students that share the priority class of g. Denote the distribution of student-types over $\widetilde{\mathcal{S}}$ as $\widetilde{\pi}^S$

Now that I have assigned the synthetic types, consider a match $x(c, s)$ that is feasible and stable with respect to $\widetilde{\pi}^S$. My goal is now to show that this match admits a rich set of path by constructing a match that is stable with respect to the modified college preferences (i.e., including the high and low priority versions of each g). When I consider the resulting match under the original college preferences, the match remains stable since I am only eliminating some of the stability restrictions imposed by the extended college preferences.

Let $\underline{x}_C, \underline{x}_S$ be defined with respect to the match under the extended preferences. I now explicitly construct the displacement paths. Consider an arbitrary student-type s. There are three cases I consider:

1. $x(\varnothing, s) > 0$, in which case $(s, \varnothing)$ is a valid displacement path.

2. $x(\varnothing, s) = 0$ and there exists c such that $x(c, s) > 0$ and $x(c, \varnothing) > 0$. In this case $(s, c, \varnothing)$ is a valid displacement path.

3. $x(\varnothing, s) = 0$ and there exists c such that $x(c, s) > 0$ and $x(c, \varnothing) = 0$. In this case, there is a grounder $g \in \underline{x}_C(c)$ such that $x(c, g) > 0$. If $x(\varnothing, g) > 0$, then $(s, c, g, \varnothing)$ is a valid displacement path.

The construction of the vacancy paths parallels the displacement paths closely. There are two cases. Case 1a mirrors case 1 for displacement paths, while case 2a mirrors case 3 above.

1a. $x(c, \varnothing) > 0$, which case $(c, \varnothing)$ is a valid vacancy path.

2a. $x(c, \varnothing) = 0$ and there there is a grounder $g \in \underline{x}_C(c)$ such that $x(c, g) > 0$. If $x(\varnothing, g) > 0$, then $(c, g, \varnothing)$ is a valid vacancy path.

Part Two: Proving the construction works generically

I now show that for a generic set of measures π^S, the construction yields cases 1 to 3 for all student-types. Since the construction of the vacancy paths is so similar to the construction of the displacement paths, I provide arguments for the displacement paths and note that essentially identical arguments apply for the genericity of the π^S such that all college-types fall into either case 1a or 2a.

First note that theorem 2 implies that the set of π^S for which I can construct a match that yields a rich set of paths is open. To show genericity, I need to prove that the set of π^S that admit a match with a rich set of paths is and open and dense set in $\Delta(S)$. Although I proceed by considering a particular s the intersection of the (open and dense) set for each s is also open and dense since $\mathcal{S}$ is finite.

Since my construction uses inequalities, the set of π^S for which the construction works is open. To show that the set of π^S that admit a match with a rich set of paths for a given s is dense, I argue that if the algorithm from part 1 does not yield a match with a rich set of paths , then the algorithm can yield such a match for a measure arbitrarily close to π^S. In other words, if a student-type s does not fall into categories 1 to 3 for the measure, then for a measure $\widetilde{\pi}^S$ arbitrarily close to π^S I can use the constructed match, $x(c,s)$, to define a new match that is very close to $x(c,s)$ and for which all of the student-types fall into categories 1 to 3.

Consider the residual case for an arbitrary student-type s:

4. $x(\varnothing, s) = 0$ and for all c such that $x(c,s) > 0$ I have: 1) $x(c, \varnothing) = 0$ and (2) for all grounders $g \in \underline{x}_C(c)$ I have $x(\varnothing, g) = 0$.

To show that the set of π^S for which cases 1 to 3 apply for student-type s is dense in $\Delta(S)$, simply note that to fall into case 4 requires that all grounders $g \in \underline{x}_C(c)$ are matched for all colleges to which student-type s is matched. If I consider $\widetilde{\pi}^S$ such that $\widetilde{\pi}^S(g) > \pi^S(g)$ and $\widetilde{\pi}^S(s) = \pi^S(s)$ for all $s \neq g$, then my construction yields a match that falls into case 3. Therefore, the π^S for which the construction yields cases 1 to 3 for student-type s are dense. Although I have proved density for a single student-type s, it is obvious that if I choose $\widetilde{\pi}^S(g) > \pi^S(g)$ for all g such that $g \in \underline{x}_C(c)$ for some c, then all of the student-types must fall into categories 1 through 3. $\square$

The following corollary follows immediately from the combination of theorems 2 and 3.

Corollary 2. *For generic π^S, there exists a match $x(c,s;\pi^S)$ that is locally Lipschitz continuous in π^S.*