

Optimizing for Distributional Goals in School Choice Problems

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

Abstract

Using data from Boston Public Schools, I investigate the tension between three goals of school choice: student welfare, encouraging neighborhood schools, and diversity within schools. I develop a novel framework for computing the optimal match for any convex combination of these goals. I find that one can make significant gains across all three metrics relative to the status quo outcome generated by the Gale-Shapley algorithm. I also find there are modest trade-offs between the goals, so one can simultaneously increase welfare and school diversity while encouraging neighborhood schools. I close by discussing new menus-and-quotas mechanisms that implement the improved matches.

1 Introduction

School choice mechanisms have been the subject of intensive mechanism design efforts over the past 15 years, which has led Boston and other large cities to employ centralized matching procedures to allocate students to public schools (Abdulkadiroğlu and Sönmez [5]). Market designers often recommend that school assignment mechanisms be built in two steps. First, given the school system’s public policy goals, the designer chooses a priority structure that gives particular students higher priorities at certain schools than other students. For example, to encourage a student to attend her neighborhood school, Boston Public Schools (BPS) gave higher priority for seats at a school to students that live close to that school. Second, given the priority structure chosen, solicit student preferences over the schools and run the student-proposing Gale-Shapley algorithm to generate the school assignment.

There is much to recommend this advice. First, the Gale-Shapley algorithm is incentive compatible, so the students have no motive to try and manipulate the assignment system. Second, the outcome is *stable*, which reflects a notion of fairness that requires that an agent with a high priority for a good be allocated that good before a lower priority agent is provided the good. However, it is far from obvious that encoding policy goals into the priority structure and then running the Gale-Shapley algorithm is the most effective method of implementing goals such as encouraging

*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.

neighborhood schools. In addition, it is not clear what the trade-offs are between goals such as student welfare and encouraging neighborhood schools.

I argue in this paper that matching problems in general, and school choice problems in particular, can be posed as computationally tractable constrained optimization problems. The constrained optimization approach has the advantage of allowing the market designer to explicitly describe the policy goals in the objective function and encode required properties such as stability and incentive compatibility in the constraints. Using my approach I can compute the global optimum for various desiderata (e.g., the student welfare maximizing school assignment) or assess trade-offs between the goals (e.g., encouraging neighborhood schools and school diversity). The optimization approach also allows me to use the shadow prices on the constraints to assess the benefits of, for example, adding capacity to a popular school.

Treating a school choice problem as a constrained optimization problem is not a novel theoretical idea. However, I combine this idea with the more recent approach of modeling large school choice problems as a match between a continuum of seats at a finite set of schools and a continuum of each of a finite set of student types.¹ Since each student has a negligible effect on the aggregate outcome, I can write the incentive compatibility conditions in a computationally tractable form. Although the resulting optimization problem has a large number of constraints, modern techniques for solving large linear and quadratic programs allow me to solve these problems in under a minute even with modest computational resources.

I use my framework to solve for the optimal school assignment using data from the Boston Public Schools (BPS) high school match for the 2011-2012 school year. During this period, BPS used the student-proposing Gale-Shapley algorithm to assign students to schools (see Roth and Sotomayor [58]). Students report a rank-ordered list of up to 10 school programs to BPS, and students are assigned a priority at each school based on whether or not they live within a school's walk-zone.² My data includes the preference lists the students submitted to the mechanism, the priority of each student for assignment to each school, and demographic information such as the student's ethnicity and whether the student participated in a free school lunch program.

I consider three potential objectives: student welfare, student body diversity, and encouraging neighborhood schools. My optimization approach allows me to compare the status quo outcome generated by the Gale-Shapley algorithm with the global optimum found for the relevant optimization problem. In addition, I study the tensions between these goals by, for example, determining to what extent fostering diverse schools impairs the achievement of high student welfare.

Although my analysis focuses on a particular year of the BPS school assignment problem, my goal is to make three broader contributions to the market design and economics of education literature. First, while welfare objectives are clearly important in any school choice setting, there are distributional goals that warrant consideration, such as the encouragement of neighborhood

¹A recent paper by Liu and Pycia [48] proves that all efficient, symmetric, and strategy proof mechanisms lead to the same allocations. We avoid this by considering mechanisms that are not "symmetric" in that students with the same preferences from different demographic groups are treated differently.

²The priority structure is actually more complex. I discuss this issue in more depth in Section 2.

schools or building ethnically and socioeconomically diverse student bodies. My optimization framework provides a method for assessing existing school assignments and designing new school assignments that are optimal with respect to these distributional goals. Second, the analysis of the BPS data shows that there is room to improve on the Gale-Shapley algorithm especially when it comes to distributional goals such as student diversity and encouraging neighborhood schools. Third, my analysis of the BPS school assignment problem serves as a proof of concept for the constrained optimization approach, and the menus-and-quotas mechanisms that I introduce provide a practical way of converting solutions to the constrained optimization problems into real-world mechanisms.

First I consider the objective of maximizing student welfare subject to the capacity, stability, and incentive compatibility constraints.³ Although the Gale-Shapley algorithm is known to satisfy these constraints, the random tie-breakers used in Boston’s Gale-Shapley algorithm can result in a welfare loss (Erdil and Ergin [24], Kesten [41]). I measure a student’s welfare as her assigned school’s rank in the student’s preference list, and the welfare generated by a school assignment is defined as the average welfare across the population of students.⁴ The Gale-Shapley algorithm generates 78% of the welfare gain from moving from the worst to the best stable and incentive compatible match. The 78% welfare gain is equivalent to moving 23.7 students to a school one rank higher in their preference list. I show that the incentive compatibility constraints sharply limit the mechanism’s ability to improve student welfare, while the stability conditions have very little impact on the mechanism’s efficacy.

I demonstrate that if one allows the mechanism to condition on the demographics of the students, then more significant welfare gains are possible. The welfare gained by moving from the Gale-Shapley outcome to the welfare-optimal, stable, and ODIC match found by a mechanism that conditions on student demographics is equivalent to moving 291 students to a school one rank higher in their preference list. The reason the mechanism’s performance improves is that allowing the mechanism to condition on the demographics of the students weakens the incentive compatibility constraints. If one considers it unfair that a mechanism treats students differently based on their demographic traits, then one can view the improved efficiency of a mechanism that does so as the *welfare cost of fairness*. In addition to computing the optimal match, I use the shadow prices of the optimization problem to identify how to reallocate resources across schools to increase welfare. For example, the shadow prices on the capacity constraints reveal that the welfare gain of an extra seat at Snowden International School is 85% higher than at any other overdemanded school.

My second goal is to assess the potential tension between student welfare and encouraging ethnically and socioeconomically diverse student bodies.⁵ I measure the diversity of a school by how closely the composition of a school’s student body matches the demographics of the entire

³In this initial analysis, I do not consider mechanisms that condition on the students’ demographics (e.g., ethnic background).

⁴Although one could interpret my metric as an assumption about the utility functions of the students, I prefer to think of it as a refinement of Pareto optimality.

⁵One might attempt to encode student diversity directly into the students’ utility functions, but I strive to make as few assumptions about the students’ utility functions as possible.

BPS student population, and the average diversity of a match is the average diversity across all of the schools. In other words, a school is diverse if the demographics of its student body closely resembles the demographics of all the students in the BPS system, and a school assignment is diverse if the schools are diverse on average.

Boston has a complicated history of school desegregation beginning with the 1974 “Garrity Decision.” In the case, *Tallulah Morgan et al. v. James Hennigan et al.*, Federal Judge Arthur Garrity determined that the Boston School Committee had intentionally segregated the schools by race. The ruling mandated that a school busing program be used to desegregate the schools, and violent protests against the busing program broke out on several occasions. The federal mandate was lifted in 1987, but BPS continued to set aside seats at the prestigious examination schools for minority students. This race-based admissions system was ruled unconstitutional in 1998 by the U.S. Court of Appeals for the First Circuit. In 1999 BPS began using the Boston mechanism (Abdulkadiroğlu and Sönmez [5]), which does not take ethnic or socioeconomic background into account when assigning students to schools. There is concern that the reforms to the BPS school assignment mechanism in the 1990s has led to a re-segregated school system.⁶

This history may suggest to the reader that there is a significant tension between student welfare and diversity in school assignment. Fortunately, my analysis shows that this intuition is incorrect. There are stable and incentive compatible school assignments that simultaneously perform better on both student welfare and diversity than the outcome generated by the Gale-Shapley algorithm. Moreover, one can increase the diversity of the average school significantly without reducing student welfare at all. My analysis structure also allows me to trace-out a possibility frontier that describes all of the diversity and welfare metric values that can be realized by a stable, incentive compatible match. The possibility frontier shows that there are only weak trade-offs between diversity and student welfare for the vast majority of the possibility frontier. In addition, the Gale-Shapley outcome is well within the frontier, showing that large improvements can be made along both dimensions relative to the status-quo.

My third analysis studies the cost of encouraging neighborhood schools in terms of both student welfare and school diversity. In a 2012 speech, Boston Mayor Thomas Menino articulated the following externality based logic for encouraging neighborhood schools:⁷

“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.”

⁶Barbara Fields, a former equity officer for BPS, has said, “Some of us fear we’re going to return to a very segregated school system.” From: <https://www.bostonglobe.com/opinion/2015/07/01/bringing-parents-strength-all-boston-public-schools/Zc8KBcQ718sSwCnmT19qbI/story.html>

⁷Quote taken from Dur et al. [19].

A school assignment encourages neighborhood schools if a high fraction of each school's student body is drawn from the school's walk-zone. However, the demographics of Boston differ widely from zip code to zip code. While some zip codes do not exhibit significant segregation, there are many zip codes that have either a very low (e.g., 1.3% in the 02126 zip code) or a very high (85.2% in the 01906 zip code) fraction of African-Americans among the school-age residents. One might assume that an assignment that encourages neighborhood schools would require that the schools reflect the segregated zip codes around each school, which would reduce the diversity of the schools.

I study the tension between encouraging neighborhood schools and other desiderata by maximizing a linear combination of the average percentage of students drawn from each school's walk-zone and my metric for the other goal.⁸ I show that one can create classes with 35.8% of their students from the school's walk-zone relative to the status-quo baseline of 22.1%, and this can be accomplished without decreasing average student welfare or making the schools less diverse. I also trace out possibility frontiers for each pair of desiderata (welfare, diversity, encouraging neighborhood schools). As in the case with diversity and student welfare, the frontiers show weak tensions between my objectives and that the Gale-Shapley outcome is well inside the frontiers.

I close my empirical analysis by solving a problem that maximizes a linear combination of the metrics for student welfare, ethnic diversity, and encouraging neighborhood schools. The solution results in schools wherein 32% of the students are drawn from the school's walk-zone, the student bodies' ethnic make-up much more closely resembles that of BPS as a whole, and student welfare is increased by the equivalent of moving almost 162 students to a school one rank higher in their preference ordering. In summary, since there are only weak tensions between my goals, it is possible to find a school assignment that simultaneously achieves significant improvements in the level of diversity, student welfare, and school cohesion relative to the status-quo assignment, all while satisfying stability and incentive compatibility constraints.

As I mentioned above, the use of a continuum model is absolutely crucial to the tractability of my constrained optimization approach. The negligible aggregate 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 in a computationally tractable form. Before accepting the results of my analyses using the continuum model, it is reasonable to ask what connection, if any, my analyses have to the real-world in which a finite set of students is assigned to a finite set of school seats. I provide two answers.

The first answer is a theoretical link between the game-theoretic properties of the model with a continuum agents and an analogous game played by a finite set of students. If one considers each of the optimization problems I solve during my analysis as a mechanism that accepts the students' ordinal preference rankings as inputs, then I show that an equilibrium of the game with a continuum of agents is an approximate equilibrium of the finite game. The core of my argument is a novel result on the continuity of the stable set, which is of independent interest.

The second link is a practical one. I provide a system for converting a solution of one of my

⁸I continue to impose the walk-zone priority that BPS used in the 2011-2012 school year.

constrained optimization problems into a *menus-and-quotas* mechanism that is incentive compatible and implements a stable match close to the solution of the limit model. Since a static version of the menus-and-quotas mechanism may leave a large number of seats open at over-demanded schools, I also provide a modification of the mechanism that dynamically adjusts the menu to limit the number of open seats. The effectiveness of these mechanisms are assessed through simulation.

I provide a summary of the data provided by BPS as well as the status quo outcome of the match in Section 2. In Section 3 I introduce the constrained optimization problem I solve in my analysis. Section 4 analyzes student welfare and the extent to which conditioning on student demographics can improve outcomes. Section 5 analyzes the possibility of optimizing for one of my distributional goals in conjunction with student welfare. Section 6 discusses how the results of the limit model pertain to the finite market that is executed in reality. In addition, Section 6 introduces the menus-and-quotas mechanisms and assesses their efficacy through simulations on bootstrap samples of my data. Placing this paper within the extant literature is greatly simplified if the discussion is placed after both the framework and empirical results are introduced. With that in mind, I conduct a literature review in Section 7, and I conclude in Section 8. All of the proofs are found in the Appendix.

2 Data

This study is based on administrative data that covers 3,479 rising 9th grade students in the Boston Public Schools system during the 2011 - 2012 school year. The data includes administrative data on student demographics as well as all of the inputs to the school assignment mechanism used by BPS at the time to assign the students to the high school programs.

BPS used the student-proposing Gale-Shapley algorithm, which I refer to as the Gale-Shapley mechanism, to assign students to school programs for the 2011 - 2012 school year. I provide a brief description of the Gale-Shapley mechanism after I describe the priority structure used by BPS. The Gale-Shapley mechanism required that each student submit to BPS a rank-order list of up to 10 school programs.⁹ Given the set of student preference profiles, the Gale-Shapley mechanism assigned the students to schools to form a stable match. Importantly, the Gale-Shapley mechanism gave the students no incentive to declare their preferences nontruthfully.¹⁰

Since the Gale-Shapley mechanism is truthful, I take the preferences declared by the students at face value. However, there is field evidence that students may not submit their preferences to a mechanism truthfully (e.g., Rees-Jones [53], Hassidim et al [33]). Some authors have suggested

⁹It is possible for a student to submit a preference ranking consisting entirely of overdemanded schools, which might result in the student not receiving a school assignment from the Gale-Shapley mechanism. In practice, if a student fails to be matched through the Gale-Shapley algorithm (e.g., the student only lists overdemanded school programs and has a low priority), then the student is assigned administratively outside of the system. Conversations with BPS indicate that the administrative assignment is usually to an underdemanded school program near the student's home.

¹⁰As it turns out, there were exactly 10 overdemanded high school programs in the subset of the data I consider, which means it is reasonable to ignore the possibility that students needed to strategically truncate their preferences to rank the overdemanded school programs.

Walk-Zone Seats	Open Seats
Guaranteed-Walk	
Sibling-Walk	Guaranteed
Walk	Sibling-Walk, Sibling,
Sibling	NoWalkZoneInGeo
NoWalkZoneInGeo	Walk, No Priority
No Priority	

Table 1: BPS Priority Structure

that when nontruthful declarations are payoff irrelevant (e.g., failing to rank a school to which you know you will not be admitted), then assuming preference declarations are truthful can result in incorrect inferences (e.g., Artemov et al. [7], Fack et al. [26]). In the status-quo Gale-Shapley mechanism BPS used, any student could potentially be assigned to any school she ranks if her tie-breaker is sufficiently high. Since her random tie-breaker is not known to the student at the time her preferences are declared, there are no payoff-irrelevant mistakes in the status-quo mechanism. Therefore I feel confident in treating the preference declarations of the students as truthful. This issue is discussed in more depth in Section 7.

As mentioned, the Gale-Shapley mechanism results in a stable match. Stability reflects a notion of fairness that requires that all students with a high priority for a seat in a given school program be allocated that seat before a lower priority student is assigned to it. BPS chose the priority structure to reflect a variety of goals: encouraging neighborhood schools, facilitating the enrollment of siblings at the same school, and allowing students to continue at their current school. To achieve these ends, BPS first divides the seats in each school program into two equally sized groups: a set of *walk-zone seats* and a set of *open seats*. The school assignment mechanism treats each pool of seats as distinct school programs with differing priority structures. Each student fits into one of either three (open seat) or six (walk-zone seat) priority classes as described by Table 1, where priority classes are ranked from highest to lowest priority. For example, a student with “Walk” priority has a higher priority for a walk-zone seat than a student with “Sibling” priority.

A student with *guaranteed* or *guaranteed-walk* priority at a school program is insured a place in that program if the student ranks it.¹¹ *Sibling* and *sibling-walk* apply if the student has a sibling already enrolled in that program. *Walk* applies if the student is in the school’s walk zone. *NoWalkZoneInGeo* is the priority of students that are not in any school’s walk-zone. Since many students share the same priority (e.g., many students are in each school’s walk zone), random

¹¹Students are typically guaranteed a spot within a school if they are currently enrolled in a middle school affiliated with the high school.

Walk-Zone Seats	Open Seats
Insured-Walk	Insured-Open
Walk	Walk, No Priority
No Priority	

Table 2: Effective Priority Structure

tie-breaking was used by BPS to assign a unique priority rank to each student at each school program. 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 a relatively high priority at these schools) and be outside the walk-zone of others (and have a relatively low priority).

Students that have a priority at a program of guaranteed, sibling, or NoWalkZoneInGeo are effectively insured admission to an open seat in the program if they rank it. Similarly, students with guaranteed-walk or sibling-walk are effectively insured admission to a walk-zone seat in the program if they rank it. This insight implies that the effective priority structure can be described by Table 2. The new priorities *insured-walk* and *insured-open* mean that the student is automatically assigned to the respective walk or open seat program if she ranks the school program and is not matched to a more preferred school.

For the majority of this paper I treat the BPS priority structure as a constraint in my optimization problems. To the extent that parents view a high priority to attend a high quality school as an inviolable property right, it may be hard to abandon these priorities when changing the school choice mechanism. Although imposing these constraints impedes my ability to find optimal assignments, I include them as a practical constraint.

The Gale-Shapley algorithm has the following timing:

1. Each student submits a rank ordered list of schools to the mechanism and begins the algorithm without a school assignment.
2. All students without a school assignment apply to their most favored school that has not yet rejected their application.
3. Each school tentatively admits students from highest to lowest priority out of the pool of applicants in the current period and the set of students tentatively assigned to that school at the end of the previous period. Once the school is at capacity enrollment, all students that remain in the pool are rejected.
4. If all students have a school assignment, then the mechanism halts and the school assignments are finalized. If there is some student without a school assignment, return to step two.

Ethnicity	Number	% of Data
African American	938	37.3
White	457	18.2
Asian	304	12.1
Hispanic	755	30.0
Native American	11	0.4
Mixed-Other	48	1.9

Table 3: Ethnicity Demographics Across BPS Student Population

Throughout the analysis I exclude students that require special education programs as these students are effectively matched through a different market. There is also a large population of English as a second language (ESL) students in the BPS system. BPS ranks ESL students on a five point proficiency scale, and the least English proficient students, those that score a three or lower on the proficiency scale, are assigned to special ESL programs. Since these low scoring ESL students are also effectively matched in a distinct market from the rest of the students, I also exclude these students from my analysis. Once I exclude these students, I am left with 2,513 students and 28 high school programs (of which 10 programs are overdemanded). Given the preference profiles and priorities of each student, I can replicate the status quo assignment generated by the Gale-Shapley mechanism used by BPS for all but 8 of the students.¹²

Finally, I have a variety of administrative data on each student. The variables of interest are the ethnicity of the student, whether the student participated in a free school lunch program, and the zip code of the student’s home residence. BPS uses the six categories shown in Table 3 to classify the ethnicity of the students. A student can be part of the free school lunch program for a variety of reasons (e.g., as part of a state program such as SNAP benefits, certified as a migrant or homeless). On average 66.5% of the BPS students receive a free lunch. I use the participation of the student in a free lunch program as an indicator of low socioeconomic status.

The demographics of the overdemanded programs do not match the overall demographics of BPS as described in Table 3. This is the result of White and Asian students disproportionately either entering examination schools or choosing to attend an underdemanded school (presumably a neighborhood school). In addition, the demographic composition of the student body varies widely from school to school.

A similar message holds with regard to the socioeconomic demographics. An average of 78% of the students at the 10 overdemanded schools participate in a free school lunch program, which implies that the students attending the overdemanded programs are poorer than the overall BPS student population. The rate of free lunch participation in the 10 overdemanded programs varies

¹²The discrepancy between the BPS outcome and my attempt to replicate it is due to the fact that some low proficiency ESL students that I exclude from the data set are matched to non-ESL programs. Conversations with BPS enrollment staff could not resolve why these students were assigned to non-ESL programs.

from as low as 65.4% to as high as 90.2%, which reveals a high degree of heterogeneity across schools in terms of the socioeconomic demographics.

There is the capacity to enroll 688 students at the 10 overdemanded high school programs. Because BPS did not provide me the capacity of each program directly, I identify overdemanded school programs as those that rejected at least one student during the Gale-Shapley algorithm. Underdemanded schools accept all of the students that apply during the Gale-Shapley algorithm. For these schools I can only infer a lower bound on the capacity. Table 4 provides the average, minimum, and maximum percentage of the student body drawn from each of the six ethnic groups across the overdemanded programs.

In reality, whether or not a school is underdemanded is an endogenous outcome. Even though it is not an ideal assumption, I need to treat the set of underdemanded schools as exogenous since I do not have capacity figures for these schools. In effect, I am assuming that there is enough excess capacity at the underdemanded schools that my constrained optimization problems yield solutions that do not result in a school program that is underdemanded in my data to become overdemanded. I have several reasons for thinking that this is a reasonable assumption. First, the underdemanded school programs tend to be large — the smallest such school program enrolled 56 students and the average number of enrollees is over 156. In addition, when I compare the number of enrolled students in the 2011-2012 and 2012 - 2013 school years, there are large increases and decreases in total enrollment even at overdemanded schools. For example, one program at Brighton High is overdemanded with 173 enrolled students during the 2011 - 2012 school year, and yet 216 students are enrolled in the same program for the 2012 - 2013 school year. This suggests that there is flexibility to adjust the capacity of individual programs within a school. The total 9th grade enrollment at Brighton High across all programs increased by 27 students, which implies that the school has not been pushed to a physical constraint on the number of students that can enroll in the school. In summary, I think it is unlikely that a solution to one of my optimization problems will violate the capacity constraints of a school that is underdemanded in the 2011 - 2012 status quo.¹³

Finally, I do not have data on the final enrollments of the students. This is primarily of concern for understanding the outside options of the students, which would be evident if a student did not enroll in the school to which he or she is assigned. Therefore, one should interpret the analysis as pertaining to the intended school assignment.¹⁴

¹³If a solution to the optimization problem implies a new school becomes overdemanded, then it would be necessary to move the newly overdemanded school to the set of overdemanded schools and re-solve the optimization problem.

¹⁴Since I focus on the popular, overdemanded school programs in the BPS system, I believe that it is less likely that the outside option dominates one of these programs. However, in the absence of data on enrollment, I do not wish to take a strong stand on this — one could sketch a story wherein a student discovered an outside option only after being assigned to a very popular BPS school program.

Ethnicity	Average	Min	Max
African American	47.7	26.9	67.6
White	8.0	2	25.0
Asian	5.4	0	10.9
Hispanic	36.2	16.7	54.8
Native American	0.6	0	3.85
Mixed-Other	2.0	0	8.3

Table 4: Demographics at Overdemanded Schools

3 Matching in the Continuum Framework

I now describe the fundamentals of the constrained optimization problems I solve. For notational clarity, I adopt the language of the college admissions problem in this section and describe one side of the market as students and the other side of the market as colleges. A school choice mechanism matches each student to one college, but each college may be matched with multiple students. Since the continuum model will be treated as the limit of a sequence of finite economies in Section 6, I use the terms *limit model* and *continuum model* interchangeably.

I assume that there exists a finite set of student types \mathcal{S} , and there is a measure $\pi^{\mathcal{S}}(s)$ continuum of each type $s \in \mathcal{S}$. Each type $s = (v, \succeq_s) \in \mathcal{S}$ is characterized by verifiable traits, $v \in \mathcal{V}$, and a preference ordering over the finite set of colleges, $\succeq_s \in \mathcal{P}_{\mathcal{S}}$. I assume the set of verifiable traits \mathcal{V} is finite and each student-type's realization v is known to the mechanism designer. In my context v includes student ethnicity, socioeconomic status, the zip code of her residence, and the student's priority at each college.

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 of each student-type for seats at that college. I do not require that the priority orderings be strict. College c has a continuum of seats of total measure $q_c \in [0, 1]$, where q_c is interpreted as the fraction of the student population that can be enrolled in college c . A generic college-type is $c = (t_c, \succeq_c, q_c) \in \mathcal{C}$ where $t_c \in \mathcal{T}_C$ and $\succeq_c \in \mathcal{P}_C$. Since the walk-zone and regular seats at each program have different priority structures, I treat the different kinds of seats at each school as distinct colleges.

For colleges, \emptyset denotes an outcome wherein a seat is left unfilled. In the BPS system all student types are acceptable to all colleges, so $s \succeq_c \emptyset$ for all colleges c and student-types s .¹⁵ When \emptyset appears in a student's preference ordering, it refers to a generic underdemanded school program. For example, $c \succeq_s \emptyset$ means that student-type s prefers college c to her most preferred underdemanded school.

¹⁵It is easy to extend the framework to allow agents on each side to find some types on the other side of the market unacceptable. Since this is not relevant to the BPS application, I do not define the additional notation.

I break the discussion of the formalities of the constrained optimization problems I solve into three subsections. First, I provide a description of the notation I use to describe a match. Second, I describe the stability constraints I employ, I derive a useful simplification that allows me to describe the stable set as a union of convex subsets, and I describe the stability conditions used in my application. Third, I explain the incentive compatibility conditions I impose.

3.1 Matching Notation

I define a *match* as a function $x : (\mathcal{C} \cup \emptyset) \times (\mathcal{S} \cup \emptyset) \times \Delta(\mathcal{S}) \rightarrow [0, 1]$ 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 .

I now provide a demonstration of how my matching notation can be applied with either a finite number or a continuum of students. Consider a 500 student economy where $\mathcal{C} = \{c_1, c_2\}$, $\mathcal{S} = \{s_1, s_2\}$, $\pi^S(s_1) = \frac{3}{5}$, and $x(c_1, s_1) = x(c_2, s_1) = x(\emptyset, s_1) = \frac{1}{5}$. 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 continuum, $x(c_1, s_1) = \frac{1}{5}$ 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 notation 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. Formally,

$$\Pr\{\text{Student of type } s \text{ matched to college } c\} = \frac{x(c, s)}{\pi^S(s)} \quad (3.1)$$

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

$$\text{For all } s \in \mathcal{S}, \sum_{c \in \mathcal{C}} x(c, s) \leq \pi^S(s) \quad (3.2)$$

$$\text{For all } c \in \mathcal{C}, \sum_{s \in \mathcal{S}} x(c, s) \leq q_c \quad (3.3)$$

Equation 3.2 insures that the measure of seats occupied by students of type s does not exceed the total measure of these students. Equation 3.3 implies that the measure of students assigned to school c , $\sum_{s \in \mathcal{S}} x(c, s)$, does not exceed the school's capacity, q_c .

3.2 Stability

A pair (c, s) of college and student types is a *blocking pair* if:

$$\sum_{\{s': s' \succeq_c s\}} x(c, s') < q_c \text{ and } \sum_{\{c': c' \succeq_s c\}} x(c', s) < \pi^S(s) \quad (3.4)$$

The first condition implies that college c either has an empty seat or that student-type s has a higher priority than some student-type that has a seat at college c . The second condition implies that some student of type s is either unmatched or matched to a college to which c is strictly preferred by the student. A match x is *stable* if it does not admit blocking pairs.^{16,17}

The set of stable matches can be described as the set of $x \in [0, 1]^{|C|*|S|}$ that satisfy the following constraints:

$$\text{For all } s \in \mathcal{S}, \quad \sum_{c \in \mathcal{C} \cup \{\emptyset\}} x(c, s) = \pi^S(s) \quad (3.5)$$

$$\text{For all } c \in \mathcal{C}, \quad \sum_{s \in \mathcal{S} \cup \{\emptyset\}} x(c, s) = q_c \quad (3.6)$$

$$\text{For all } s \in \mathcal{S}, \quad \sum_{\{c \in \mathcal{C} : \emptyset \succ_s c\}} x(c, s) = 0 \quad (3.7)$$

$$\text{For all } (c, s), \quad \left(\pi^S(s) - \sum_{\{c' \in \mathcal{C} \cup \{\emptyset\} : c' \succeq_s c\}} x(c', s) \right) * \quad (3.8)$$

$$\left(q_c - \sum_{\{s' \in \mathcal{S} \cup \{\emptyset\} : s' \succeq_c s\}} x(c, s') \right) = 0$$

The first two conditions are feasibility constraints. Equation 3.5 insures that an individual student is matched with at most one college seat. Equation 3.6 implies that college c enrolls at most a fraction q_c of students. Including the outcome of being unmatched in the summations allows me to write equations 3.5 and 3.6 as equalities. Equation 3.7 requires that each student is matched to colleges that are weakly preferred to her most preferred underdemanded college. Equation 3.8 encapsulates all of the restrictions imposed by stability as per equation 3.4.

Proposition 1. *Equations 3.5, 3.6, 3.7, and 3.8 are necessary and sufficient for a match x to be stable.*

The following proposition reveals that there is at least one feasible and stable match. The existence result is known for markets with a finite number of agents. When the distribution of student-types and the college capacities 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 compute a feasible and stable match using a continuum analog of the deferred acceptance algorithm used by the Gale-Shapley mechanism. When the distribution of student-types or college capacities have irrational values, I use limit analysis techniques to argue that feasible and stable matches must exist since

¹⁶Stronger notions of stability such as *strong stability* and *super stability* have been proposed for matching models with indifferences. A stable match in my setting (in general) does not exist under these stronger notions of stability.

¹⁷My definition of stability is identical to the notion of *ex ante stability* proposed by Kesten and Unver [42]. In the limit model or an economy with multiple agents of the same type, the two notions coincide.

matches are known to exist in arbitrarily close economies with rational-valued distributions of students and colleges.

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

The description of the stable set provided by equations 3.5 - 3.8 is concise, but it presents practical challenges since the stability constraint is nonconvex and, in general, solving nonconvex optimization problems is NP-hard. Although I will not be able to avoid the nonconvexity, I can rewrite the stability condition as a family of sets of linear constraints, which in turn allows me to use powerful numerical methods to solve the problem.

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

$$\underline{x}_S(s) = \max\{c : c \preceq_s c' \text{ for all } c' \text{ such that } x(c', s) > 0\} \quad (3.9)$$

where the maximum is taken with respect to \succeq_s .¹⁸ If $\sum_{c \in \mathcal{C}} x(c, s) < \pi^S(s)$, then $\underline{x}_S(s) = \emptyset$. 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 $\underline{x}_C(c) \in \mathcal{S} \cup \{\emptyset\}$ to describe the lowest priority student-type assigned to college c in match x . If $\sum_{s \in \mathcal{S}} x(c, s) = q_c$, define \underline{x}_C as

$$\underline{x}_C(c) = \max\{s : s \preceq_c s' \text{ for all } s' \text{ such that } x(c, s') > 0\} \quad (3.10)$$

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

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 constrained optimization problems.

Lemma 1. *$(c, s) \in \Gamma$ is a blocking pair (i.e., equation 3.8 holds) if and only if $c \succ_s \underline{x}_S(s)$ and $s \succ_c \underline{x}_C(c)$.*

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 $s \succ_c s' \succeq_c \underline{x}_C(c)$. This in turn means Equation 3.8 must hold. The intuition for the reverse direction is similar.

Lemma 1 implies that I can decompose the stable set into convex subsets, each of which is defined by a choice of \underline{x}_C and \underline{x}_S that rules out blocking pairs. Given such a choice of \underline{x}_C and \underline{x}_S ,

¹⁸The 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 students of type s are matched.

I can write the stability conditions for the associated stable subset as a linear system:

$$\text{For all } s \in \mathcal{S}, \quad \sum_{c \in \mathcal{C} \cup \{\emptyset\}} x(c, s) = \pi^S(s) \quad (3.11)$$

$$\text{For all } c \in \mathcal{C}, \quad \sum_{s \in \mathcal{S} \cup \{\emptyset\}} x(c, s) = q_c \quad (3.12)$$

$$\text{For all } s \in \mathcal{S}, \quad \sum_{\{c \in \mathcal{C} : \emptyset \succ_s c\}} x(c, s) = 0 \quad (3.13)$$

$$\text{For all } s \in \mathcal{S}, \quad \sum_{\{c : c \succeq_s \underline{x}_S(s)\}} x(c, s) = \pi^S(s) \quad (3.14)$$

$$\text{For all } c \in \mathcal{C}, \quad \sum_{\{s : s \succeq_c \underline{x}_C(c)\}} x(c, s) = q_c \quad (3.15)$$

Equations 3.11 - 3.13 are unchanged and represent the feasibility and individual rationality constraints. Equations 3.14 and 3.15 represent the stability conditions implied by Lemma 1 for that choice of \underline{x}_S and \underline{x}_C . Equations 3.11 - 3.15 represents one set of stability constraints out of the family of such sets that define the full stable set.

Three things are of note. First, Equation 3.11 is weaker than Equation 3.14 and Equation 3.12 is weaker than Equation 3.15, so I do not include the first two equations in my constrained optimization problems. Second, solving an optimization problem with a convex objective subject to these linear constraints is a fast process. Third, there is no guarantee that Equations 3.11 - 3.15 admit a feasible point for every choice of \underline{x}_C and \underline{x}_S , but Proposition 2 insures that one such choice does.

My final task is to describe the stability conditions at work in my particular school choice application. There are four kinds of stability constraints I need to impose. First, many of the school programs (both open and walk-zone) have students that are insured a seat at the program. Denote the set of student-types with insured seats at school program c as $\mathcal{I}(c)$. Since there are not enough students with insured seats to fill any of the school programs, stability requires that $s \succ_c \underline{x}_C(c)$ for all $s \in \mathcal{I}(c)$. This means that for each college c and each $s \in \mathcal{I}(c)$ I have a single constraint of the form:

$$\text{For all } s \in \mathcal{I}(c), \quad \pi^S(s) = \sum_{\{c' \in \mathcal{C} \cup \{\emptyset\} : c' \succeq_s c\}} x(c', s)$$

This constraint requires that all student-types s that are insured a seat at college c must be assigned to a college that is weakly preferred by that student-type to college c .

Since the overdemanded schools are, by definition, overdemanded, we must impose the following constraint on all such colleges, which requires that all of the seats at the overdemanded schools are assigned to some student-type.

$$\sum_{s \in \mathcal{S}} x(c, s) = q_c$$

If one instead imposes the capacity conditions as inequalities, then it is possible that the best ODIC match might require leaving seats at an overdemanded program empty. This can be optimal if it slackens the incentive compatibility conditions, but such an outcome would violate stability.

The remaining two kinds of stability constraints only apply to walk-zone school programs. The form of the constraint depends on whether a walk-zone program is filled with students with walk-zone priority or if a student without walk-zone priority is admitted to the walk-zone program. Let $\mathcal{W}(c)$ refer to student-types that have walk-zone priority at school program c . If the walk-zone seats are filled by students with walk-zone priority, then $\underline{x}_C(c) = \mathcal{W}(c)$ and there is no restriction on \underline{x}_S . In this case the constraint for college c takes the form:

$$\text{For a walk-zone school program } c, \quad \sum_{s \in \mathcal{W}(c) \cup \mathcal{I}(c)} x(c, s) = q_c \quad (3.16)$$

If a walk-zone seat is assigned to a student from outside the walk-zone, then I have $\underline{x}_C(c) = \{s : s \notin \mathcal{W}(c)\}$ and $\underline{x}_S(s) \succeq_s c$ for $s \in \mathcal{W}(c)$. The relevant constraints take the form:

$$\text{For all } s \in \mathcal{W}(c), \quad \pi^S(s) = \sum_{\{c' \in \mathcal{C} \cup \{\emptyset\} : c' \succeq_s c\}} x(c', s) \quad (3.17)$$

Equation 3.17 requires that any student-type s with walk-zone priority at college c must be assigned to a school that she weakly prefers to college c .

Summarizing, for each of the two possible realizations of $\underline{x}_C(c)$ I get a different set of stability constraints. If $\underline{x}_C(c) = \mathcal{W}(c)$, I have

$$\begin{aligned} \text{For all } s \in \mathcal{I}(c), \quad \pi^S(s) &= \sum_{\{c' \in \mathcal{C} \cup \{\emptyset\} : c' \succeq_s c\}} x(c', s) \\ \sum_{s \in \mathcal{S}} x(c, s) &= q_c \\ \text{For walk-zone school program } c, \quad \sum_{s \in \mathcal{W}(c) \cup \mathcal{I}(c)} x(c, s) &= q_c \end{aligned}$$

If instead $\underline{x}_C(c) = \{s : s \notin \mathcal{W}(c)\}$ (and hence $\underline{x}_S(s) \succeq_s c$ for $s \in \mathcal{W}(c)$), I have

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

Given 10 overdemanded school programs, this means that there are 1,024 possible combinations of \underline{x}_C and \underline{x}_S (i.e., sets of stability conditions) I need to consider. I find that 24 of the 1,024 sets of constraints are feasible. The global optimum for my problem is then the optimum over these

24 subproblems.

3.3 Incentive Compatibility Constraints

I use *Ordinal Dominance Incentive Compatibility* (ODIC) as my definition of incentive compatibility. ODIC is based on the ordinal ranking of the students' preferences, which is convenient given my data on the students' ordinal preferences. Since the verifiable components of the students' types are known to BPS, I only need to impose incentive compatibility conditions that insure it is optimal for students to declare the nonverifiable components of their type truthfully. The ODIC conditions are defined as follows:

Definition 1. *The match $x(c, s)$ satisfies **ordinal dominance incentive compatibility** if for all types $s = (v, \succeq_s) \in \mathcal{S}$, $s' = (v, \succeq_{s'}) \in \mathcal{S}$ and $c \in \mathcal{C} \cup \{\emptyset\}$ I have*

$$\sum_{\{c': c' \succeq_s c\}} x(c', s; \pi^S) \geq \sum_{\{c': c' \succeq_{s'} c\}} x(c', s'; \pi^S) \quad (3.18)$$

Note that π^S appears on both sides of Equation 3.18, which means that I have assumed that a single agent's deviation has no effect on the aggregate distribution of actions in my limit game. This formulation of Equation 3.18 reflects an intuition that a single student's choice of a preference declaration should have a small effect on the match when there are many participants, and a single agent's action should have no effect on π^S in the limit game with an infinite number of participants. Justifying this intuitive argument is the key take-away from Section 6.1, wherein I prove that definition 3.18 is equivalent to an approximate equilibrium in the finite game with the approximation becoming arbitrarily precise as the number of students grows.

I have defined student preferences as ordinal rankings of the colleges, but one might have alternatively assumed students had a cardinal utility for being matched to each college and that the ordinal ranking was generated by these underlying cardinal utilities.^{19,20} In such a cardinal model, one might instead impose Bayesian incentive compatibility conditions that require that the expected utility from a truthful report exceeds the expected utility of any misreport. The ODIC conditions are equivalent to requiring that truthful reporting of the student's ordinal ranking of colleges be optimal for any cardinal utility function that could have generated the ordinal ranking. If one is willing to assume that the cardinal utilities of the students have full support, then the ODIC and the Bayesian incentive compatibility conditions are identical.

Before discussing the details of the numerical implementation of the ODIC constraints, I would like to take a moment to discuss why the ODIC constraints make the constrained optimization problem computationally intractable for the analogous problem with a finite number of students.

¹⁹The analysis of the relationship between the continuum model and large, finite school choice problem in Section 6 takes this cardinal view.

²⁰Given the stringency of the definition, one might question whether any feasible and stable match satisfies the ODIC condition. Helpfully, since any outcome of the Gale-Shapley mechanism with random tie-breakers satisfies ODIC, there must be a match that satisfies ODIC.

In a model with N agents, the ODIC constraint would be:

$$\text{For all } c \in \mathcal{C}, \quad \sum_{\{c':c' \succeq_s c\}} x(c', s; \pi^S) \geq \sum_{\{c':c' \succeq_s c\}} x(c', s'; \pi^S + \frac{1}{N} [\delta_{s'} - \delta_s]) \quad (3.19)$$

where δ_s is a measure placing probability 1 on student type s . The right-hand side of Equation 3.19 reflects the two effects of a nontruthful declaration when there is a finite set of students. First, as in Equation 3.18, a nontruthful declaration means the student receives the assignment of another type. The second effect, absent from Equation 3.18, is that the distribution of declared types changes.

To determine if Equation 3.19 is satisfied, I would need to compute the stable and incentive compatible match when the distribution of declarations is $\pi^S + \frac{1}{N} [\delta_{s'} - \delta_s]$. However, to compute a stable and incentive compatible match $x(c, \circ; \pi^S + \frac{1}{N} [\delta_{s'} - \delta_s])$, I would then need to solve another constrained optimization problem with ODIC constraints of the form:

$$\text{For all } c \in \mathcal{C}, \quad \sum_{\{c':c' \succeq_s c\}} x(c', s; \pi^S) \geq \sum_{\{c':c' \succeq_s c\}} x(c', s'; \pi^S + \frac{1}{N} [\delta_{s'} - \delta_s] + \frac{1}{N} [\delta_{s''} - \delta_{s'''}])$$

for all possible choices of $s'', s''' \in \mathcal{S}$. If this thought experiment is taken to its logical conclusion, one finds that solving the school assignment problem with a finite number of agents requires simultaneously finding an ODIC and stable match for *every* possible realization of π^S . This task is impossible with even a handful of students, and the key benefit of the continuum framework is that I avoid this computational problem.

Now let me turn to the details of how I implement the ODIC constraints in my constrained optimization problems. The key issue to address is minimizing the number of ODIC constraints I need to impose as these constraints make up the majority of the constraints for my problems. I make two simplifications. First, although the 28 school programs can be ranked a huge number of ways, by imposing stability I know that the outcome for a student is determined by how she ranks her most preferred underdemanded school program and any overdemanded programs preferred to that program. For example, suppose a student submits the ranking $c_3 \succ_s c_1 \succ_s c_4$ where c_1 is underdemanded and the other school programs are overdemanded. Since the student can do no worse than be admitted to c_1 in a stable match, I can recode this student's preference ranking as $c_3 \succ_s \emptyset$ where \emptyset denotes her most preferred underdemanded school. In the BPS context this simplification is particularly useful since many students only rank a few overdemanded school programs. Similarly, if the same student is insured a seat at c_1 , I can recode her preference ranking as $c_3 \succ_s c_1$ since she must be placed in one of these two school programs.

My second simplification is to relax the incentive constraints and argue that the relaxed constraints would be slack had they been imposed. Let \mathcal{S}_\exists denote the types of students that exist in the data (i.e., $s \in \mathcal{S}_\exists$ if $\pi^S(s) > 0$). For all types $s \in \mathcal{S}_\exists$, 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 $s' \notin \mathcal{S}_\exists$ or for

type $s' \notin \mathcal{S}_\exists$ to declare her preferences nontruthfully.

Before formally justifying this relaxation, consider what would happen if I solved the model with the full $|\mathcal{S}|(|\mathcal{S}| - 1)10$ set of ODIC constraints. My objective functions are based only on the assignment of student-types $s \in \mathcal{S}_\exists$. Consider a type of student $s' \notin \mathcal{S}_\exists$. The solution to my optimization problem will assign an outcome to type s' that slackens the incentive constraints for the student-types $s \in \mathcal{S}_\exists$ as much as possible.

So what value of $x(c, s')$, $s' = (v', \succeq_{s'}) \notin \mathcal{S}_\exists$, does my optimization problem implicitly set? Recalling that the students cannot nontruthfully declare their verifiable traits, let $\mathcal{S}_\exists(v') = \{s = (v, \succeq_s) \in \mathcal{S}_\exists : v = v'\}$ denote the set of types in \mathcal{S}_\exists that can be mimicked by s' . If $\mathcal{S}_\exists(v')$ is nonempty, then

$$x(c, s') = \max \{x(c, s) : s \in \mathcal{S}_\exists(v')\} \quad (3.20)$$

where the maximum is with respect to $\succeq_{s'}$. There are two key points. First, if $s \in \mathcal{S}_\exists$ mimics $s' \notin \mathcal{S}_\exists$, then she could obtain the same outcome by mimicking some type $s'' \in \mathcal{S}_\exists$. Second, Equation 3.20 implies that any type $s'' \notin \mathcal{S}_\exists$ that can mimic s' must receive an outcome at least as desired as that assigned to s' . Together, these points imply that my relaxation of the ODIC constraints is without loss of generality.

If $\mathcal{S}_\exists(v')$ is empty, then in a sense I do not need to define the assignment for these types since none of the participants in the mechanism can mimic these types. For completeness however, I now describe a possible choice of $x(c, s')$ for $s' = (v', \succeq)$ if $\mathcal{S}_\exists(v')$ is empty. Let $\mathcal{C}(s')$ denote the union of the set of underdemanded colleges; any colleges to which s' is insured a seat; and any colleges at which s' has walk-zone priority and where the college admits students without walk-zone priority to the walk-zone seats. Assign any student with type s' to her favorite school in $\mathcal{C}(s')$. Any type s'' that can mimic s' must have $\mathcal{C}(s') = \mathcal{C}(s'')$ since $\mathcal{C}(s')$ is determined by the verifiable component of s' . This means that s'' can only be hurt by nontruthfully claiming to have type s' .

As a final point, even with these simplifications, the set of ODIC constraints is large. I solved the constrained optimization problems using the Gurobi optimization package for Matlab.²¹ This software speeds up the solution of the problems over classical operations research algorithms by removing many redundant constraints and encoding the problem in a form favorable for the solver algorithms. Even on a laptop, solving one of the 24 subproblems takes less than 30 seconds.

However, the size of the set of ODIC constraints does raise the important of whether the constrained optimization problem scales well. I would be skeptical of attempts to apply my techniques to the New York City school choice problem given the hundreds of programs.²² However, there are many school systems (e.g., San Francisco, Chicago) that have a small enough number of schools that one could apply my model. For example, San Francisco has only 17 high schools and some of these may be underdemanded. Another example is Chicago, which currently has 10 selective enrollment high schools. In school systems of this size, I see no reason why my methods could not be applied.

²¹Gurobi version 6.5.1 and 7.0 and Matlab 2016a.

²²That being said I do not know how many NYC school programs are overdemanded in practice.

4 Welfare Maximization in the Boston High School Match

The goal of this section is to find the student-welfare optimal school choice mechanism for the BPS high school match and understand the extent to which the Gale-Shapley mechanism maximizes student welfare. Throughout I use the average rank of the school to which the students are assigned as my metric for student welfare. Section 4.1 compares the student welfare achieved by an optimal mechanism to the welfare achieved by the Gale-Shapley mechanism employed by BPS. In Section 4.2, I study the benefits of allowing the mechanism to condition on the students' demographics. I also provide a detailed breakdown of the distribution of the ranks of schools to which the students are assigned. All of my statistics are accompanied by 95% confidence intervals computed using bootstrap sampling.²³

There are several classes of student that I do not include when computing my welfare statistics. First, BPS uses formal exams for admissions to three exam schools: Boston Latin Academy, Boston Latin School and the John D. O'Bryant School of Mathematics and Science. Of the 2,513 students I consider, 753 students are assigned through the examination process. An additional 461 students rank an underdemanded school with empty seats as their most preferred option. Another 116 of the students rank a school to which they are insured admission as their top choice. Stability requires that these students be automatically assigned to these schools. In order to better describe the impact I am having on the students whose assignment I can influence, the statistics I provide in this section are based on the 1,183 remaining students. In Section 5 I return to considering the full set of students.

4.1 Constrained Welfare Maximization and The Gale-Shapley Outcome

The goal of this section is to evaluate the scope for improving student welfare through the use of a stable, ODIC mechanism. Since the Gale-Shapley mechanism does not condition assignments on student demographics (other than priorities), the stable, ODIC mechanisms that I analyze do not condition on student demographics either. In addition to finding the level of welfare at the optimum, a second goal is to understand which of the constraints restricts my ability to achieve higher levels of welfare. At the close of this section, I use the shadow prices of the constraints to assess how additional capacity ought to be distributed as it becomes available.

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

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

²³I computed 550 bootstrap samples in total, and 9 of these samples were infeasible because more students were insured a seat in a small school program than there was capacity to admit. All of the bootstrap results were computed using the same set of 541 feasible bootstrap samples.

Best Stable, ODIC match	Gale-Shapley	Worst Stable, ODIC match
2.597	2.617	2.686
[2.531, 2.678]	[2.547, 2.690]	[2.610, 2.7626]

Table 5: Welfare Analysis

Best Match	Best Stable Match	Best ODIC match	Best Stable, ODIC match
2.215	2.250	2.588	2.597
[2.155, 2.290]	[2.187, 2.327]	[2.519, 2.665]	[2.531, 2.678]

Table 6: Welfare Analysis

My welfare objective is the average expected rank over the set of students:

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

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 .

Table 5 describes the average welfare of the best and worst stable and ODIC matches as well as the welfare generated by the match produced by the Gale-Shapley mechanism.²⁴ 95% confidence intervals are provided beneath the relevant statistic. The difference between the best and worst stable and ODIC matches is 0.1, which is equivalent to moving 118 students to a school one rank higher on their preference lists. The Gale-Shapley mechanism recovers 78% percent of the welfare gap between the best and worst stable and ODIC matches.

I now turn to analyzing the importance of the different constraints I have imposed on my mechanism. Table 6 displays the effect of removing the stability constraints and/or ODIC constraints. In all cases I continue to insist that the capacity constraints hold with equality, which I impose as a minimal notion of stability. One can conclude from Table 5 that it is the ODIC constraints, rather than stability conditions, that limit my ability to improve student welfare. In fact, the stability conditions appear to have almost no effect on student welfare, which is reflected in two statistics in Table 5. First, one can impose stability on top of the ODIC constraint at a welfare cost of less than 0.01, which is the equivalent of moving 11 students to a school one rank higher in their preference

²⁴I compute the Gale-Shapley outcome in a limit economy using the techniques developed by Azevedo and Leshno [10].

list. Second, the gap between the best match and the best stable match is less 0.04, which is equivalent to moving 47 students to a school one rank higher in their preference list. In contrast, the difference between the welfare achieved in the best match and the best ODIC match is 0.373, a gap that is equivalent to moving 441 students to a school one rank higher in their preference list.

As mentioned in the introduction, it is known that using random tie-breakers can result in a welfare loss. 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. In order to assess the cost of strategy-proofness, the authors compare the status quo outcome with the outcome that is generated by applying Erdil and Ergin’s [24] manipulable stable improvement cycles to the status quo outcome. The authors conclude that the welfare costs of incentive compatibility are small since the improvement generated by using stable improvement cycles is small. The authors then compare the status quo outcome with a match generated by applying Gale’s top-trading cycles algorithm to the status quo outcome, which yields an efficient, strategy-proof, and (potentially) unstable outcome. Again, the authors find that few students are able to improve their assignment, so the authors conclude that there is a low cost of stability in the BPS school assignment system.

My results support the conclusion of Abdulkadiroğlu et al. [4] that the welfare costs of stability are low, but unlike Abdulkadiroğlu et al. [4] I find that incentive compatibility has significant costs. It is possible that the difference between our conclusions stems from the fact that I am analyzing data from years after the publication of Abdulkadiroğlu et al. [4]. Another explanation for our differing conclusions is that I use a stricter welfare metric. Abdulkadiroğlu et al. [4] find that moving from the Gale-Shapley outcome to an efficient match results in a change equivalent to moving 9.7 students to a school one rank higher in their preference list. I find that the difference between the Gale-Shapley mechanism and the welfare optimal (non-stable, non-ODIC) mechanism is equivalent to moving more than 475 students to a school one rank higher in their preference list. It could be that the Gale-Shapley results in an assignment near the Pareto frontier, but that assignment could still be quite suboptimal with respect to my average rank metric.

Finally, one can use the shadow prices of the constraints to evaluate how to allocate 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 programs offered by Snowden International. This means that if there are resources available to expand a school’s capacity, 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.

4.2 The Costs and Benefits of Conditioning on Demographics

The Gale-Shapley mechanism’s outcome depends only on the priorities of and preferences declared by the students along with the random tie-breakers assigned to each student. Now suppose the mechanism treats the zip codes of the students’ residences as a verifiable trait and conditions the

	Unconditional	Ethnicity	Zip Code	All Demographics
Best ODIC Match	2.588 [2.519, 2.665]	2.5076 [2.437, 2.565]	2.464 [2.369, 2.514]	2.346 [2.247, 2.390]
Best Stable, ODIC Match	2.600 [2.531, 2.678]	2.5234 [2.442, 2.583]	2.485 [2.388, 2.535]	2.371 [2.274, 2.424]
Worst Stable, ODIC Match	2.686 [2.61, 2.762]	2.751 [2.697, 2.852]	2.831 [2.788, 2.934]	2.962 [2.923, 3.072]
Worst ODIC Match	2.711 [2.639, 2.785]	2.776 [2.728, 2.878]	2.854 [2.822, 2.963]	3.006 [2.974, 3.120]

Table 7: Effect of ODIC Conditioning

outcomes on this trait. Since the class of mechanisms that condition on the students' zip codes is strictly larger than the class of mechanisms that do not, the optimal mechanism that conditions on the students' zip codes must perform (weakly) better than a mechanism that does not condition on the student demographics at all. The goal of this section is to study the improvement in student welfare that can be reaped by allowing the mechanism to condition on different aspects of the students' verifiable demographics.

I consider four classes of mechanisms in Table 7. All of these mechanisms condition outcomes on the students' priorities and declared preferences, but the mechanisms differ in the extent to which outcomes are also conditioned on student demographics. The first column (*Unconditional*) analyzes mechanisms that do not condition on the student demographics. The second (*Ethnicity*) and third (*Zip Code*) columns consider mechanisms that condition on student ethnicity (as recorded by BPS) and the zip code of the student's home residence. The final column (*All Demo.*) considers mechanisms that can condition on the students' ethnicity, zip code, and free school lunch status.

Table 7 presents the best and worst ODIC matches as well as the best and worst stable and ODIC matches for each form of conditioning. Recall from Table 6 that one could achieve a welfare of 2.250 under the best stable, non-ODIC match. Our analysis implies that if the mechanism conditions on the full set of demographic information available, then the resulting welfare of 2.371 recovers over 70% of the gap between the welfare provided by the status quo Gale-Shapley mechanism and the welfare possible if one completely ignores the ODIC conditions. The welfare improvement from implementing the best stable, ODIC match that conditions on all the demographics relative to the outcome of the Gale-Shapley mechanism is equivalent to moving 291 students to a school one rank higher in their preference list.

Now let us turn to a comparison of the different mechanisms in Table 7. Since the second and third columns consider mechanisms that condition on strictly more information than the first column, the outcomes from the optimal mechanism described in the second and third columns must represent a welfare improvement relative to mechanisms that do not condition on these demographic

Condition on ...	None	Ethnicity	Zip Code	All Demographics
None	1	0.963	0.965	0.940
Race	0.963	1	0.948	0.949
Zip Code	0.965	0.948	1	0.965
All Demo.	0.940	0.949	0.965	1

Table 8: Correlation between Welfare of Stable, ODIC matches

variables. Similarly, since the final column describes the outcome of an optimal mechanism that conditions on more information than any other mechanism in Table 7, the optimal mechanism from this group must improve on all of the others I consider.

Despite being able to partially order the effectiveness of the mechanisms based on the information on which the outcomes are conditioned, it is still of interest to determine whether these differences are significant. Paying attention to the point estimates alone, it appears that a sizable decrease in (i.e., improvement of) the welfare metric can be had by moving from a default mechanism that does not condition on any demographics to a mechanism that conditions on either ethnicity or zip code. Another significant improvement of between 0.12 and 0.16 can be made by moving from conditioning on either race or zip code alone to conditioning on all of the demographic information.

A direct reading of Table 7 might suggest that while the point estimates differ, the standard errors are sufficiently large that the choice of mechanisms does not yield a statistically significant difference in welfare. Table 8 provides the correlation coefficients of the welfare statistics across the bootstrap runs, and the high correlation shows that this naive reading is misleading. To get a more complete picture of the benefit of conditioning on student demographics, I compare the outcomes of different mechanisms within each bootstrap sample. I find that the optimal mechanism that conditions on student zip code does better than the optimal mechanism that conditions on race in every bootstrap run. This is likely due to the fact that there are 30 zip codes covered by BPS and only six ethnic groups, and zip codes are relatively good predictors of ethnic group given the segregation across Boston. Roughly speaking, this means that conditioning on zip codes provides a finer partition of the students than conditioning on ethnicity.

While the statistics presented above focused on the average rank received by the students, Table 9 displays the number of students who are assigned to schools at a particular rank of their preference relation under different mechanisms. The mechanisms I consider are the status-quo Gale-Shapley (GS) mechanism, the optimal mechanisms that does not condition on demographics (NC), and the optimal mechanism that conditions on all demographics (AD).²⁵ The first three columns of Table

²⁵Since all of the mechanisms were solved using continuum methods, none of the mechanisms produce an integer

\succ Rank	GS	NC	AD	$\Delta(\text{NC} - \text{GS})$	$\Delta(\text{AD} - \text{GS})$
1st	443.8 [426.1, 458.8]	444.1 [421.6, 456.5]	467.5 [443.3, 487.4]	0.3 [-13.3, 5.2]	23.7 [7.5, 38.4]
2nd	196.3 [166.5, 224.7]	204.3 [178.0, 238.7]	249.0 [219.1, 298.3]	8.04 [1.1, 23.1]	52.8 [42.2, 82.7]
3rd	248.9 [228.1, 269.3]	247.1 [228.9, 269.1]	249.9 [223.2, 273.7]	-1.80 [-4.6, 5.5]	0.96 [-14.4, 13.5]
4th	155.7 [137.6, 174.4]	151.3 [133.0, 169.9]	141.2 [115.1, 160.3]	-4.42 [-8.8, -0.8]	-14.51 [-31.0, -5.2]
5th	81.1 [69.6, 94.9]	81.7 [68.8, 95.4]	53.1 [40.2, 68.0]	0.66 [-2.6, 2.0]	-28.03 [-37.8, -18.4]
6th	41.2 [33.7, 50.0]	39.5 [31.0, 48.3]	16.6 [8.0, 21.4]	-1.66 [-4.6, -0.2]	-24.58 [-34.4, -20.0]
7th	11.6 [7.2, 15.9]	11.0 [6.4, 15.2]	4.1 [0., 6.0]	-0.60 [-2.6, 0.2]	-7.47 [-12.3, -4.8]
8th	4.76 [2.2, 7.6]	4.2 [1.8, 6.8]	2.0 [0, 3.0]	-0.54 [-1.7, -0.1]	-2.76 [-6.2, -1.5]
9th	0.1 [0, 0.4]	0.1 [0, 0.4]	0 [0, 0]	0.01 [0, 0.03]	-0.11 [-0.4, 0]

Table 9: Rank Assignments Under Different Mechanisms in Terms of the Number of Students

9 describe the raw number of student assigned a school at each respective rank of their preference list. The fourth column displays the difference between the Gale-Shapley mechanism and the optimal mechanisms that does not condition on demographics, and the fifth column displays the difference between the Gale-Shapley mechanism and the optimal mechanisms that conditions on all demographics.

The mechanism that conditions on all of the student demographics causes a significant improvement in the rank of the assignment for many students. Relative to the Gale-Shapley mechanism, 23.7 additional students obtain their top choice and 52.8 students achieve their second choice, and many of these students received their fifth or sixth ranked choice under the Gale-Shapley mechanism. The optimal mechanism that does not condition on student demographics yields an improvement over the status quo assignment, but the effect is much smaller — less than 9 extra students get their first or second ranked choice, and most of these students received their third or fourth place choice under the Gale-Shapley mechanism.

There are two concerns regarding allowing the mechanism to condition on the students' demographics. First, I allowed the mechanisms to condition on the demographics of the students under the principle that these traits are verifiable and need not be revealed by the students. Un-

number of students at each rank.

fortunately, there are problems with the verifiability of these data. For example, parents could choose to reside in a zip code that would, in the parents’ opinion, lead to favorable treatment by the mechanism. Parents already had an incentive to engage in behavior like this under the status quo system since the student’s residence determines at which schools that student has walk-zone priority. If families adapt their choice of residence to “manipulate” the mechanism, the distribution of preferences and demographics will change, meaning that the results provided here may not provide a reliable estimate of the long-run distribution of outcomes.²⁶ This is one of several potential nonstationarities in the market, which are discussed at greater length in Section 6.3.

There is also a fairness issue that ought to be discussed when considering conditioning on demographic data. Making the school assignment depend on immutable and exogenous characteristics of a student (e.g., ethnicity) for reasons (e.g., welfare) that have a tenuous intrinsic connection to the characteristic violates an intuitive notion of equal-treatment-of-equals.²⁷ One can view the welfare gains that conditioning on demographics allows as the *Welfare Cost of Fairness* — if the designer is willing to violate equal-treatment-of-equals, the welfare cost of fairness is the degree to which student welfare can be enhanced.

Certain violations of equal-treatment-of-equals are more acceptable than others. For example, conditioning on the student’s zip code alone would insure equal-treatment-of-equals within a zip code, but would result in violations of this fairness property across zip codes. While zip code information may also be viewed as tenuously connected to the goal of welfare, there is a precedent for using geographical data in the school assignment process since the walk-zone priority system already treats students from different parts of Boston differently. Since zip-code based assignment may be more acceptable than conditioning on all of the demographics, I also provide results for mechanisms that satisfy equal-treatment-of-equals within each zip code when analyzing racial diversity goals (see Section 5, Table 10 and Figure 2).

5 Distributional Goals

In this section we study three distributional goals: ethnic diversity, socioeconomic diversity, and the encouragement of neighborhood schools. Section 5.1 defines my metrics for these goals, and Section 5.2 analyzes the extent to which we can optimize with respect to one of our goals while doing no worse than the Gale-Shapley mechanism on the remaining objectives. The point of this analysis is to determine how much we can improve on the status-quo mechanism at no cost with respect to our remaining desiderata. I close with an analysis of how effectively one can optimize with respect to a convex combination of student welfare, ethnic diversity, and the encouragement of neighborhood schools. I then provide a graphical analysis of the possibility frontiers for achieving different

²⁶Ethnicity presents additional issues. First, ethnicity is a vaguely defined concept, which means that even individuals without interest in manipulating the mechanism could disagree about what ethnicity to assign to a student. Second, it is difficult to imagine a school district recording an ethnicity for a student that did not conform to the ethnicity declared by a parent, which means that the ethnicity data is potentially manipulable.

²⁷In addition, conditioning on ethnicity stands on shaky legal footing given the previous court ruling that sharply limit how BPS may use information about the students’ ethnicities in the school assignment process.

combinations of our goals. Section 5.3 provides the possibility frontiers that describe the tradeoffs between student welfare and our metrics for ethnic and socioeconomic diversity, and Section 5.4 considers the trade-offs between encouraging neighborhood schools, diversity, and welfare.

5.1 Metrics

I define a school c as diverse if the demographics of the student body at school c closely resemble the demographics of the student population of BPS as a whole, and a school becomes more diverse if its demographics shift to more closely resemble the BPS population at large. My objective function is the average diversity across all of the overdemanded BPS schools.²⁸

Let \bar{p}_g denote the fraction of the BPS population in demographic group g , where g is one of the six ethnic groups recorded on Table 3. Let $p_{g,c}(x)$ denote the fraction of the students at school c in demographic group g under school assignment x . Let $e(s)$ denote the ethnic group that student type s belongs to, and $\mathbb{1}\{e(s) = g\}$ be an indicator that takes the value 1 if $e(s) = g$ and 0 otherwise. With this notation in hand, I can write:

$$p_{g,c}(x) = \sum_{s \in \mathcal{S}} \mathbb{1}\{e(s) = g\} \frac{x(c, s)}{q_c}$$

I then define the diversity of school c , denoted $d_c(x)$, as the average of the squared difference between $p_{g,c}(x)$ and \bar{p}_g across groups g ²⁹:

$$d_c(x) = \sum_g \bar{p}_g (p_{g,c}(x) - \bar{p}_g)^2$$

Lower values of $d_c(x)$ mean a school's student population more closely resembles that of BPS as a whole. Finally, my metric of aggregate school diversity, $D(x)$, is the average of d_c over the set of schools.

$$D(x) = \frac{1}{\sum_c q_c} \sum_c q_c d_c(x) = \frac{1}{\sum_c q_c} \sum_c q_c \sum_g \bar{p}_g (p_{g,c}(x) - \bar{p}_g)^2$$

In my tables and figures I have, for expositional clarity, used the following average absolute deviation metric for diversity:³⁰

$$D_{Abs}(x) = \frac{1}{\sum_c q_c} \sum_c q_c \sum_g \bar{p}_g \|p_{g,c}(x) - \bar{p}_g\| \quad (5.1)$$

The metric I use to measure socioeconomic diversity is similar in form to the metric for ethnic diversity. First I compute the variable $f_c(x)$ that denotes the fraction of students at school c that receive a free school lunch in school assignment x . I use the notation $l(s) = 1$ to denote the event

²⁸Enforcing stability severely restricts my ability to manipulate the student bodies of underdemanded schools since any student can claim a seat at any underdemanded school in lieu of her current assignment. Because of the restrictions stability imposes, I focus on the diversity of the overdemanded schools.

²⁹Alternatively, if one thinks of $\bar{\mathbf{p}}$ and \mathbf{p}_c as six-element vectors, one might have used the square of the Euclidean norm as metric of diversity at a school. My results would remain qualitatively the same.

³⁰While I would like to have used the absolute difference criterion in the optimization problem, it is computationally easier to solve a quadratic program.

that students of type s are eligible for a free lunch, so I can write:

$$f_c(x) = \sum_{s \in \mathcal{S}} \mathbb{1}\{l(s) = 1\} \frac{x(c, s)}{q_c}$$

$\bar{f} = 0.665$ is the percentage of students in the BPS system that receive a free school lunch. The socioeconomic diversity of school c , denoted $s_c(x)$, is defined as the squared difference between $f_c(x)$ and \bar{f} .

$$s_c(x) = (f_c(x) - \bar{f})^2$$

Lower values of $s_c(x)$ mean a school's student population more closely resembles that of the BPS as a whole in terms of socioeconomic diversity. Finally, my metric of aggregate socioeconomic diversity of school assignment x , $S(x)$, is the average of $s_c(x)$ over the set of schools.

$$S(x) = \frac{1}{\sum_c q_c} \sum_c q_c s_c(x) = \frac{1}{\sum_c q_c} \sum_c q_c (f_c(x) - \bar{f})^2$$

For expositional clarity, my tables and figures use the following absolute deviation metric of socioeconomic diversity:

$$\frac{1}{\sum_c q_c} \sum_c q_c \|f_c(x) - \bar{f}\|$$

The final distributional goal I consider is the encouragement of neighborhood schools. Let $n_c(x)$ denote the fraction of school c 's student body drawn from the school's walk-zone. My metric for the encouragement of neighborhood schools is the average of $n_c(x)$ across the set of schools:

$$N(x) = \frac{1}{\sum_c q_c} \sum_c q_c n_c(x)$$

5.2 How Much Can One Improve On Gale-Shapley?

Our first goal in this section is to assess the degree of suboptimality of the Gale-Shapley mechanism by computing matches that are optimal with respect to one of our goals — welfare, diversity, and encouraging neighborhood schools — while doing no worse than the Gale-Shapley mechanism on the remaining two. Second, I maximize a convex combination of all three goals at once to better understand the tensions between these goals.

Turning to my first analysis, I compute the mechanism that is optimal with respect to one of my metrics while performing no worse than the status quo Gale-Shapley mechanism on the remaining two. I refer to these mechanisms as *constrained optimal mechanisms* because I literally add constraints to the optimization problems defining each mechanism that ensure that the solution performs no worse than the Gale-Shapley assignment on all three metrics. Table 10 presents the outcomes for the constrained optimal mechanism for each metric — note that each metric has a distinct constrained optimal mechanism. Recall that lower values represents an improvement for all of my metrics except for encouraging neighborhood schools.

Metric	Gale-Shapley (GS)	Constrained Optimal Mechanism (CON)	Δ (GS - ON)
Welfare	2.62 [2.55, 2.69]	2.38 [2.28, 2.43]	0.24 [0.288, 0.239]
Ethnic Diversity	0.0994 [0.092, 0.113]	0.047 [0.037, 0.063]	0.052 [0.063, 0.043]
Ethnic Diversity Using Zip Code	0.0994 [0.092, 0.113]	0.048 [0.038, 0.063]	0.051 [0.063, 0.043]
Socioeconomic Diversity	0.120 [0.104, 0.140]	0.008 [0, 0.035]	0.112 [0.124, 0.086]
Encouraging Neighborhood Schools	0.221 [0.199, 0.241]	0.358 [0.323, 0.391]	-0.137 [-0.111, -0.1582]

Table 10: Optimal Mechanisms For Each Metric

What our analysis reveals is that it is possible to make a difference that is both economically and statistically significant on any one of our metrics while continuing to do no worse than the Gale-Shapley mechanism on the remaining two. For example, Table 7 shows that the optimal outcome from a welfare maximizing, ODIC, and stable mechanism is an average rank of 2.37. Table 10 shows that if we wish to do no worse than the Gale-Shapley mechanism on our diversity metric, we can achieve an average rank of 2.38. The gap of 0.01 is equivalent of moving 11.8 students to a school one rank more preferred in their preference list, and this represents the cost of maintaining the Gale-Shapley level of diversity.

The gains along the other dimensions are even more sizable. Under the Gale-Shapley mechanism an average of 78% of the students enrolled in the schools participate in a free school lunch program, whereas the average across the BPS population is only 66.5%, and the fraction of students receiving a free school lunch at the overdemanded schools ranges from 66.4% to 90.2% of the entering student body. The outcome created by the constrained optimal mechanism that maximizes socioeconomic diversity results in a match where 66.4% of the students at each school receive a free school lunch at 9 of the 10 programs. In other words, the socioeconomic demographics at these schools exactly match those of the overall BPS student population. The lone exception is the TechBoston Academy. Since the majority of the seats at this school are allocated to students that are insured a seat, there is little that my mechanism can do about the demographics at this school. As a result, 71.5% of the students at this school receive a free lunch in the constrained optimal mechanism.

To get some sense of what the improved ethnic diversity means in terms of the individual schools, Table 11 presents statistics on the demographics at each school under the constrained optimal mechanism for maximizing ethnic diversity.³¹ The first columns refer to the constrained optimal mechanism, while the second set of columns refers to the Gale-Shapley mechanism.³² Confidence

³¹Tables 19 and 20 in the Appendix provide raw counts of the number of students from each demographic assigned to each of the 10 overdemanded programs under each of the mechanisms.

³²The statistics for the status quo Gale-Shapley mechanism are not quite the same as Table 4 because we computed

Ethnicity	BPS Avg.	Constrained Optimal Mechanism			Gale-Shapley Mechanism		
		Average	Min	Max	Average	Min	Max
African American	37.3	43.2 [41.1, 45.6]	36.5 [25.0, 38.4]	60.8 [49, 73.5]	48.3 [45.9, 50.9]	32.5 [13.1, 34.2]	70.7 [64.0, 77.8]
White	18.2	12.5 [10.4, 14.0]	7.2 [2.6, 9.7]	18.3 [17.1, 25.0]	7.7 [6.2, 9.1]	1.8 [0.3, 3.2]	17.3 [12.4, 33.1]
Asian	12.1	8.6 [7.3, 10.5]	4.2 [0, 6.6]	12.1 [10.8, 25]	6.4 [5.1, 7.9]	1.5 [0, 2.7]	13.9 [11.7, 28.0]
Hispanic	30.0	32.5 [30.6, 34.3]	26.9 [18.9, 29.9]	37.2 [34.2, 41.7]	35.4 [32.7, 37.9]	20.4 [12.7, 25.2]	51.7 [45.5, 61.2]
Native American	0.4	1.1 [0.3, 1.5]	0 [0, 0]	2.3 [1.0, 4.8]	0.6 [0.1, 1.0]	0 [0, 0,]	1.2 [0.6, 3.3]
Mixed- Other	1.9	2.2 [1.3, 3.1]	0 [0, 0]	8.3 [3.1, 25]	1.6 [1.0, 2.3]	0.3 [0, 0.3]	8.3 [2.9, 25.0]

Table 11: Demographics at Overdemanded Schools Under the Constrained Optimal Mechanism

intervals were computed using bootstrap sampling. Table 19 in the appendix provides school-by-school demographic statistics that provide a different perspective on the outcome generated by the constrained optimal mechanism for maximizing ethnic diversity.

The constrained optimal mechanism for maximizing ethnic diversity does so through two channels. First, the mechanism tries to make the average fraction of the students in each demographic group that are assigned to the overdemanded schools match the fraction of that demographic group in the population of BPS students. I refer to this as the *gap in means* for each demographic group. For example, the constrained optimal mechanism results in classes that are 43.2% African-American on average as compared to 48.3% under the Gale-Shapley mechanism and 37.3% for the BPS student population. This means that the gap in means has dropped from 11%(= 48.3% – 37.3%) to 5.9%(= 43.2% – 37.3%). The gap in means for the other three large ethnic groups are reduced by a similar fraction.

The second channel through which the optimal mechanism improves diversity is to limit the extremity of the difference between schools with the highest and lowest fraction of their student bodies drawn from each ethnic group. I refer to this as the *across school gap*, and the demographics of the schools resemble each other more closely as this gap shrinks. This reduction of the across school gap is most visible for the Hispanic student population. The constrained optimal mechanism results in schools with student bodies that range from 26.9% to 37.2% hispanic. Under the Gale-Shapley mechanism this range is more than three times larger (20.4% to 51.7%). The fraction of white students ranges from 7.2% to 18.3% under the constrained optimal mechanism, which is almost a 50% decrease of the range generated by the Gale-Shapley mechanism. A similar magnitude

the numbers for Table 11 using the continuum analog of the discrete Gale-Shapley mechanism that is used in reality.

of effect is also found for the proportion of Asian and African-American enrollees at these schools.

Following the U.S. District court ruling in *McLaughlin v. Boston School Committee* (1996), BPS cannot use race as a factor in determining admission to the prestigious Boston Latin school. BPS stopped using ethnicity as a factor in any component of the assignment mechanism in 1999. One might be concerned that achieving the diversity gains shown in Table 11 requires conditioning the assignment on the students' ethnicity, which would go against the 1996 court ruling. In order to assuage this concern, the row labeled in "Ethnic Diversity Using Zip Code" in Table 10 maximizes racial diversity under the requirement that the assignment mechanism treat all students from the same zip code symmetrically. As expected, requiring symmetric treatment of students within a given zip code slightly reduces my ability to achieve racially diverse schools. In other words, using zip code as a proxy for ethnic background is an effective, if indirect, tool for encouraging school diversity.

Finally, consider the encouragement of neighborhood schools. BPS believed this objective was sufficiently important to warrant making the ability to enroll in neighborhood schools an important component of the priority structure. Even after imposing the BPS walk-zone priorities, the Gale-Shapley mechanism only manages to fill 22% of the seats at these schools with students from the walk-zones. The walk-zone seats are filled with students from the walk-zone for seven of the ten schools, but these include the five smallest programs and one larger school, TechBoston Academy. The optimal constrained mechanism manages to fill 36% of the seats with students from the walk-zones. The difference between the outcome under the Gale-Shapley outcome and the optimal constrained mechanisms is so large that there is a 95% chance the optimal mechanism fills at least 11% more seats with students from the schools' walk-zones.

Now I would like to consider how well one can do on all three metrics at the same time. I study this by choosing an arbitrary convex combination of the three metrics to maximize, and I refer to the solution to this problem as the *solution to the convex combination*. From a market design perspective, the ideal weighting between these metrics is a matter for public policy makers. Although this means that the exact weight placed on each metric is ad hoc, I did experiment to find a combination that yielded significant improvements across all three relative to the status quo.³³

The values across the three metrics of interest as well as bootstrap standard errors are presented in Table 12. There is some tension between the metrics, which is the reason why I cannot quite achieve the performance on the individual metrics displayed in Table 10. That being said, the solution to the convex combination generates values closer to the values of the constrained optimal mechanisms than the values of the metrics realized at the Gale-Shapley assignment. Because of the relatively small tension between the goal of increasing student welfare and attaining the distributional goals, the market designer can implement a school assignment that "has it all" — an outcome that realizes high welfare and diversity while at the same time encouraging neighborhood schools.

³³An alternative method for describing the tensions between these objectives would be to trace out the three dimensional surface that describes the possibility frontier across all three metrics. This would have been computationally intensive and difficult to interpret graphically.

Metric	Optimum of Convex Combination of Metrics	Gale-Shapley Mechanism
Welfare	2.48 [2.40, 2.54]	2.62 [2.55, 2.69]
Ethnic Diversity	0.063 [0.052, 0.075]	0.0994 [0.092, 0.113]
Encouraging Neighborhood Schools	0.319 [0.274, 0.342]	0.221 [0.199, 0.241]

Table 12: Maximizing a Convex Combination of the Metrics

	BPS Avg.	Convex Combination of Metrics		
Ethnicity		Average	Min	Max
African American	37.3	45.0 [42.7, 47.0]	33.4 [25.0, 39.0]	60.8 [49.0, 75.3]
White	18.2	10.1 [8.7, 12.1]	2.0 [0, 5.0]	18.5 [16.2, 25.0]
Asian	12.1	8.5 [6.9, 10.3]	4.2 [0, 5.0]	16.7 [12.4, 27.5]
Hispanic	30.0	33.8 [31.6, 35.4]	29.0 [18.7, 30.8]	37.4 [35.3, 41.7]
Native American	0.4	0.7 [0.3, 1.2]	0 [0, 0]	3.7 [1.0, 7.4]
Mixed- Other	1.9	1.8 [1.0, 2.6]	0 [0, 0]	8.3 [3.7, 25]

Table 13: Demographics at Overdemanded Schools In Convex Combination Solution

Table 13 presents statistics about the demographics across the overdemanded schools at the solution to the convex combination, and Table 21 in the appendix displays school-by-school demographic data. Recall that the diversity of the overdemanded schools is improved by (1) reducing the gap in means and (2) reducing the across school gap. Comparing Tables 11 and 13, one sees that at least half of the reduction of the gap in means achieved by the constrained optimal mechanism relative to the Gale-Shapley mechanism is retained by the solution to the convex combination. For the smallest ethnic groups, the gap in means is even smaller in the solution to the convex combination than under the constrained optimal mechanism. The same pattern holds for the maximum and minimal fractions of each school composed of each demographic group. The extremal statistics for most of the demographic groups are very close to those that obtain under the constrained optimal mechanism. The major exception to this is the White student population, which now has one school with just 2% of the student population being white.

5.3 Diversity Possibility Frontiers

The goal of this section is to explore the trade-offs between my ethnic diversity metric and student welfare by creating a possibility frontier that describes the possible realizations of these goals that are compatible with stability and ODIC. To this end, I minimize the following objective function:

$$\alpha R(x) + (1 - \alpha)D(x)$$

where $R(x)$ is the average expected rank of the school assignment x and $D(x)$ is my metric for ethnic diversity. Varying $\alpha \in [0, 1]$ changes the relative importance of the welfare and diversity criteria and allows me to trace out the welfare-diversity possibility frontier. Figure 1 presents the results of my analysis where $D_{Abs}(x)$ is used as the metric for diversity.

Each “dot” represents an optimal, feasible, and ODIC match for some value of α and set of stability constraints.³⁴ Matches with greater levels of student welfare are located to the left of the plot, while matches towards the bottom of the plot have higher levels of diversity. The possibility frontier is the convex hull of the set of stable and ODIC matches.

There are two primary takeaways from my analysis. First, while there are trade-offs between welfare and diversity, almost all of the gains in terms of diversity can be reaped at a welfare cost of less than 0.1 in terms of the average rank. Second, the Gale-Shapley outcome is well inside the frontier and lacking in both welfare and diversity relative to most of the matches I compute. For example, there exist matchings that achieve essentially the maximum possible level of diversity and provide an average welfare improvement of more than 0.15 in terms of average rank relative to the Gale-Shapley mechanism, which is equivalent to moving 177.5 students to a school one rank higher on their preference list.

Figure 2 replicates the diversity analysis under the requirement that the assignment mechanism

³⁴Recall that to define the stability constraints one must decide which overdemanded schools admit students from outside the walk-zone into the walk-zone seats. For any given choice of α there are 24 sets of stability constraints to consider.

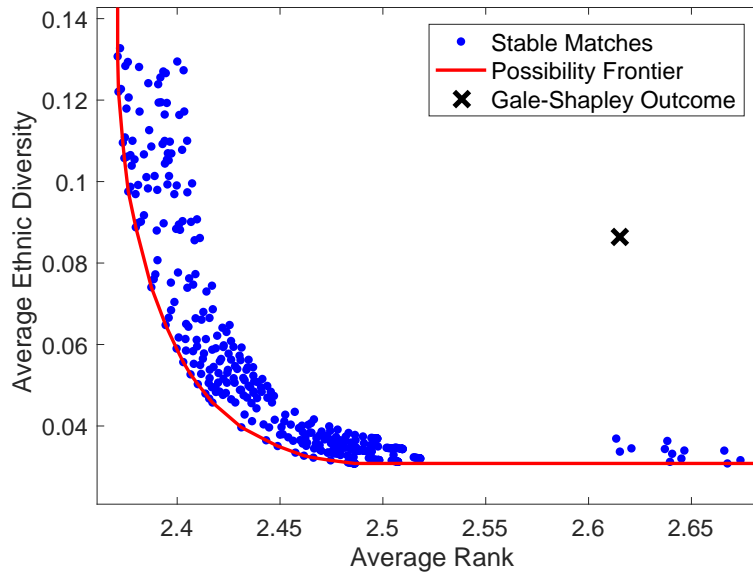


Figure 1: Diversity vs Welfare

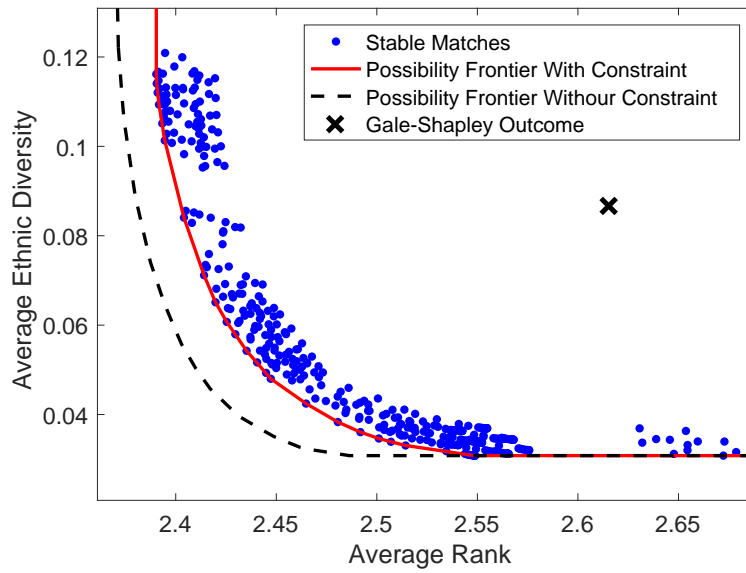


Figure 2: Diversity vs Welfare

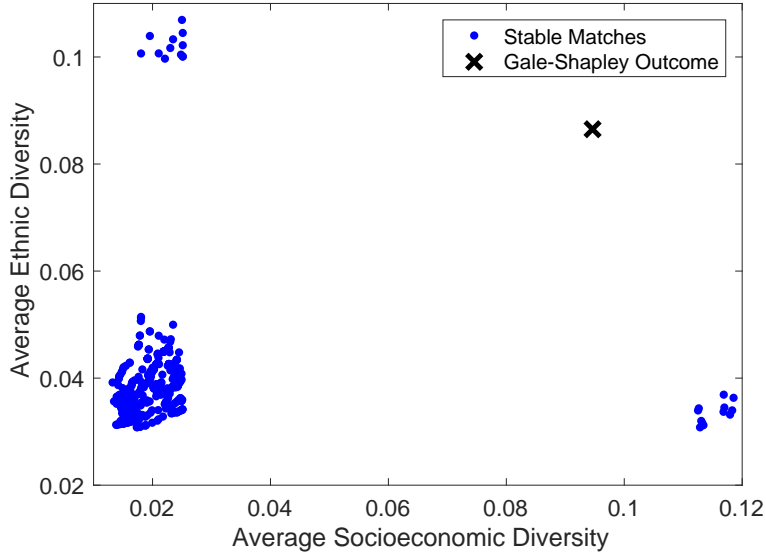


Figure 3: Ethnic vs Socioeconomic Diversity

treat all students from the same zip code symmetrically. I have plotted the stable matches for the more constrained problem as well as the welfare-diversity possibility frontier for the less constrained model (i.e., as shown in Figure 1). As expected, requiring symmetric treatment of students within each zip code reduces my ability to achieve either high welfare or diversity. However, I am still able to improve significantly on the Gale-Shapley outcome on both metrics even given this constraint. Using zip code as a proxy for ethnic background is an effective, if indirect, tool for encouraging school diversity.

I close this section by creating a possibility frontier that describes the the tension, if any, between encouraging ethnic and socioeconomic diversity. The possibility frontier is defined by the minima of $\alpha S(x) + (1 - \alpha)D(x)$ for all possible values of $\alpha \in [0, 1]$. Figure 3 presents the possible combinations of racial and ethnic diversity, where each dot represents an optimal feasible, ODIC match for some choice of α and set of stability constraints. Points close to the left-hand side of Figure 3 denote assignments with a high degree of socioeconomic diversity, and points towards the bottom of the figure are school assignments with a high degree of ethnic diversity.

I have not bothered to draw the convex hull of the possibility frontier since it appears that there is essentially no trade-off between ethnic and socioeconomic diversity. To see this, first note the two clusters of matches, one with a high degree of socioeconomic diversity and a low level of ethnic diversity, and the second cluster with the opposite characteristics. These two clusters are generated by maximizing one form of diversity and completely neglecting the other. The fact that these two clusters are far apart on both dimensions of diversity shows that there is room to separately affect each measure of diversity.

The second point of note is the thick cluster of matches that achieve a high degree of both ethnic and socioeconomic diversity. These are the matches that place a positive weight on each diversity

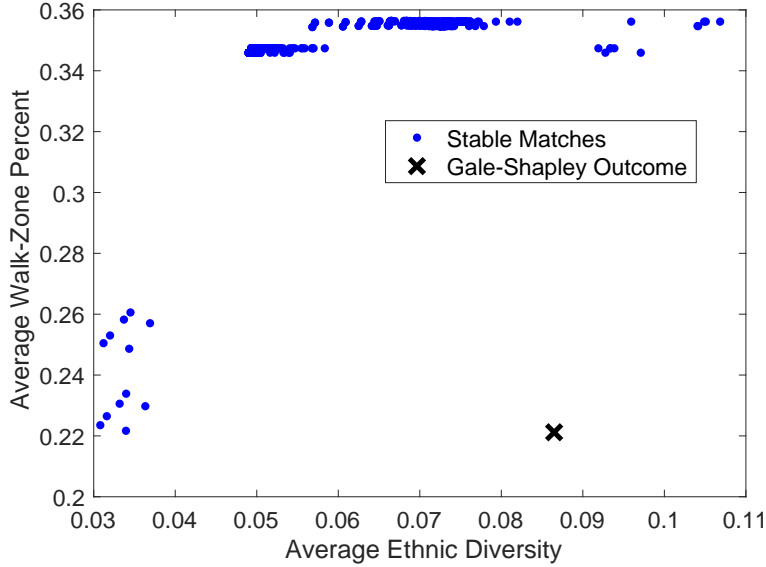


Figure 4: Diversity vs Neighborhood Schools

metric. The tight clustering of these points means that there is essentially no trade-off between these metrics. Since I can achieve diversity along both dimensions easily, it does not matter how much weight I put on each metric when solving the optimization problem. What this means in practice is that it is possible to find a match where the student bodies of the schools resemble the aggregate student population in terms of both ethnic and socioeconomic diversity.

5.4 Encouraging Neighborhood Schools Possibility Frontiers

The origin of the walk-zone priority system was a desire to encourage the enrollment of students in their respective neighborhood schools. I provide three analyses of how effectively one can achieve this goal. First, I assess the trade-offs between encouraging neighborhood schools and diverse schools, and I evaluate whether the segregation of the zip codes necessarily means choosing between these goals. Second, I assess the tension between encouraging neighborhood schools and student welfare. Third, I examine the tensions between these goals without imposing the walk-zone priority structure. I view this final analysis as an *ex ante* perspective in the sense that this would be the natural comparison to use when evaluating potential mechanisms in the absence of any preexisting school choice system.

First I consider the trade-offs between encouraging neighborhood schools and ethnic diversity with the walk-zone priority system in place. Figure 4 reports the result of solving my problem for various weightings of the diversity and walk-zone metrics, $D(x)$ and $N(x)$ respectively. As in Figure 1, Figure 4 uses the absolute deviation metric to describe the diversity of the schools. Matches at the top of the plot encourage stronger neighborhood schools, while points towards the left of the plot encourage diverse student bodies in each school. There are two features to note. The outlying clusters of solutions in the lower-left and upper-right of the figure, one denoting solutions

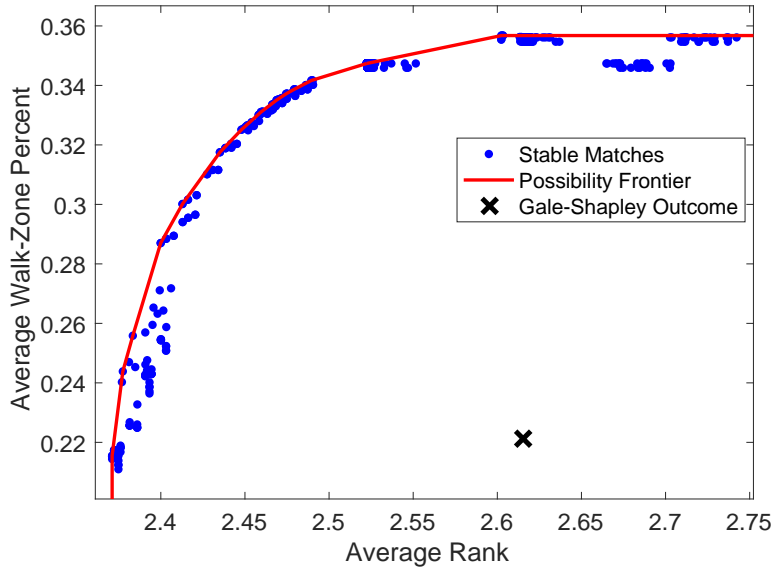


Figure 5: Welfare vs Neighborhood Schools

with a high degree of diversity and few students drawn from the walk-zones and the other having the reverse qualities, are the result of optimization problems that do not place any weight on one of the objectives. This proves that it is possible to significantly influence my diversity and neighborhood school metrics given the stability and ODIC constraints. The second feature is that there exist solutions near the top-left corner of the plot, which represent school assignments that are both diverse and encourage neighborhood schools. The primary takeaway is that it is possible to choose a school assignment that both encourages diversity and creates strong neighborhood schools. Note that the status quo outcome generated by the Gale-Shapley mechanism does not perform well either in terms of diversity or encouraging neighborhood schools even with the walk-zone priorities that BPS applies.

Now I examine the tension between encouraging neighborhood schools and student welfare. The results of my analysis are displayed in Figure 5. Points towards the left-hand side of the plot denote assignments with a high level of student welfare, while points towards the top of the figure represent school assignments that encourage strong neighborhood schools. The outcome generated by the Gale-Shapley mechanism is far from the frontier in terms of both student welfare and encouraging neighborhood schools.

My third set of results provides the exact same analysis of the tensions between encouraging neighborhood schools, the diversity of the schools, and student welfare without imposing the walk-zone priority constraints. From a purely mathematical perspective, it is obvious that eliminating the walk-zone priority relaxes the constraint set and allows the solution of my optimization problems to better achieve the goals I set. The fact that eliminating the walk-zone priority might yield stronger neighborhood schools is, however, probably counterintuitive to a layperson. This confusion comes of trying to achieve an objective (strong neighborhood schools) through the application of a

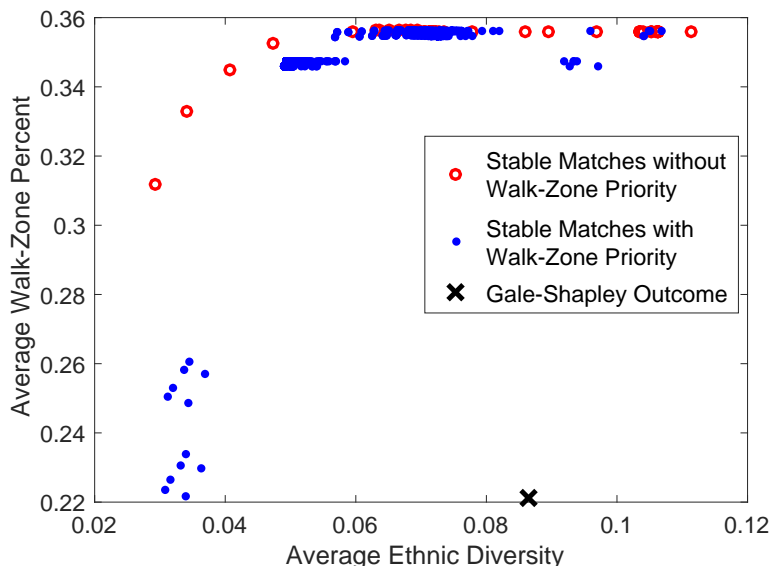


Figure 6: Diversity vs Neighborhood Schools without Walk-Zone Priorities

constraint (walk-zone priority).

Figure 6 presents my results on the tension between encouraging neighborhood schools and school diversity without imposing walk-zone priority constraints. The figure uses the same format as Figure 4. I have plotted the same set of stable matches displayed in Figure 4 along with the stable matches I compute without imposing the walk-zone priority constraints. Note that if I try to choose schools with the maximal level of diversity, the walk-zone priority constraint prevents me from enrolling students in their neighborhood school. To see this, note the circle at the left side of the image, which represents a school assignment with the highest possible level of ethnic diversity that also assigns over 31% of the students to their neighborhood school. If I impose the walk-zone priority constraints and consider only the blue dots, then I can assign at most 25% of the students to their neighborhood school if I insist on also achieving the highest possible levels of ethnic diversity.

Figure 7 displays the neighborhood school - welfare possibility frontier when I do not impose the walk-zone constraints. I have also plotted the possibility frontier under the walk-zone priority constraints from Figure 5. The story is largely similar to that told by Figure 6. If the primary goal is to encourage neighborhood schools, then the walk-zone priority constraints do not interfere with my ability to maximize welfare as a secondary goal. One can deduce this from Figure 7 by noting that the possibility frontiers overlap in the region of the plot where a high percentage of each school's student body is drawn from its walk-zone. If the primary goal is to maximize welfare, however, then the walk-zone priority constraints do restrict my ability to also encourage neighborhood schools. These welfare-optimal matches are towards the left side of Figure 7, which is where there is a significant gap between the possibility frontiers.

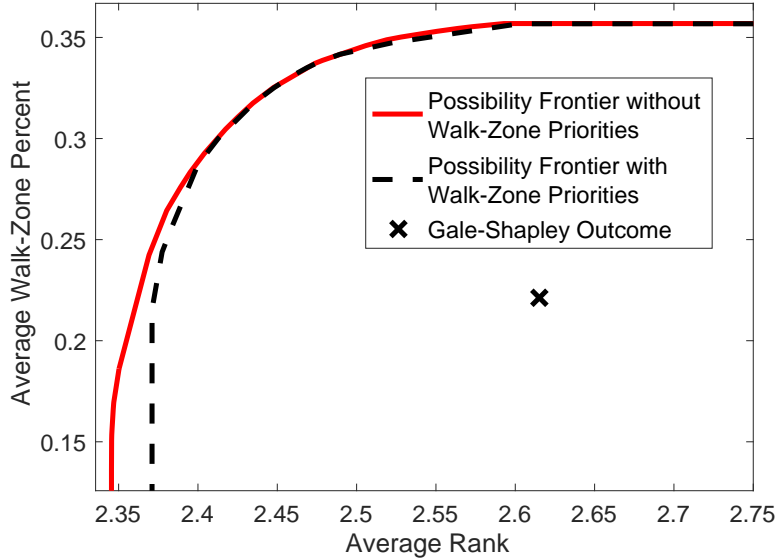


Figure 7: Welfare vs Neighborhood Schools without Walk-Zone Priorities

6 Relation to the Real-World

The linchpin of my analytical framework is that I assume there is a continuum of students of each type matching with a continuum of seats at each college. This allows me to write the ODIC constraints in a form that avoids the curse of dimensionality. However, unless I can establish a link between the continuum model and the finite real-world in which school assignment takes place, it is unclear what the results generated in the continuum framework actually teach us. Section 6.1 provides a theoretical justification for working with the limit model by establishing a link between the game-theoretic properties of the limit model and the analogous school-choice game played by a large, but finite, number of agents. In Section 6.2 I provide a technique for converting a solution to one of my constrained optimization problems into a real-world mechanism that is feasible, stable, and incentive compatible.

6.1 Game-Theoretic Links

In this section I explore the theoretical relationship between the finite and limit models when the optimization problems are thought of as mechanisms. Throughout this section I consider a sequence of N student economies. The goal of each of my approaches is to justify the continuum model as the limit as $N \rightarrow \infty$. I assume the set of colleges is fixed across this sequence of economies. 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 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 . Throughout I assume that the parent distribution from which student

types are drawn is common knowledge, but students are only aware of their own type when they declare their preferences to the mechanism.³⁵

In either the finite or the continuum setting, the mechanism takes the distribution of preference declarations by the students as an input, solves a constrained optimization problem, and provides the solution to the optimization problem in the form of a school assignment. By definition, the resulting school assignment is feasible and stable under the declared preferences. The more interesting question is whether the mechanism is incentive compatible in a finite setting since a change in one student's declaration can influence π^S . In this section I argue that in school choice settings with many students, the mechanism is approximately incentive compatible for all students and, with high probability, strictly incentive compatible for any given student.

There are two assumptions needed to establish the link between the finite and continuum settings. Section 3.3 argued that the ODIC constraints used in my limit model would be appropriate if an agent had a small effect on $x(\circ, \circ; \pi^S)$ if she altered her preference declaration. Assumption 1 formalizes this intuition.

Assumption 1. *In the setting with a continuum of agents, $x(\circ, \circ; \pi^S)$ is continuous in an open neighborhood of π^S .*

The following example shows that Assumption 1 does not hold for all choices of π^S .

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 $q_{c_1} = q_{c_2} = \frac{1}{2}$. The following four cases illustrate that the match fails to be continuous as $\pi^S(s_1) + \pi^S(s_2)$ crosses $\frac{1}{2}$.*

- *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 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.*

Example 1 raises the question of whether Assumption 1 holds for “generic” choices of π^S and $(q_c)_{c \in \mathcal{C}}$. Let $S((q_c)_{c \in \mathcal{C}}, \pi^S)$ denote the set of stable matches given college capacities $(q_c)_{c \in \mathcal{C}}$ and a student-type distribution π^S . A property of $S((q_c)_{c \in \mathcal{C}}, \pi^S)$ is *topologically generic* if it holds over a dense, open set of π^S and $(q_c)_{c \in \mathcal{C}}$.³⁶ Theorem 1 implies that the stable set is continuous for a topologically generic set of $(q_c)_{c \in \mathcal{C}}, \pi^S \gg \mathbf{0}$. Theorem 1 implies that one can choose a stable match $x(\circ, \circ; \pi^S)$ that satisfies Assumption 1 for generic choices of π^S and $(q_c)_{c \in \mathcal{C}}$.

³⁵We could replicate our results in a complete-information setting, but we would need to include caveats precluding nongeneric cases where an agent is pivotal in the iterative menus-and-quotas scheme discussed in Section 6.2.

³⁶I employ the usual topology over Euclidean spaces.

Theorem 1. For topologically 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 $(q_c)_{c \in \mathcal{C}}, \pi^S$.

For the bulk of my analysis I have remained agnostic regarding the cardinal utilities underlying the preference rankings of the students. I now assume that each student has an underlying cardinal utility for each school, and I represent the cardinal utilities of a particular student by a vector $\mathbf{u} \in [0, 1]^{|\mathcal{C}|}$. I use the notation \succeq_u to denote the ordinal ranking consistent with \mathbf{u} .³⁷ I require that the distribution of \mathbf{u} generates a joint distribution of (v, \succeq_u) denoted $\pi^{\mathbf{u}}$ that is consistent with π^S in the sense that for any preference ordering \succeq over the colleges and verifiable trait $v \in \mathcal{V}$ I have:

$$\pi^S((v, \succeq)) = \pi^{\mathbf{u}}(\{(v, u) : \succeq_u = \succeq\})$$

Theorem 2 uses Assumption 1 to argue that in a finite market with many students, the ODIC conditions will only be violated by a small amount if the ODIC conditions hold exactly in the limit game. I am also interested in finding conditions under which the violations of the IC conditions are both small in magnitude and limited to a small set of the students. Intuitively, incentive compatibility will be violated following a small change in x only if some type of student s is almost indifferent between $x(c, s; \pi^S)$ and $x(c, s'; \pi^S)$, $s' \neq s$. I use the cardinal utilities that underly the ordinal preferences to make the notion of “almost indifferent” concrete. Assumption 2 insures that only a small measure of students have a near indifference in their preferences.

Assumption 2. The support of \mathbf{u} includes an open subset of $[0, 1]^{|\mathcal{C}|}$

The mechanism generates a game of incomplete information between the students where the action is a declaration of a preference ranking to the mechanism.³⁸ The von Neumann-Morgenstern utility of an agent with cardinal utility vector \mathbf{u} declaring type s is:

$$\sum_{c \in \mathcal{C}} \mathbb{E} \left[\frac{x(c, s; \pi_E^{S,N})}{\pi_E^{S,N}(s)} \right] u_c$$

where the expectation is taken over the realization of $\pi_E^{S,N}$. My goal is to prove that if truth-telling is an exactly optimal strategy in the mechanism with a continuum of agents, then in the mechanism with a finite number of agents truth-telling is an approximately optimal strategy for all of the agents and exactly optimal for all but a small measure of them. By “approximately optimal” I mean truthfulness is an (ε, ρ) -Bayesian Nash equilibrium.

Definition 2. Truthful declaration of the students’ preference rankings is an (ε, ρ) -Bayesian Nash

³⁷There are many realizations of \mathbf{u} that generate the same ordinal ranking of the school programs.

³⁸If one set out to define a revelation mechanism, the more natural assumption would be that students declare their cardinal utility vector to the mechanism. We assume that preference rankings are submitted for consistency with our existing framework. An additional benefit to declaring a preference ranking is that less information needs to be elicited from the students, making the task easier for the participants.

Equilibrium for $\varepsilon, \rho > 0$ if for all \mathbf{u} and $s = (v, \succeq_u)$ I have:

$$\text{For all } s' = (v, \succeq_{s'}), \sum_{c \in \mathcal{C}} \mathbb{E} \left[\frac{x(c, s(u); \pi_E^{S,N})}{\pi_E^{S,N}(s)} \right] u_c + \varepsilon > \sum_{c \in \mathcal{C}} \mathbb{E} \left[\frac{x(c, s'; \pi_E^{S,N})}{\pi_E^{S,N}(s')} \right] u_c \quad (6.1)$$

and for a measure $1 - \rho$ of students Equation 6.1 can be satisfied with $\varepsilon = 0$.

I can now state my result on the relationship between the limit and finite mechanisms.

Theorem 2. *Under Assumption 1, for any $\varepsilon > 0$ there exists N^* such that if $N > N^*$ truthful declaration is an $(\varepsilon, 1)$ -Bayesian Nash Equilibrium. If Assumption 2 also holds, then for any $\varepsilon, \rho > 0$ there exists N^* such that if $N > N^*$ truthful declaration is an (ε, ρ) -Bayesian Nash Equilibrium.*

There are three intuitions underlying Theorem 2. First, $\pi_E^{S,N}$ will be close to π^S with high probability for large N . Second, if $\pi_E^{S,N}$ is close to π^S , then Assumption 1 implies that $x(c, \circ; \pi_E^{S,N})$ is close to $x(c, \circ; \pi^S)$. If truthfulness is incentive compatible under $x(c, \circ; \pi^S)$, then there can be at most a small benefit to a nontruthful declaration under $x(c, \circ; \pi_E^{S,N})$. Finally, an agent could only benefit from a deviation under $x(c, \circ; \pi_E^{S,N})$ if that agent was nearly indifferent between the declarations under $x(c, \circ; \pi^S)$. Assumption 2 implies that the probability of such an indifference vanishes as N grows (i.e., as $\varepsilon \rightarrow 0$).

Whether the match can be implemented is the final issue I address. The output of the constrained optimization problem is a probability that each student of type s is assigned to each college c , and I refer to this as a *stochastic assignment*. To formalize a stochastic assignment in the N -student mechanism, 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 is assigned to her most preferred underdemanded college (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 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

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

where each \mathbf{X}_i satisfies the feasibility and stability requirements.

My proof uses Theorem 1 of Budish et al. [14], 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 (and not of the particular form studied in Budish et al. [14]), I use Lemma 1 to rewrite the constraints in a way that fits within the framework of Budish et al. [14].

Proposition 3. *Any assignment satisfying equations 3.5 through 3.8 can be implemented.*

6.2 Menus-and-Quotas Mechanisms

I now provide a mechanism that is designed to implement a solution of one of the constrained optimization problems solved above. The mechanism has two parts. First, I restrict the schools that students can rank to a *menu* that is based on their verifiable traits. Second, I reserve a *quota* of seats for students with each possible verifiable trait realization at each of the schools. The combination of these two components results in a mechanism that restricts where students can enroll using the menus and uses the quotas to insure that the number of students of each type enrolled in each program approximately matches the solution to the optimization problem.³⁹

A menu for the group of students with verifiable trait $v \in \mathcal{V}$, denoted $\mathcal{M}(v)$, includes all of the underdemanded schools as well as any overdemanded schools from the following set:

$$\cup_{\succeq} \{c : x(c, s) > 0 \text{ where } s = (\succeq, v)\} \quad (6.2)$$

Equation 6.2 says that $\mathcal{M}(v)$ includes all of the overdemanded schools to which students with trait v are assigned with positive probability by the solution to the optimization problem.

As a practical matter, it may be useful to use a “hidden menu” structure. In such a mechanism, the students can report any preference ranking of any of the colleges. Once this ranking is submitted, the mechanism removes from the preference ranking any schools that are not on that student’s menu. The disadvantage of this structure is that students may waste effort ranking schools that are not on their menus. The benefit is that the school district can collect information on the student’s true preferences, a feature we return to when discussing the adaptability of the mechanism in Section 6.3.

The second component of the mechanisms is the quotas, which are implemented through a priority system. In the limit game, for each verifiable trait $v \in \mathcal{V}$, a measure of seats equal to:

$$Q^\infty(v, c) = \sum_{\succeq} x(c, (\succeq, v)) \quad (6.3)$$

is reserved for students with trait v at school c . I have defined the menus such that the students

³⁹My analysis of the quotas-without-menus mechanism in Table 16 demonstrates that the menus are a necessary addition to the quotas if one wishes to generate an outcome that is close to the solution of the constrained optimization problem.

can only rank schools to which they are assigned under the optimization solution that I want to implement. In the continuum setting the students with each verifiable trait exactly fill the quotas their types have been allocated. As a result, with a continuum of agents, the menus-and-quotas system exactly replicates the solution to the optimization problem.

In the game with a finite number of students, the quota reserves a number of seats at college c for students with trait v equal to:

$$Q(v, c) = N \sum_{\succeq} x(c, (\succeq, v)) \quad (6.4)$$

In a finite setting, one needs to round the number of seats defined by Equation 6.4 to the nearest integer since $x(c, s)$ generates a fractional assignment. This rounding will result in the total number of seats in the quotas to be above the capacity at some schools and below the capacity at others.⁴⁰ When the number of seats assigned via quotas exceeds a school’s capacity, I randomly remove seats from the quotas at that school. When a school has fewer seats assigned via quotas than the capacity, I randomly add seats to the quotas at that school. In the end, every seat at every school is part of a quota for some $v \in \mathcal{V}$.

The quota for each type defines a priority structure for the seats at the school. Students that are insured a seat at c have the highest priority for any seat at that school. Since being insured a seat at a school is a component of v , I assume that students that are insured seats are assigned seats from $Q(v, c)$ before occupying any other seats. Students with verifiable trait v have the second highest priority for seats in $Q(v, c)$, while all other students fall into a third, lowest priority class.

The menus-and-quotas mechanism can be viewed as a matching system with slot-specific priorities (Kominers and Sönmez [45]). By definition, providing types with verifiable trait v a quota $Q(v, c)$ is equivalent to giving these agents a high priority at these seats.⁴¹ The menu can also be thought of in terms of priorities. Consider a student type $s = (\succeq, v)$ and suppose college c is not on her menu (i.e., $c \notin \mathcal{M}(v)$). Removing c from $\mathcal{M}(v)$ is equivalent to making students of type s unacceptable in the priority ranking of school c .

Since the menus-and-quotas mechanism is, in effect, trying to fill a series of lower-bound constraints on the fraction of seats at each school assigned to each type of student, it is worth taking a moment to consider the differences between the menus-and-quotas mechanism and previous mechanisms that used quotas and/or reserves. Previous works have studied how one can design a matching mechanism that tries to satisfy both lower and upper bounds on the fraction of each college that is occupied by each type of student (e.g., Erdil and Kumano [25], Halafir et al. [32], Ehlers et al. [22], Kamada and Kojima [38], Fragiadakis and Troyan [29]).⁴² In order to satisfy

⁴⁰The gap between the total number of seats in the quotas and the capacity of the schools ranged from -2 to $+3$, which represents less than 2% of the seats at each school.

⁴¹One can consider an even stronger notion of a quota where only agents with a particular verifiable trait are acceptable for the seats reserved for that trait realization. Such a system would exacerbate the problem of unfilled seats that we discuss below.

⁴²Some papers differentiate between hard quotas that must be satisfied for a feasible match and soft quotas that merely serve as goals. In this vernacular, our quotas take the form of soft goals.

Metric	Menus-and-Quotas	Convex Combination of Metrics, Limit Model	Gale-Shapley Mechanism, Limit Model
Welfare	2.54 [2.47, 2.60]	2.48 [2.40, 2.54]	2.62 [2.55, 2.69]
Ethnic Diversity	0.082 [0.068, 0.094]	0.063 [0.052, 0.075]	0.0994 [0.092, 0.113]
Encouraging Neighborhood Schools	0.427 [0.403, 0.453]	0.319 [0.274, 0.342]	0.221 [0.199, 0.241]

Table 14: Maximizing a Convex Combination of the Metrics

lower bound constraints, many of these mechanisms consider artificial capacities that are weakly below the true capacity of the college. The idea is that the artificially low capacity will prevent students from enrolling in a popular school, and the newly rejected students will hopefully help satisfy lower bound constraints by enrolling elsewhere. In the menus-and-quotas mechanism, our menus control where students can enroll and are designed to guide the students towards enrolling in colleges where their enrollment helps satisfy a lower bound constraint without manipulating the capacities at the schools. It is not clear to me how one could have designed effective menus without solving some underlying optimization problem, which none of the prior works in the quotas and reserves literature has done to my knowledge.

The timing of the menus-and-quotas mechanism proceeds as follows:

1. Students are provided menus by the mechanism.
2. Each student submits a rank ordered list of schools from his or her menu to the mechanism.
3. The mechanism assigns each student a random number to break ties in priority.
4. The Gale-Shapley algorithm is executed using the priority system defined by the the quotas (i.e., priorities).

The menus-and-quotas mechanism will satisfy the capacity constraints of the schools, as students can only be assigned to seats at a school that actually exist as in the standard Gale-Shapley mechanism. The mechanism is also incentive compatible — the restrictions imposed by the menu lead to the same outcome as if the student were deemed unacceptable by every school not on his or her menu and I had run the usual Gale-Shapley algorithm. The remaining concern is the stability of the mechanism, which I address after describing the performance of the menus-and-quotas mechanism on my three metrics.

The first column of Table 14 presents the realizations of our three metrics for the menus-and-quotas implementation of the solution to the convex combination (Section 5.2, Tables 12 and 13). I chose the menus and quotas based on the status quo distribution of types and did not adjust them across bootstrap runs. The second column of results refers to the solution of the optimization problem using a convex combination of my metrics as the objective with a continuum of students.

The solution to the convex combination is recomputed for each bootstrap sample of types, which gives a sense for the improvement yielded by adjusting the menus-and-quotas scheme based on the realized π^S in the bootstrap sample. The final column includes the outcomes that would be realized if one ran the Gale-Shapley mechanism with a continuum of students on each bootstrap sample.

One should not expect to do as well on the welfare or ethnic diversity metrics using a menus and quotas system in the finite setting as one might naively predict from the limit model. The main reason that the performance drops relative to that realized in the limit model is that I am not adapting the menus and quotas for the distribution of types realized in each bootstrap sample. This means that in a particular bootstrap sample, there may not be enough students with a particular verifiable trait $v \in \mathcal{V}$ to fill the respective quota or so many students with the trait that the quota cannot accommodate all of them. In the later case, the excess students will compete for seats from other quotas that do not have enough students with the corresponding verifiable trait v to fill them. Both of these effects will distort the school assignment in the menus-and-quotas system away from the exact solution of the limit model. A second reason the menus-and-quotas outcome will differ from the limit model solutions is that one needs to round the fractional matchings off, which will cause larger discrepancies at smaller schools.⁴³

Based on the preceding discussion, one might assume that the menus-and-quotas system must also do worse than the limit model solution on my metric for encouraging neighborhood schools, but in fact the opposite occurs. The menus-and-quotas system does significantly better than the limit model on the walk-zone metric because the menus the students are provided tend to prevent students from ranking a school unless it is in their walk-zone. This results in a disproportionate number of walk-zone students being assigned to seats in their local overdemanded school. One could interpret the menus-and-quotas mechanism's poorer performance on the welfare and racial diversity metrics as the price one must pay for the improvement in the percentage of the school's student body drawn from the walk-zone.⁴⁴

By design, none of the schools can have an enrollment over capacity in a menus-and-quotas scheme. However, it is possible that a school will have empty seats that would result in violations of stability. To get a sense for the scale of this problem, Table 15 provides the mean number of empty seats in each of the 10 overdemanded school programs, the probability that the school is filled to capacity, and the probability that the school is filled to at least 90% of capacity. With the exception of Snowden International, all of the schools are filled to capacity most of the time. In addition, with the exception of the smallest schools (Snowden International and Lyon), all of the

⁴³If I run the menus-and-quotas mechanism in a finite setting with the status-quo distribution of student types, then the outcome is only affected by the rounding of the fractional matchings. When I do this, I achieve an average rank of 2.49, a diversity metric of 0.0704, and an average fraction 0.434 of the student body of each school is drawn from the walk-zone. Since these values are much closer to the outcome of the solution of the convex combination in the limit model than they are to the average outcome of applying the menus-and-quotas mechanism to the bootstrap samples, I conclude that most of the failure of the finite mechanism to replicate the outcomes of the continuum solution is the choice to not adapt the mechanism to variation in π^S .

⁴⁴Technically there probably exists solutions that do slightly better than the menus-and-quotas mechanism's outcome on all three metrics by the simple fact that the menus and quotas system is not the exact solution to an optimization problem.

School Name	Number of Empty Seats	Probability 100% Full	Probability 90% Full
Brighton High	2.47 [15, 0]	0.737	0.967
Excel High	1.94 [11, 0]	0.695	0.944
Lyon	0.91 [5, 0]	0.675	0.768
Snowden International, All Programs	7.28 [19, 0]	0.109	0.766
Another Course College	1.38 [7, 0]	0.692	0.871
Urban Science Academy	4.74 [20, 0]	0.532	0.898
Tech Boston Academy	2.56 [15, 0]	0.633	0.913

Table 15: Empty Seats Under the Menus-And-Quotas Mechanism

schools are at 90% of capacity or more at least 85% of the time.

Any time there is an empty seat at an overdemanded school, there will be a blocking pair. The fact that the average number of empty seats in a school is low might not be very useful if parents realize that a rare event has occurred and a very popular school like Snowden International is 10% of more empty. I explore two modifications of the menus-and-quotas mechanism to avoid these high rates of under-capacity schools. The first modification is to allow the students to rank any school in the preferences submitted to the mechanism, but a student that ranks a school that is not in her menu has the lowest priority at that school (but is acceptable to that school). I refer to this as a *Quotas-without-Menus* mechanism since the menus do not define which schools can be ranked. Bootstrap simulations show that this mechanism fills all of the overdemanded schools with the exception of the Urban Science Academy, which is full 87.5% of the time and at least 90% full 99% of the time.

A second modification to the menus-and-quotas system is to allow more students to rank a school when that school falls below a capacity threshold. Since the mechanism iterates the static menus-and-quotas mechanism, I refer to this as the *Iterative Menus-and-Quotas* mechanism. The iterative mechanism proceeds as follows:

1. Choose a lower bound on the percentage of the seats at each overdemanded school that must be assigned.
2. Build the menus $\mathcal{M}(v)$ as per Equation 6.2.
3. Run the menus-and-quotas mechanism.
4. Identify any schools that do not satisfy the lower bound on the fraction of its seats that are

Metric	Iterative Menus-and-Quotas	Menus w/o Quotas	Static Menus-and-Quotas
Welfare	2.52 [2.45, 2.59]	2.55 [2.48, 2.63]	2.54 [2.47, 2.60]
Ethnic Diversity	0.084 [0.071, 0.098]	0.097 [0.085, 0.108]	0.082 [0.068, 0.094]
Encouraging Neighborhood Schools	0.403 [0.370, 0.436]	0.349 [0.326, 0.372]	0.427 [0.403, 0.453]

Table 16: Maximizing a Convex Combination of the Metrics

assigned.

- If there is such a school, add this school to the menu of all of the students and go to step 3.
- If there is no such school, finalize the assignment.

The iterative menus-and-quotas mechanism will eventually terminate with all of the schools satisfying the constraint on the percentage of assigned seats. To see this, simply note that the iterative process eventually allows all students to rank all of the schools, at which point all of the overdemanded schools (except possibly Urban Science Academy) must be full. The more interesting question is whether the mechanism is incentive compatible, since one might think a student could manipulate her ranking to cause another school to be added to her menu.

In the finite setting, a student has an incentive to declare her preferences nontruthfully only if she is pivotal for determining whether a school meets the bound on the percentage of assigned seats and can thus alter the future menu of schools provided to each participant.⁴⁵ However, the probability of this event vanishes as the mechanism grows, and so the mechanism will be approximately incentive compatible for most agents. Across more than 500 bootstrap samples, there was never an event where a single agent was pivotal in this way.⁴⁶

Table 16 shows the distribution of outcomes achieved by the Iterative menus-and-quotas mechanism, the quotas-without-menus mechanism, and the static menus-and-quotas mechanism. I require that the schools be at least 95% filled under the iterative design. There are two primary messages from Table 16. First, menus are important for encouraging racial diversity within the schools. Since students probably care little for their aggregate effect on the demographics of the BPS schools, it

⁴⁵In the continuum version of the iterative menus-and-quotas mechanism, a single student's deviation cannot alter the menu he is offered at future stages, which means the mechanism is incentive compatible in the limit model. It is easy to prove using techniques similar to the proof of Theorem 2 that truthfulness is an (ϵ, ρ) -Bayesian Nash equilibrium.

⁴⁶If one insists on an exactly incentive compatible mechanism, then for each agent i one could run the iterative menus-and-quotas mechanism with agent i excluded. One can deduce the menu that agents that share agent i 's verifiable type received in the counterfactual mechanism without agent i , and denote this menu \mathcal{M}_i . We can then run the static menus-and-quotas mechanism where agent i receives menu \mathcal{M}_i . Since agent i can no longer influence the menu she is offered, the mechanism is incentive compatible.

Ethnicity	BPS Avg.	Convex Combination of Metrics		
		Average	Min	Max
African American	37.3	47.0 [44.4, 50.1]	25.1 [8.3, 37.5]	64.0 [54.9, 73.5]
White	18.2	9.2 [7.3, 11.0]	1.1 [0, 3.0]	23.2 [15.4, 33.3]
Asian	12.1	7.7 [6.0, 9.5]	1.1 [0, 3.7]	20.0 [11.9, 33.3]
Hispanic	30.0	33.1 [30.1, 36.1]	19.9 [8.3, 26.8]	48.2 [39.2, 61.5]
Native American	0.4	0.7 [0.2, 1.2]	0 [0, 0]	3.4 [1.0, 7.4]
Mixed- Other	1.9	1.8 [1.0, 2.6]	0 [0, 0]	11.2 [3.7, 25.0]

Table 17: Demographics at Overdemanded Schools Under Iterative Menus-and-Quotas

is not surprising that the schools become more segregated when I do not use menus to restrict the preference rankings that the students can submit. However, even without menus I am able to improve student welfare and encourage neighborhood schools at the same time by crafting the priority system around the solution to the constrained optimization problem.⁴⁷

Second, and relatedly, using an iterative mechanism that tunes the menus to the capacities of the overdemanded schools does a good job of encouraging diverse student bodies while also eliminating the spare capacity at the overdemanded programs. The outcome of the iterative menus-and-quotas system is close to the outcome of static menus-and-quotas system across all three metrics, but without the large number of unfilled seats that we observe in the static version. Table 17 provides statistics on the student bodies of the overdemanded schools when the iterative menus-and-quotas mechanism is used. Referring to Table 13, the reduction of the gap in mean achieved under the iterative menus-and-quotas scheme is roughly the same as that achieved by the solution to the convex combination in the limit model. The minimum and maximum statistics are volatile due to changes in the student bodies of the smallest school programs, which suggests that the diversity of these smaller school programs may benefit from adapting the menus-and-quotas scheme to the actual realization of student types in each year. Table 18 presents school-by-school demographic data as well in terms of the number of students from each ethnic group enrolled in each program.

⁴⁷It is difficult to compare the results of Table 16 with the previous solutions to the constrained optimization problem since (1) I am not directly enforcing the walk-zone stability constraints in the menus-and-quotas system and (2) I allow the capacity constraints to be satisfied weakly in the menus-and-quotas system.

School Name	African-American	White	Asian	Hispanic
Brighton High	78.1 [68, 89]	10.9 [5, 17]	13.2 [7, 19]	63.5 [53, 75]
Excel High	43.1 [35, 51]	13.2 [8, 19]	14.9 [9, 20]	28.4 [21, 35]
Lyon	3.8 [1, 7]	1.3 [0, 3]	1.7 [0, 4]	4.2 [1, 7]
Snowden International, Chinese	9.8 [6, 13]	2.6 [1, 5]	2.4 [0, 5]	7.6 [4, 11]
Snowden International, French	11.6 [8, 16]	4.5 [2, 7]	1.5 [0, 4]	9.3 [5, 14]
Snowden International, Japanese	10.6 [7, 14]	4.0 [1, 7]	2.0 [0, 4]	9.3 [5, 13]
Snowden International, Spanish	10.7 [7, 14]	3.3 [1, 6]	0.9 [0, 2]	10.6 [7, 14]
Snowden International, All Languages	42.7 [35, 51]	14.3 [9, 20]	6.7 [3, 11]	36.8 [28, 46]
Another Course College	18.2 [13, 23]	6 [2, 11]	3.9 [1, 7]	13.6 [8, 18]
Urban Science Academy	72.5 [62, 83]	16.0 [10, 22]	5.8 [2, 10]	54.1 [44, 64]
Tech Boston Academy	64.9 [55, 75]	1.7 [0, 4]	6.8 [1, 12]	26.9 [19, 35]

Table 18: Demographics at Overdemanded Schools Under Iterative Menus-and-Quotas Mechanism in Number of Students

6.3 Practical Issues

In this section I address three potential difficulties with implementing the mechanism of Section 6.2. The first issue involves whether the mechanism can adapt to changes to the school capacities, the introduction of new schools, or shocks to the aggregate distribution of preference declarations. The second issue involves the transparency of the mechanism to parents. The final issue is the accessibility of the schools.

The first difficulty in implementing the menus-and-quotas mechanism is deciding how to adapt the mechanism to changes to the school capacities, the introduction of new schools, or shocks to the aggregate distribution of preferences. Some changes are planned by the school system, such as the founding of a new school or the introduction of a new program within an existing school.⁴⁸ Unplanned changes include aggregate shocks to the distribution of student preferences caused by an unexpectedly poor outcome of a school evaluation or if a school expands its offerings of college-level courses. If a new school is introduced, the designer could allow all of the students to rank the new school program. Once information about the new distribution of student preferences has been collected, the optimization problem can be re-solved and the menus-and-quotas mechanism can be rebuilt. If there is a shock to the distribution of preferences, then the iterative menu-and-quota system will yield an outcome that is distorted relative to the optimization solution on which the mechanism is founded. This would suggest that it may be appropriate to re-solve the optimization problem given the new distribution of student preferences and build a new menus-and-quotas system around the new solution.

The second issue to discuss is the transparency of the mechanism. Two notions of transparency should be considered: *procedural transparency* and *validation transparency*. Procedural transparency refers to the degree to which parents understand how the mechanism operates. Validation transparency refers to the ease with which parents can verify that their children would not have benefited by providing a false preference ranking. Both forms of transparency are oriented towards reassuring families that truthful declarations of their preferences are optimal. In contexts where the Gale-Shapley mechanism is used, a great deal of effort is spent trying to explain the procedures used by the mechanism to generate the assignment. In some contexts this is relatively successful. For example, Rees-Jones [53] uses survey responses from participants in the medical residency match program to argue that only 5.4% of medical students submit nontruthful preference ranking in an effort to manipulate the outcome.

The menus-and-quotas mechanism has a high degree of validation transparency. The challenge facing parents that wish to confirm that their child was assigned to the best school consistent with stability is essentially the same as under the Gale-Shapley mechanism. To do so, the parents need to know the priorities of the students assigned to each school and the priority of their own child. The procedural transparency of a menus-and-quotas system is limited by the fact that the origins of the menus and quotas is an underlying optimization problem. Since for most people, perhaps

⁴⁸A good example of this is the Margarita Muniz Academy, which started accepting students for the first time in the Fall of 2012.

even most economists, the computational method used to solve a complex constrained optimization problem is a black box, explaining the source of the assignment on which the menus-and-quotas system is based may not be possible.

The final concern I would like to address is school accessibility, which refers to the ability of parents to enroll their child in the school of their choice. The menus-and-quotas mechanism may make it impossible for some students to enroll in certain schools. There are multiple ways that this could be addressed, all with upsides and downsides. The first solution would be to reserve a block of seats at each school for open-enrollment, allow any student to rank open-enrollment seats, and give each student a priority for these seats determined by random tie-breaking. Simulation studies could be used to test the trade-offs between increasing accessibility by increasing the size of the open-enrollment seats and distorting the assignment away from the underlying solution of the optimization problem. A second solution would be to include a subset of the overdemanded schools in each student’s menu, which is a milder version of the “quotas without menus” mechanism analyzed in Section 6.2, Table 16. A final solution would be to modify the optimization problem underlying the menus-and-quotas system by (1) identifying the student types that rank overdemanded schools and (2) apply *school accessibility restrictions* that require that these types have a positive probability of enrolling at one of the overdemanded schools they have ranked. One benefit of this approach is that the optimization problem would automatically adjust the solution to account for the effect of the school accessibility restrictions.

7 Related Literature

The closest analogs to my work are the contemporaneously developed papers of Ashlagi and Shi [8], and Shi [61].⁴⁹ Both of these papers share my view that matching problems can be effectively tackled using constrained optimization problems featuring a continuum of students matching to a continuum of college seats. The focus of both Ashlagi and Shi [8] and Shi [61] is optimizing a convex combination of average and minimal student welfare subject to busing cost constraints as well as capacity and incentive compatibility constraints. Instead of solving the school assignment problem directly (as I do), Ashlagi and Shi [8] and Shi [61] prove that one can reformulate the optimal stable matching problem without priorities as a random assortment problem in which the mechanism offers students randomized menus of school assignments from which to choose. Importantly, the random assortment problem can be solved in polynomial time if one assumes that student utility has a multinomial logit form.

One of the biggest differences between this project and Ashlagi and Shi [8] and Shi [61] is the objectives I am trying to achieve. I am interested in encoding distributional concerns such as school diversity or encouraging neighborhood colleges into my objective function, while Ashlagi and Shi [8] and Shi [61] are interested in maximizing welfare metrics. Another major difference is that I can

⁴⁹Featherstone [27] 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 [27] does not include incentive compatibility constraints.

incorporate priority structures (e.g., walk-zone priority) into my framework since I solve the school assignment problem directly, which can be practically important if parents view these as inviolable property rights.⁵⁰ The priority structures render my problem nonconvex, however, which means that the computational efficiency of solving my model may fall rapidly as more schools are added. I also provide a link between the continuum and finite markets by proving the stable set is generically continuous, which I believe is a novel result in the matching literature and of independent interest. Finally, I propose a novel menus-and-quotas system for implementing a solution to a constrained optimization problem, and I study the practicality of such a method using bootstrap simulation.

The idea of describing the stable set using constraints is not new and appears in the economics literature as early as Roth et al. [56]. These ideas were extended to many-to-one matching in Baïou and Balinski [12]. Both of these prior works consider a finite set of agents, which means that, while the formulations are useful for proving properties of matching problems, they cannot be used to generate numerically tractable optimization problems. See Section 3.3 for more details on the incentive compatibility conditions for a matching mechanism with a finite number of agents.

My framework is based on a continuum of each of a discrete set of student-types that are matched to a discrete set of colleges with a continuum of seats. An early example of a matching paper featuring a continuum of agents is Azevedo and Leshno [10], which assumes an exogenous distribution of student preferences and priority numbers that do not admit indifferences. The authors 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. Later works used continuum models to provide weak conditions under which large matching markets are approximately stable (e.g., Azevedo and Hatfield [11], Che, Kim and Kojima [16]). There is also a vein of the large market literature that tries to establish the equivalence (or lack thereof) between different mechanisms (e.g., Kojima and Manea [43], Manea [49], and Liu and Pycia [48])

The primary difference between my paper and the extant literature on large matching markets is the goal of the research projects. The goal of my paper is to provide a pragmatic framework for assessing and designing optimal school assignment mechanisms with a particular emphasis on distributional goals of school choice (e.g., diversity) as well as student welfare. The existing literature on large matching economies is focused on proving useful technical properties of large matching markets (e.g., approximate stability).

One of the goals of this paper is to provide a practical method for accommodating diversity goals into stable, incentive compatible matching systems. There have been a number of recent theoretical 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. [14], Halafir et al. [32], Kominers and Sönmez [45]).⁵¹ Echenique and Yenmez [21] axiomatize several notions of diversity that are compatible with stability.

⁵⁰I also need make no assumptions on the students' utility functions, but this is perhaps of lesser benefit given the prominence of multinomial logit models in empirical work across economics.

⁵¹Boston's efforts to provide walk-zone students with priority at their local schools has recently been critiqued by Dur et al. [19] as a de facto quota.

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. The use of random tie-breakers, however, can result in an inefficient outcome. Erdil and Ergin [24] 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 [41] provides an algorithm for identifying situations in which an agent's priority can be altered without affecting the agent's outcome, thus insuring incentive compatibility while allowing for welfare improvements.

There also exists a literature focusing on limit properties of a finite matching market as it grows without working with a limit model. Roth and Peranson [55] points out that the set of stable matches in the National Residency Matching Program is small, which implies approximate strategy-proofness. Immorlica and Mahdian [37] and Kojima and Pathak [44] analyze models of one-to-one and many-to-one matching markets (respectively) using the Gale-Shapley algorithm, and show that the markets become approximately incentive compatible as they grow. Lee [47] 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. [20], which uses a model with multiple copies of each type of agent to study the identification restrictions imposed by stability. Echenique et al. [20] assumes all agents have strict preferences and does not study the limit as the finite model approaches a continuum game. Menzel [50] uses a large market model where agent types can take values in a continuum and preferences are strict. The primary results of Menzel [50] show that the finite model converges to a limit game wherein the equilibrium objects of interest are identified.

While most of the empirical literature on matching is focused on the Gale-Shapley mechanism due to its common usage, there have been a number of papers that study nontruthful matching mechanisms such as the Boston mechanism (e.g., Agrawal and Somaini [6], Calsamiglia et al. [15], He [35], and Kapor et al. [39]). Estimating the preferences of agents participating in a nontruthful mechanism requires modeling and structurally estimating the underlying cardinal utility function as these nontruthful mechanisms typically expose the participants to risk. If I had estimated a model of the cardinal utilities underlying the preference declarations of the students, I would be able to calibrate utility in concrete terms. For example, I might discuss the value of moving a student to a more preferred school in terms of the student's disutility from the distance of his or her commute. The major downside of estimating such a model is that one is required to parameterize a model of the agent's utility function and (often) the distribution of utilities in the population. In my judgment the benefits of such an estimation exercise are outweighed by the costs of imposing parametric assumptions on the agents' preferences.

Since the Gale-Shapley mechanism is incentive compatible, I have taken the student's preference declarations at face value. However, there is field evidence that students may not submit their

preferences to a mechanism truthfully.⁵² As mentioned, Rees-Jones [53] finds that 5% of participants in the U.S. National Residency Matching Program misreport their preferences for strategic reasons. Hassidim et al [33] find that a significant fraction of participants in the Israeli Psychology Master’s Match do not report truthfully.⁵³

Artemov et al. [7] shows that many participants in the Australian college admissions system make payoff-irrelevant mistakes since they have precise information about their priorities at the different colleges. The authors point out that an estimator that assumes the preferences are declared truthfully can be biased by these mistakes, and the authors argue for an identification strategy based on stability. Fack et al. [26] show that nontruthful preference rankings can happen in equilibrium when (1) preferences and priorities are her private information and (2) there is a cost to ranking a school. In the status-quo Gale-Shapley mechanism BPS used, any student could potentially be assigned to any school she ranks if her tie-breaker is sufficiently high. Since her random tie-breaker is not known to the student at the time her preferences are declared, there are no payoff-irrelevant mistakes in the status-quo mechanism.

8 Conclusion

In this paper I introduce a novel analysis tool for school choice problems based on constrained optimization problems 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 stability and incentive compatibility constraints in a computationally tractable form. One major advantage of the constrained optimization approach is that a market designer can explicitly encode his or her goals into an objective function, and the solution to the resulting constrained optimization problem is ensured to achieve the global optimum subject to the stability and incentive compatibility 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 study the benefits of relaxing the constraints by for example, increasing a school’s capacity.

In the context of the BPS high school match, I find that the Gale-Shapley algorithm reaps 78% of the welfare benefit of moving from the worst to the best stable and ODIC match. I also find that the ODIC constraints limit the welfare that can be achieved. Finally, I show that if one is willing to let the mechanism condition on student demographics, one can achieve a welfare gain that is equivalent to moving 291 students to a school one rank higher in their preference lists. I refer to this welfare improvement as the welfare cost of fairness to highlight that achieving the gain requires treating students differently based on welfare-irrelevant traits like the zip code in which the student lives.

⁵²There is also a rich experimental literature studying violations of truthful behavior in incentive compatible mechanisms.

⁵³Hassidim et al. [33] survey the students who made an “obvious” misreport. A theme of the survey results are comments from the participants that suggest they either did not understand the mechanism or believed the misreport would increase their chance of admission at a preferred institution. It is unclear whether such a thing could happen in the context of BPS, an institution that most parents are familiar with and had interacted with previously.

I then maximize student body diversity with respect to ethnicity and socioeconomic status. I show that there are stable matches that are significantly more diverse than the Gale-Shapley outcome, and the cost in terms of welfare is modest. In addition, I show that there is essentially no trade-off between the two forms of diversity - it is easy to find a stable and ODIC match that is diverse in both senses.

For my third analysis, I study the extent to which one can both encourage neighborhood schools while at the same time increase student welfare or school diversity. I again find that one can do significantly better than the Gale-Shapley outcome on both metrics. One can increase the average percentage of students drawn from a school's walk-zone by 50% (relative to the status quo outcome) with no loss of student welfare. One might have assumed diversity and encouraging neighborhood schools would be at odds with each other because many zip codes in Boston are segregated. On the contrary, I find that there is little tension between these goals - one can find stable matches that simultaneously maximize the fraction of the students drawn from the schools' walk-zones and achieve high levels of diversity.

My final task is to relate the optimization problems in the continuum model to the real-world market with a finite number of agents. First, I show that if one thinks of an optimization problem as a mechanism that requires strategic agents their ordinal preference rankings, then since truthfulness is incentive compatible when a continuum of students participates, truthfulness is approximately incentive compatible when a finite number of agents participate. The foundation of this result is a novel theorem on the continuity of the stable set, which is of independent interest.

Next I propose a class of menus-and-quotas mechanisms to implement solutions to my optimization problems. First, I design a menu for each student-type that restricts the schools students of that type can rank based on where students of that type are assigned in the solution to the optimization problem. The menus then guide the students to enroll in schools where they can, for example, increase the diversity of the student body. The quota component gives students with each verifiable trait a high priority for a block of seats at each school, and the size of the quota is derived from a solution to one of my optimization problems. These quotas help control how many students of each type can enroll in each college. Since a static version of this mechanism can leave a large number of seats unassigned at the overdemanded schools, I propose a dynamic version of this mechanism that leaves a small fraction of the seats at these schools unassigned while retaining most of the improvements in terms of student welfare, increasing school diversity, and encouraging neighborhood schools.

My hope is that the techniques used here make three specific contributions to the school choice and market design literatures. First, while the welfare of the market participants is an obviously relevant goal, there are other goals that are distributional in nature (e.g., school diversity) that also merit consideration. My analytical framework provides a way of implementing distributional objectives in a school choice market design problem. Second, I find that one can make significant improvements on the Gale-Shapley outcome with respect to welfare and my distributional goals without sacrificing either stability or incentive compatibility. Third, the application of my frame-

work to the BPS data, while of independent interest, provides a proof of concept for the constrained optimization approach. Although there are some school systems (e.g., New York City) that have so many school programs that my approach may be impractical, the constrained optimization approach is tractable for many moderately large school systems such as BPS.

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] N. Agrawal and P. Somaini (2014) “Demand Analysis Using Strategic Reports: An Application to a School Choice Mechanism,” *mimeo*.
- [7] Artemov, G.; Y-K Che; and Y. He (2017) “Strategic “Mistakes:” Implications for Market Design Research,” *mimeo*.
- [8] Ashlagi, I. and P. Shi (2014) “Optimal Allocation without Money: an Engineering Approach,” *Management Science*, 62 (4), pp. 1078 - 1097.
- [9] Azevedo, E. and J. Hatfield (2012) “Complementarity and Multidimensional Heterogeneity in Matching Markets,” *mimeo*.
- [10] Azevedo, E. and J. Leshno (Forthcoming) “A Supply and Demand Framework for Two-Sided Matching Markets,” *Journal of Political Economy*.
- [11] Azevedo, E. and J. Hatfield (2012) “Complementarity and Multidimensional Heterogeneity in Matching Markets,” *mimeo*.
- [12] Baiou, M. and M. Balinski (2000) “The stable admissions polytope,” *Mathematical Programming*, 87, pp. 427 - 439.

- [13] Bodoh-Creed, A. (2013) “Efficiency and information aggregation in large uniform-price auctions,” *Journal of Economic Theory*, 148 (6), pp. 2436 - 2466.
- [14] 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.
- [15] Calsamiglia, C.; C. Fu; and M. Güell, “Structural Estimation of a Model of School Choices: the Boston Mechanism vs. Its Alternatives,” *mimeo*.
- [16] Che, Y-K; J. Kim; and F. Kojima (2014) “Stable Matching in Large Economies,” *mimeo*.
- [17] Chen, Y. and O. Kesten (Forthcoming) “Chinese College Admissions and School Choice Reforms: A Theoretical Analysis,” *Journal of Political Economy*.
- [18] Curtis, C. (1984) *Linear Algebra*, Springer-Verlag: New York.
- [19] Dur, U.; S. Kominers; P. Pathak; and T. Sönmez (2014) “The Demise of Walk Zones in Boston: Priorities vs. Precedence in School Choice,” *mimeo*.
- [20] 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.
- [21] Echenique, F and M.B. Yenmez (2014) “How to Control Controlled School Choice,” *The American Economic Review*, 105 (8), pp. 2679 - 2694.
- [22] Ehlers, L; I. Hafalir, M.B. Yenmez, and M. Yildirim (2014), “School choice with controlled choice constraints: Hard bounds versus soft bounds,” *Journal of Economic Theory*, 153, 648683.
- [23] Embrechts, P.; C. Klüppelberg; and T. Mikosch (1997) *Modelling Extremal Events for Insurance and Finance*, Springer: New York.
- [24] Erdil, A. and H. Ergin (2008) “What’s the Matter with Tie-Breaking? Improving Efficiency in School Choice,” *American Economic Review*, 98, pp. 669-689.
- [25] Erdil, A. and T. Kumano (2013), “Prioritizing diversity in school choice,” *mimeo*.
- [26] Fack, G.; J. Grenet; and Y. He “Beyond Truth-Telling: Preference Estimation with Centralized School Choice,” *mimeo*.
- [27] Featherstone, C. (2006) “What’s the Matter with Tie-Breaking? Improving Efficiency in School Choice,” *American Economic Review*, 98, pp. 669-689.
- [28] Featherstone, C. and M. Niederle (2013) “Improving on Strategy-proof School Choice Mechanisms: An Experimental Investigation,” *mimeo*.

- [29] Fragiadakis, D. and P.Troyan (2017) “Improving matching under hard distributional constraints,” *Theoretical Economics*, 12, pp. 863 - 908.
- [30] Gale, D. and L.S. Shapley (1962) ”College Admissions and the Stability of Marriage,” *The American Mathematical Monthly*, 69, pp. 9-15.
- [31] Haeringer, G. and F. Klijn (2009) ”Constrained School Choice,” *Journal of Economics Theory*, 144, pp. 1921-1947.
- [32] Halafir, I.; M. Yenmez; and M. Yildirim (2013) “Effective affirmative action in school choice,” *Theoretical Economics*, 8, pp. 325-363.
- [33] Hassidim, A.; A. Romm; and R. Shorrer (2016) ““Strategic” Behavior in a Startegy-proof Environment,” *EC’16 Proceedings*.
- [34] Hatfield, J. and Milgrom, P. (2005) ”Matching with Contracts,” *The American Economic Review*, 95 (4), pp. 913-935.
- [35] Y. He “Gaming the Boston School Choice Mechanism in Beijing”, *mimeo*.
- [36] Hitsch, G.; A. Hortagsu; and D. Ariely (2010) “Matching and Sorting in Online Dating,” *The American Economic Review*, 100 (1), pp. 130-163.
- [37] Immorlica, N. and M. Mahdian (2005) “Marriage, Honesty, and Stability,” *SODA 2005*, pp. 53-62.
- [38] Kamada, Y. and F. Kojima (2015), “Efficient matching under distributional concerns: Theory and application,” *American Economic Review*, 105, 6799.
- [39] A. Kapor; C. Neilson; S. Zimmerman (2017) “Heterogeneous Beliefs and School Choice Mechanisms,” *mimeo*.
- [40] Kelso, A. and V. Crawford (1982) ”Job Matching, Coalition Formation, and Gross Substitutes,” *Econometrica*, 50 (6), pp. 1483-1504.
- [41] Kesten, O. (2010) “School Choice with Consent,” *The Quarterly Journal of Economics*, 125 (3), pp. 1297 - 1348.
- [42] Kesten, O. and U. Unver (2015) “A theory of school-choice lotteries,” *Theoretical Economics*, 10, pp. 543 - 595.
- [43] Kojima, F. and M. Manea (2006) “Strategy-Proofness of the Probabilistic Serial Mechanism in Large Random Assignment Problems,” *mimeo*.
- [44] Kojima, F. and P. Pathak (2009) “Incentives and Stability in Large Two Sided Matching Markets,” *American Economic Review*, 99 (3), pp. 608 - 627.

- [45] Kominers, S. and T. Sönmez (2016) “Matching with Slot-Specific Priorities: Theory,” *Theoretical Economics*, 11 (2), pp. 683-710.
- [46] 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.
- [47] Lee, S. (2012) “Incentive Compatibility of Large Centralized Matching Markets,” *mimeo*.
- [48] Q. Liu and M. Pycia (2016) “Ordinal Efficiency, Fairness, and Incentives in Large Markets,” *mimeo*.
- [49] Manea, M. (2009) “Asymptotic ordinal inefficiency of random serial dictatorship,” *Theoretical Economics*, 4, pp. 165 -197.
- [50] Menzel, Konrad (2015) “Large Matching Markets as Two-Sided Demand Systems,” *Econometrica*, 83 (3), pp. 897 - 941.
- [51] 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.
- [52] 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.
- [53] Rees-Jones, A. (2017) “Suboptimal behavior in strategy-proof mechanisms: Evidence from the residency match,” *Games and Economic Behavior*, *forthcoming*.
- [54] 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.
- [55] 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.
- [56] Roth, A.; U. Rothblum; and J. Vande Vate (1993) “Stable Matchings, Optimal Assignments, and Linear Programming,” *Mathematics of Operations Research*, 18, pp. 803-828.
- [57] Roth, A.; T. Sönmez; and M.U. Unver (2007) “Efficient Kidney Exchange: Coincidence of Wants in Markets with compatibility-Based Preferences,” *The American Economic Review*, 97 (3), pp. 828-851.
- [58] Roth, A. and M. Sotomayor (1990) *Two Sided Matching: a Study in Game-Theoretic Modeling and Analysis*, Cambridge University Press: Cambridge, England.
- [59] Roth, A. and J. Vande Vate (1990) “Random Paths to Stability in Two-Sided Matching,” *Econometrica*, 58 (6), pp. 1475 - 1480.

- [60] Roth, A. and J. Vande Vate (1991) “Incentives in Two-Sided Matching with Random Stable Matches,” *Economic Theory*, 1, pp. 31 - 44.
- [61] Shi, P. (2016) “Assortment Planning in School Choice,” *mimeo*.
- [62] Sönmez, T. and T. Switzer (2013) “Matching with (Branch-of-Choice) Contracts at the United States Military Academy” *Econometrica*, 81, pp. 451-488.
- [63] van der Vaart, A. and J. Wellner (1996) *Weak Convergence and Empirical Processes*, Springer-Verlag: New York.

A Proofs

Proposition 1. *Equations 3.5, 3.6, 3.8, and 3.7 are necessary and sufficient for a match x to be stable.*

Proof. Equations 3.5 and 3.6 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 3.7 holds. Equation 3.8 is a compressed formulation of my definition of stability. \square

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 $x(c, s) = m * \#\{(c, s) \text{ matches in the finite economy}\}$, I obtain a feasible and stable match in the continuum model. \square

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 $\pi_n^S \rightarrow \pi^S$ and $q_C^n \rightarrow q_C$. 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 $(\pi^S, (q_C)_{c \in \mathcal{C}})$. \square

Lemma 1. *Equation 3.4 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 3. *Any assignment satisfying equations 3.5 through 3.7 can be implemented.*

Proof. I rely on theorem 1 of Budish et al. [14]. 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. [14] 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:

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

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

where I have implicitly used $x(c, \emptyset)$ and $x(\emptyset, s)$ 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 3.5 and A.1 form a hierarchy, while equations 3.6 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. [14] implies \mathbf{X} can be implemented. \square

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

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

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 \emptyset (i.e., the agent is left unmatched). I define A as follows, where I use s to denote a generic student-type and c to denote a generic college. Unless stated otherwise, the entry of A is 0.

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

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 [18], 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 [18] 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 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 $\mathbf{q}_i \rightarrow \mathbf{q}$ I can choose \mathbf{y}_i such that $A\mathbf{y}_i = \mathbf{q}_i$ and $\mathbf{y}_i \rightarrow \mathbf{y}$. 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 $\varepsilon \in \mathbb{R}^{|\mathcal{C}|+|\mathcal{S}|}$ of \mathbf{q} and wish to solve

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

which is equivalent to solving

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

I solve the equivalent problem

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

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

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

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

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

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 $\mathbf{S}(\mathbf{q}; \underline{x}_C, \underline{x}_S)$ 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 $\mathbb{R}_{SC} \subset \mathbb{R}^{(|C|+1)(|S|+1)}$ 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

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

which insures that $\mathbf{S}(\mathbf{q}; \underline{x}_C, \underline{x}_S)$ 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 \mathbf{q} for a fixed choice of \underline{x}_C and \underline{x}_S .

Lemma 4. *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 $\mathbf{q} \in \mathbb{R}_+^{(|C|+|S|)}$ 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 $\delta \in \mathbb{R}_{SC}$ small and strictly positive in all dimensions where that is possible.⁵⁴ Let $\varepsilon = -A\delta < \mathbf{0}$, and note that as $\delta \rightarrow \mathbf{0}$ I have $\varepsilon \rightarrow \mathbf{0}$.

⁵⁴Recall that \mathbb{R}_{SC} has some dimensions without any nonzero elements.

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

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

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

Lemma 4 leads almost immediately to Theorem 1.

Theorem 1. *For topologically 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 $(q_c)_{c \in \mathcal{C}}, \pi^S$.*

Proof. First note that $S((q_c)_{c \in \mathcal{C}}, \pi^S) = \bigcup_{(\underline{x}_C, \underline{x}_S)} \mathbf{S}(\mathbf{q}; \underline{x}_C, \underline{x}_S)$ where the set of $(\underline{x}_C, \underline{x}_S)$ is finite. From the finiteness of this set, I have that, for a topologically generic set of $((q_c)_{c \in \mathcal{C}}, \pi^S)$, $S((q_c)_{c \in \mathcal{C}}, \pi^S)$ is either empty or continuous. However, $S((q_c)_{c \in \mathcal{C}}, \pi^S)$ cannot be empty by Proposition 2, which yields my claim. \square

Finally I can now prove my result on the approximate equilibria.

Theorem 2. *Under Assumption 1, for any $\varepsilon > 0$ there exists N^* such that if $N > N^*$ truthful declaration is an $(\varepsilon, 1)$ -Bayesian Nash Equilibrium. If Assumption 2 also holds, then for any $\varepsilon, \rho > 0$ there exists N^* such that if $N > N^*$ truthful declaration is an (ε, ρ) -Bayesian Nash Equilibrium.*

Proof. From the continuity of the mechanism (Assumption 1) and the linearity of the student utilities in probability, 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 [13] I know that for any $\varepsilon > 0$, I can choose N^* so that for $N > N^*$ I have that, if truthfulness is incentive compatible in the limit game, then truthfulness is an ε -Bayesian Nash equilibrium of the N -agent game.

For an agent to have a profitable deviation in an ε -Bayesian Nash equilibrium, her type must be such that her utility from two colleges c and c' are within ε of one another. From Assumption 2, the measure of the set of such agents vanishes as $N \rightarrow \infty$, which establishes my claim. \square

C Additional Tables

School Name	African-American	White	Asian	Hispanic
Brighton High	71.0 [66.8, 75.1]	14.9 [9.0, 22.0]	19.3 [14.0, 24.3]	59.9 [55.0, 64.5]
Excel High	38.5 [36.7, 40.7]	17.0 [11.0, 19.7]	12.2 [9.4, 15.4]	31.3 [29.4, 33.5]
Lyon	4.4 [3.0, 4.7]	2.0 [1.0, 3.0]	1.1 [0.1, 3.0]	3.5 [2.5, 5]
Snowden International, Chinese	9.0 [8.6, 9.4]	4.4 [3.8, 4.7]	2.4 [1.9, 3.1]	7.2 [6.9, 7.7]
Snowden International, French	10.1 [9.7, 10.7]	4.9 [4.5, 5.4]	2.7 [2.0, 3.4]	8.1 [7.8, 8.8]
Snowden International, Japanese	9.7 [9.4, 10.3]	4.8 [4.3, 5.3]	2.6 [1.9, 3.4]	8.1 [7.5, 8.6]
Snowden International, Spanish	10.1 [9.7, 10.6]	4.9 [4.4, 5.3]	2.6 [2.0, 3.4]	7.8 [7.8, 8.7]
Snowden International, All Languages	38.9 [37.4, 41.1]	19.0 [17.2, 20.6]	10.2 [8.4, 13.2]	31.3 [29.9, 33.8]
Another Course College	15.7 [15, 16.5]	7.7 [6.8, 8.4]	4.1 [3.1, 5.8]	12.6 [12.1, 14.9]
Urban Science Academy	66.5 [62.7, 70.5]	18.0 [11.0, 23.9]	8.1 [4.0, 13.9]	57.4 [52.3, 62.0]
Tech Boston Academy	62.0 [50.0, 75.0]	7.3 [2.9, 12.0]	4.3 [0.5, 9.4]	27.4 [21.0, 31.1]

Table 19: Enrollment Under the Constrained Optimal Mechanism for Diversity in Terms of the Number of Students

School Name	African-American	White	Asian	Hispanic
Brighton High	82.3 [75.3, 91.2]	9.0 [5.0, 13.1]	7.9 [4.9, 11.5]	69.9 [61.5, 77.7]
Excel High	47.8 [41.1, 53.7]	11.0 [7.1, 15.7]	14.0 [9.4, 19.7]	26.0 [20.9, 32.0]
Lyon	4.0 [1.6, 5.9]	2.1 [0.6, 4.0]	1.3 [0, 3.4]	3.6 [1.6, 5.9]
Snowden International, Chinese	8.3 [6.5, 10.1]	2.5 [1.4, 3.8]	3.0 [1.8, 4.4]	9.1 [7.3, 11.0]
Snowden International, French	8.8 [6.4, 11.8]	3.8 [2.1, 5.7]	0.4 [0.1, 0.9]	13.9 [11.2, 16.4]
Snowden International, Japanese	10.2 [8.2, 12.4]	2.7 [1.5, 4.0]	1.9 [1.0, 3.0]	10.8 [8.7, 12.8]
Snowden International, Spanish	12.0 [10.0, 13.9]	2.2 [0.8, 4.2]	0.9 [0.5, 1.5]	11.0 [9.1, 12.8]
Snowden International, All Languages	39.3 [34.9, 44.0]	11.2 [8.0, 14.3]	6.3 [4.5, 8.2]	44.9 [40.1, 50.1]
Another Course College	15.4 [12.2, 19.5]	3.6 [1.5, 5.8]	3.4 [1.2, 6.3]	19.4 [15.0, 23.5]
Urban Science Academy	71.4 [62.7, 81.5]	14.1 [9.2, 19.6]	6.3 [2.9, 9.8]	59.0 [50.6, 68.6]
Tech Boston Academy	72.1 [65.2, 79.3]	1.8 [0.3, 3.4]	5.1 [2.4, 8.0]	20.8 [14.8, 27.3]

Table 20: Enrollment Under the Gale-Shapley Mechanism in Terms of the Number of Students

School Name	African-American	White	Asian	Hispanic
Brighton High	74.8 [70.2, 78.2]	13.0 [8.2, 19.8]	15.4 [10.3, 21.4]	62.8 [57.4, 66.2]
Excel High	41.6 [38.5, 43.3]	13.3 [10, 16.4]	15.2 [10.8, 18.9]	29.9 [27.9, 33.3]
Lyon	4.0 [3.0, 4.9]	1.3 [0.1, 3.0]	2 [0, 3.3]	3.7 [2.4, 5]
Snowden International, Chinese	9.8 [8.9, 10.2]	3.0 [2.1, 4.4]	2.0 [1.1, 3.4]	8.0 [7.3, 9.0]
Snowden International, French	10.9 [10.1, 11.5]	5.0 [3.9, 5.9]	2 [0.7, 3.5]	9.1 [8.5, 10.0]
Snowden International, Japanese	10.3 [9.4, 10.9]	4.4 [3.4, 5.3]	2.5 [1.6, 3.9]	8.8 [7.8, 9.5]
Snowden International, Spanish	10.6 [10.0, 11.3]	4.0 [3, 5.6]	1.3 [0, 2.6]	8.9 [8.2, 9.6]
Snowden International, All Languages	39.9 [37.4, 41.7]	15.8 [13.3, 18.9]	7.5 [5.6, 11.0]	33.4 [31.3, 35.8]
Another Course College	16.8 [15.3, 17.7]	6.6 [4.8, 8.2]	4.0 [2, 7]	14.6 [12.6, 17.0]
Urban Science Academy	69.0 [64.7, 74.1]	17.0 [11.2, 23.0]	6.4 [2.6, 11.0]	57.6 [52.2, 62.0]
Tech Boston Academy	62.0 [50, 75]	2.0 [0, 5.4]	7.5 []	29.5 [22.0, 34.3s]

Table 21: Demographics at Overdemanded Schools In Convex Combination Solution in Terms of the Number of Students