

OPTIMIZING FOR DISTRIBUTIONAL GOALS IN SCHOOL CHOICE PROBLEMS

ABSTRACT. I investigate three goals of school choice: welfare, encouraging neighborhood schools, and diversity. I use optimization problems to find the best stable and incentive compatible match for any combination of these objectives. These problems assume there is a continuum of students and school seats, which allows me to describe the incentive compatibility conditions in a tractable form. I prove that the set of stable matchings is generically continuous in the distribution of students and the school capacities, which implies that the characterization of the possible stable matches in the continuum model approximates the set of stable matches in a matching market with a large, but finite, number of students. I then apply my framework to data from Boston Public Schools. If the mechanism conditions on demographics, the improvement (relative to the status quo) in student welfare is equivalent to moving 291 students (out of 3,479) to schools one rank higher in their preference lists. In contrast, if the mechanism does not condition on demographics, the welfare improvement is equivalent to moving only 25.1 students to schools one rank higher. Improvements in the distributional goals can be made (e.g., increasing enrollment in neighborhood schools by 50%) without reducing welfare or diversity.

1. INTRODUCTION

School choice mechanisms have been the subject of intensive market 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 [3]). 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 (i.e., stronger property rights over seats) than other students. For example, Boston Public Schools (BPS) gave higher priority for seats at a school to students that lived close to that school to encourage students to attend their neighborhood schools.¹ 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 (Gale and Shapley [31], Roth and Sotomayor [50]). There is much to recommend this advice. First, the Gale-Shapley algorithm is incentive compatible, so the students have no motive to try to manipulate the assignment system. Second, the outcome is stable, which reflects a notion of fairness that requires that an applicant with a high priority for a seat at an overdemanded school be allowed to claim a seat at that school before any lower priority applicants are enrolled.²

Key words and phrases. Matching; School Assignment; Diversity
JEL subject classification: D47, I24, C78

¹This policy arose in 1999 following the end of the use of ethnic criteria for school assignment. Explicit walk-zone priorities have recently been discontinued following the introduction of a new school assignment mechanism.

²Stability also requires that any student be able to claim an empty seat at a school if she prefers that school to the school to which she was assigned by the mechanism.

In principle, one can describe an optimal assignment of the 3,479 students in the BPS high school match to the school programs as the solution of a constrained optimization problem that takes the students' preferences as an input and treats stability and incentive compatibility (IC) as constraints. However, the IC constraints make this problem infeasible to solve. To verify that the IC constraints are satisfied, one must compare the utility each student receives given she truthfully reveals her preferences with the utility she would receive in the stable, IC match generated by the solution to the optimization problem given one deviation. Confirming that the match generated following the deviation is IC entails computing the solution to the optimization problem following two deviations. Taken to its logical conclusion, this implies that solving for a stable, IC match requires computing the stable, IC match that would result from *any* distribution of declared preferences. The curse of dimensionality renders this exercise infeasible.

To avoid this problem, this paper combines the idea of finding optimal matches through the use of constrained optimization problems with the 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.³ Since each student has a negligible effect on the aggregate outcome, I can write the incentive compatibility (IC) conditions in a tractable form. My framework opens up the possibility of systematically exploring the set of stable and IC matches to answer novel questions such as: Can one significantly improve on the outcome of the Gale-Shapley algorithm with respect to desiderata such as student welfare, school diversity, or encouraging neighborhood schools? How strong are the tensions between these goals?

I use my framework to solve for the optimal school assignment using data from the BPS high school match for the 2011-2012 school year, a time period during which BPS used the student-proposing Gale-Shapley school assignment algorithm. My data set includes information on student preferences, ethnic and socioeconomic background, and the location of the students' homes. I first consider the objective of maximizing student welfare subject to the capacity, stability (under the BPS priority structure), and IC constraints without conditioning the outcomes on student demographics.⁴ I measure a student's welfare as her assigned school's rank in her preference list, and the welfare generated by a school assignment is defined as the average student welfare across the population of students.⁵ Moving from the status quo assignment to the best stable and IC match that does not condition on student demographics generates a welfare gain equivalent to moving 25.1 (out of 3,479) students to a school one rank higher in their preference lists. I can also 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

³Liu and Pycia [44] proves that all efficient, symmetric, and strategy proof mechanisms lead to the same allocations. I consider mechanisms that are not "symmetric" in that student demographics affect outcomes.

⁴Although 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 [26], Kesten [38]).

⁵Although one could interpret my metric as an assumption about the utility functions of the students, I think of it as a social welfare function that selects a Pareto optimal outcome.

reveal that the welfare gain of an extra seat at Snowden International School is 85% higher than at any other overdemanded school.

I then demonstrate that if one allows the mechanism to condition on the verifiable demographics of the students, then a welfare increase equivalent to moving 291 students to a school one rank higher in their preference lists can be obtained in a stable, IC match. The reason the mechanism's performance improves is that allowing the mechanism to condition on the verifiable demographics of the students weakens the IC constraints as the students cannot misreport this verifiable information. If one considers it unfair that a mechanism treats students differently based on their demographics, then one can view the improved efficiency of a mechanism that does so as the **welfare cost of fairness**. In contrast, the welfare cost of stability is quite small. To fully explore the range of stable and IC matches, for the bulk of the paper I consider mechanisms that condition outcomes on student demographics.

My second goal is to assess the trade-offs between student welfare and encouraging ethnically and socioeconomically diverse student bodies.⁶ I measure the diversity of a school by how closely the composition of a school's student body matches the demographics of the entire BPS student population, and the average diversity of a match is the average diversity across all of the schools. 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, resulting in violent protests against the program. 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 [3]), 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 since the 1990s has led to a re-segregated school system.⁷

This history may suggest that there is tension between student welfare and diversity in school assignment. My analysis shows that this intuition is incorrect and that there are stable and incentive compatible school assignments that feature higher student welfare and more 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. For example, 35.4% percent of the students in overdemanded school programs are Hispanic in the status quo

⁶One 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.

⁷Barbara 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>

outcome, while 30% of the broader BPS student population is Hispanic. In addition, the percentage of Hispanic students enrolled in each overdemanded school program in the status-quo assignment ranges from 20.4% to 51.7%. In the match that maximizes diversity without reducing student welfare, 32.5% of the students in overdemanded schools are Hispanic, and the percentage of Hispanic students in each school ranges from 26.9% to 37.2%. Increasing diversity makes the student bodies of overdemanded schools look more like the BPS student population and shrinks the differences between the schools. I also trace-out a possibility frontier that describes all of the diversity and student welfare values that can be realized in a stable, IC match.⁸ There are only weak trade-offs between diversity and student welfare for the majority of the possibility frontier, which means that large increases in student welfare or diversity are possible without seriously impacting the other objective. In addition, the Gale-Shapley outcome is within the frontier.

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:⁹

"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."

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. For example, 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. Given the geographical segregation of Boston, one might conjecture that encouraging neighborhood schools would necessarily require that the schools be segregated. However, I show that one can create classes with 35.8% of their students drawn 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 again find that the Gale-Shapley outcome is well inside the frontiers and there are weak tensions between my objectives.

In Appendix F, I solve 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 demographics of the overdemanded schools more closely resemble that of BPS as a whole, and student welfare

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

⁹Quote taken from Dur et al. [23].

is increased by the equivalent of moving almost 162 students to a school one rank higher in their preference orderings. In summary, since there are only weak tensions between my goals, it is possible to find a school assignment that simultaneously achieves significant improvements across all three desiderata while satisfying stability and IC constraints.

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 literatures. First, while welfare objectives are clearly important in any school choice setting, there are distributional goals such as the encouragement of neighborhood schools or building ethnically and socioeconomically diverse student bodies that warrant consideration. 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. Third, my analysis of the BPS school assignment problem serves as a proof of concept for the constrained optimization approach.

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 problems 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 my distributional goals and the tension between my desiderata. Proofs are in Appendix A.

1.1. Related Literature. Descriptions of stability using constraints appears in the economics literature as early as Roth et al. [49]. These ideas were extended to many-to-one matching in Baïou and Balinski [13]. These prior works consider a finite set of agents, which means that, while useful for proving properties of matching problems, they cannot be used to tractably solve optimization problems (see Section 3.3).

The closest analogs to my work are Ashlagi and Shi [7] and Shi [51].¹⁰ Both papers share my view that matching problems can be tackled using constrained optimization problems featuring a continuum of students and school seats. The focus of Ashlagi and Shi [7] and Shi [51] is optimizing a convex combination of average and minimal student welfare subject to busing cost, capacity, and IC constraints. Instead of solving the school assignment problem directly (as I do), Ashlagi and Shi [7] and Shi [51] prove that one can reformulate an optimal matching problem without priorities as a random assortment problem in which the mechanism offers students randomized menus of school assignments.

There are two major difference between my work and Ashlagi and Shi [7]. First, I can incorporate priority structures (e.g., walk-zone priority) into my framework, which can be practically important if parents view these as inviolable property rights. Second, since I work directly with

¹⁰Featherstone [29] 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 [29] does not include incentive compatibility constraints.

the students' ordinal rankings of the school programs, I do not need to impose parametric restrictions on the students' utility functions. In contrast, Ashlagi and Shi [7] assume the students have multinomial logit (MNL) utilities to make the optimization problem tractable in the presence of a large number of school programs.¹¹ Although my techniques may not be feasible for very large school systems (e.g., New York City), they can be applied in many school assignment settings.

A sense of the restrictiveness of the MNL utility assumption can be gleaned by a comparison with the functional forms assumed in structural estimates of matching models, which were chosen to capture the underlying distribution of preferences without regard for the tractability of a follow-on optimization problem. Early papers in the structural school choice literature often adopted an MNL error structure for tractability (e.g., Hastings et al. [33], He [34], Ajayi [5], Burgess et al. [17], and Fack, Grenet, and He [28]).¹² The recent literature has used more flexible models of unobserved preference components that include distributions of preference shocks that are independent across students and correlated across schools, preference shocks that are correlated across students and independent across schools, and preference shocks that are heteroskedastic (e.g., Calsamiglia et al. [18], Pathak and Shi [46], Kapor et al. [37], Abdulkadiroğlu et al. [1], Agrawal and Somaini [4], Che and Tercieux [20]). These papers model the unobserved component of preferences as random coefficients that are drawn from flexibly specified multivariate normal distributions, which allows them to be estimated by Markov Chain Monte Carlo or Gibb's sampler techniques.¹³ For example, Abdulkadiroğlu, Agarwal, and Pathak [1] includes preference shocks that are independent across students and arbitrarily correlated across schools as well as interactions between student-specific unobserved shocks and school program attributes that generate heteroskedasticity.

An early example of a matching paper featuring a continuum of agents is Azevedo and Leshno [11], which assumes an exogenous distribution of student preferences and priorities that do not admit indifferences. The authors characterize a stable match in terms of priority number cutoffs, and show that the continuum and large finite-agent models yield approximately the same equilibrium. Because Azevedo and Leshno [11] assumes strict priorities, the results of that work do not apply to my setting. Later works used continuum models to provide conditions under which large matching markets are approximately stable (e.g., Azevedo and Hatfield [12], Che, Kim and Kojima [19]). There is also a vein of the large market literature that establishes the equivalence (or lack thereof) between different mechanisms (e.g., Kojima and Manea [40], Manea [45], and Liu and Pycia [44]). Che and Tercieux [20] uses a large market environment to

¹¹Shi [51] weakens the assumption of MNL utilities (e.g., a nested logit form can be used), but the tractability continues to turn on the availability of high-speed algorithms for solving optimization problems given these preferences.

¹²Many of these papers imposed functional forms in order to infer underlying preferences from strategic declarations to nontruthful school choice mechanism (e.g., the original Boston mechanism).

¹³There is some debate as to whether the added flexibility is useful when school specific dummy variables are included (see Pathak and Shi [46]).

develop the Deferred Acceptance with Circuit Breakers mechanism that is approximately stable and approximately IC in large market settings.

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 [48] points out that the set of stable matches in the National Residency Matching Program is small, which implies approximate strategy-proofness. Immorlica and Mahdian [35] and Kojima and Pathak [41] 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 IC as they grow. Lee [43] generates a similar result without a restriction to short preference lists.

There have been a number of theoretical papers that study the limits on incorporating diversity into school choice schemes, some of which add either quotas or reserves into the mechanism (Budish et al. [16], Ehlers et al. [25], Erdil and Kumano [27], Fragiadakis and Troyan [30], Hafalir et al. [32], Kamada and Kojima [36], Kominers and Sönmez [42]). Echenique and Yenmez [24] axiomatize notions of diversity that are compatible with stability.

While encouraging neighborhood schools is an explicit goal of BPS, less has been written on it. Ashlagi and Shi [8] studies using correlated priority tie-breakers to place students from the same neighborhood in the same school while leaving assignment probabilities unchanged. In contrast, I seek to increase cohesion within neighborhood schools (as opposed to some school), and I explicitly consider altering the assignment probabilities to achieve this. Dur et al. [23] show that the effect of walk-zone priorities on encouraging neighborhood schools depends heavily on the order in which students apply to the walk-zone and open seats at each school program.

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 is used to find a match, but the use of random tie-breakers can result in an inefficient outcome. Erdil and Ergin [26] provides an algorithm for computing the student-optimal stable match once an initial match has been computed using the Gale-Shapley algorithm, but this algorithm is not IC. Kesten [38] 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.

2. DATA

I use administrative data that covers 3,479 rising 9th grade students in the BPS system during the 2011 - 2012 school year. 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 in Appendix B. The data include all of the inputs to the school assignment mechanism. Since the Gale-Shapley

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

mechanism is IC, I take the preferences declared by the students at face value.¹⁴ BPS chose the priority structure to reflect several goals: encouraging neighborhood schools, facilitating enrolling siblings at the same school, and allowing students to continue at their current school.

The Gale-Shapley mechanism used by BPS divides the seats in each school program into two equally sized groups: a set of **walk-zone seats** and a set of **open seats**. Each group of seats is treated as a distinct program 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. 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 school as an inviolable property right, it may be hard to abandon these priorities when changing the school choice mechanism.

A student with **guaranteed** or **guaranteed-walk** priority at a school program is insured a place in that program if the student ranks it.¹⁵ **Sibling** and **sibling-walk** apply if the student has a sibling already enrolled in that program. **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 tie-breaking was used by BPS to assign a unique priority rank to each student at each program.

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. Therefore, the effective priority structure can be described by Table

¹⁴No student ranked more than 8 overdemanded school programs before ranking an underdemanded program, so truthful revelation of the rank-order list is both feasible and incentive compatible for the students in my data set.

¹⁵Students 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	
No Priority	Walk, No Priority

TABLE 2. Effective Priority Structure

Ethnicity	BPS Overall		Overdemanded Schools		
	Number	% of Data	Average	Min	Max
African American	938	37.3	47.7	26.9	67.6
White	457	18.2	8.0	2	25.0
Asian	304	12.1	5.4	0	10.9
Hispanic	755	30.0	36.2	16.7	54.8
Native American	11	0.4	0.6	0	3.85
Mixed-Other	48	1.9	2.0	0	8.3

TABLE 3. Ethnicity Demographics Across BPS Student Population

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

To the best of my knowledge, there is neither an explicit nor an implicit goal of using the priority structure to favor students of demographic groups that are perceived to be disadvantaged.

Throughout the analysis I exclude students that require special education programs as these students are effectively matched through a different market. There are also English as a second language (ESL) students in the BPS system. BPS ranks ESL students on a five point proficiency scale, and the least proficient students, those that score a three or lower on the scale, are assigned to ESL programs. Since these low scoring ESL students are also effectively matched in a distinct market, I exclude these students from my analysis. In the end, I have 2,513 students and 28 high school programs (of which 10 programs are overdemanded).

Finally, I have administrative data on each student including ethnicity, whether the student participated in a free school lunch program, and the zip code of the student's home residence. Table 3 shows the demographics of the BPS student population as a whole as well as the average, minimum, and maximum percentage of the student body drawn from each of the six ethnic groups across the overdemanded programs.¹⁶ 66.5% of the BPS students receive a free lunch, and I use participation in a free lunch program as an indicator of low socioeconomic status. 78% of

¹⁶The demographics of the overdemanded programs do not match the overall demographics of BPS. This is the result of White and Asian students disproportionately either entering examination schools or choosing to attend an underdemanded school.

the students at the 10 overdemanded schools participate in a free school lunch program, and the rate of free lunch participation in the 10 overdemanded programs varies from as low as 65.4% to as high as 90.2%. I assume that demographic traits are verifiable and need not be revealed by students. In reality, some of these data may be imperfectly verifiable. For example, parents could choose to reside in a zip code that would, in the parents' opinion, lead to favorable treatment by the mechanism. The results provided here may not provide a reliable estimate of the long-run distribution of outcomes if the parents' endogenous residency choices change significantly.¹⁷

There is the capacity to enroll 688 students at the 10 overdemanded high school programs. I identify overdemanded school programs, schools that some student prefers to his or her current assignment, as those that rejected at least one student during the Gale-Shapley algorithm, which implies a binding capacity constraint. Underdemanded schools, programs that no student prefers to his or her current assignment, 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. In reality, whether or not a school is underdemanded is an endogenous outcome. Even though it is not an ideal assumption, I 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 becoming overdemanded if it is underdemanded in my data. 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. Second, 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 the school has not reached a physical constraint on the number of students that can enroll.

3. MATCHING IN THE CONTINUUM FRAMEWORK

I now describe the constrained optimization problems I use in my analysis. Since the continuum model will be treated as the limit of a sequence of finite economies in Section 3.4, 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 $\pi^S(s)$ continuum of each type $s \in \mathcal{S}$ and a measure one continuum of students in total. Each type $s = (v, \succeq_s) \in \mathcal{S}$ is characterized by a verifiable trait, $v \in \mathcal{V}$, and a preference ordering over

¹⁷This problem also exists in the mechanism used at the time the data I use was collected since a student's residence determines at which schools that student has walk-zone priority.

the finite set of school programs, $\succeq_s \in \mathcal{P}_S$. I assume the set of verifiable traits \mathcal{V} is finite and each student-type's realization of v is known to the mechanism designer. In my context v can, and for most of the paper does, include student ethnicity, socioeconomic status, the zip code of her residence, and the student's priority at each school program.¹⁸

Each school program is defined by a unique trait that identifies the school program (the program'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 school program.¹⁹ I do not require that the priority orderings be strict. School program c has a continuum of seats of measure $q_c \in [0, 1]$, where q_c denotes the fraction of the student population that can be enrolled in school program c . A generic school program-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 school programs.

For school programs, \emptyset denotes an outcome wherein a seat is left unfilled. In the BPS system all student types are acceptable to all school programs, so $s \succeq_c \emptyset$ for all school programs c and student-types s .²⁰ 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 school program c to her most preferred underdemanded school.

The action set of the agents is a message describing an ordinal ranking of the school programs.²¹ Prior to sending a message to the mechanism, the students know their own type, the capacities of the school programs, the distribution of student-types, and how the mechanism maps the distribution of declared types and the student's own message into a distribution of possible assignments for that student. The students do not know the realization of this uncertainty prior to submitting their message to the mechanism.

3.1. Matching Notation. A **match** is a function $x : (\mathcal{C} \cup \emptyset) \times (\mathcal{S} \cup \emptyset) \times \Delta(\mathcal{S}) \rightarrow [0, 1]$ where $\Delta(\mathcal{S})$ denotes the space of probability measures over \mathcal{S} . $x(c, s; \pi^S)$ denotes the fraction of the student population comprised of students of type s matched with school program c given a distribution of student-types π^S . Implicitly, this notation allows each student to match with one school program, but each school program may be matched with multiple students. My preferred interpretation is that a match represents a stochastic assignment:

$$(3.1) \quad \Pr\{\text{Particular student of type } s \text{ matched to school program } c\} = \frac{x(c, s)}{\pi^S(s)}.$$

¹⁸Section 4.1 includes only the student's priority at each program in her realization of v . Section 4.2 gradually includes other pieces of verifiable information into v to study the importance of conditioning on demographics.

¹⁹The objective functions of my optimization problems capture violations responsiveness of the schools' "preferences."

²⁰It 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.

²¹A typical revelation mechanism would require each student to report her underlying cardinal utilities. I use this alternative message space as it accords better with my data.

3.2. **Stability.** A pair (c, s) of school program and student type is a **blocking pair** if:

$$(3.2) \quad \sum_{\{s':s' \succeq_c s\}} x(c, s'; \pi^S) < q_c \text{ and } \sum_{\{c':c' \succeq_s c\}} x(c', s; \pi^S) < \pi^S(s).$$

The first condition implies that program c either has an empty seat or that student-type s has a higher priority than some student-type that has a seat at program c . The second condition implies that some student of type s is either unmatched or matched to a program to which c is strictly preferred by the student. A match x is **stable** if it does not admit blocking pairs.^{22,23}

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

$$(3.3) \quad \text{For all } s \in \mathcal{S}, \quad \sum_{c \in \mathcal{C} \cup \{\emptyset\}} x(c, s; \pi^S) = \pi^S(s)$$

$$(3.4) \quad \text{For all } c \in \mathcal{C}, \quad \sum_{s \in \mathcal{S} \cup \{\emptyset\}} x(c, s; \pi^S) = q_c$$

$$(3.5) \quad \text{For all } s \in \mathcal{S}, \quad \sum_{\{c \in \mathcal{C} : \emptyset \succ_s c\}} x(c, s; \pi^S) = 0$$

$$(3.6) \quad \text{For all } (c, s), \quad \left(\pi^S(s) - \sum_{\{c' \in \mathcal{C} \cup \{\emptyset\} : c' \succeq_s c\}} x(c', s; \pi^S) \right) \left(q_c - \sum_{\{s' \in \mathcal{S} \cup \{\emptyset\} : s' \succeq_c s\}} x(c, s'; \pi^S) \right) = 0$$

The first two conditions are feasibility constraints. Equation 3.3 insures that an individual student is matched with at most one school program seat. Equation 3.4 implies that school program c enrolls at most a fraction q_c of students. Equation 3.5 requires that each student is matched to a program that is weakly preferred to her most preferred underdemanded school program. Equation 3.6 requires that the match is stable (equation 3.2).

Proposition 1. *Equations 3.3 - 3.6 are necessary and sufficient for a match $x(\circ, \circ; \pi^S)$ to be stable.*

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 school program c . Given a match x , denote the worst outcome for students of type $s \in \mathcal{S}$ 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 $\sum_{c \in \mathcal{C}} x(c, s; \pi^S) = \pi^S(s)$, let $\underline{x}_S(s) = \max\{c : c \preceq_s c' \text{ for all } c' \text{ such that } x(c', s) > 0\}$ where the maximum is taken with respect to \succeq_s .²⁴ If $\sum_{c \in \mathcal{C}} x(c, s; \pi^S) < \pi^S(s)$, then $\underline{x}_S(s) = \emptyset$. If \succeq_s admits indifferences, then \underline{x}_S can be a multivalued correspondence. I let $c \succeq_s \underline{x}_S(s)$ denote that c is weakly preferred by s to all members of $\underline{x}_S(s)$.

²²Stronger 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.

²³My definition of stability is identical to the notion of **ex ante stability** proposed by Kesten and Unver [39]. In the limit model or an economy with multiple agents of the same type, the two notions coincide.

²⁴The awkward definition of $\underline{x}_S(s)$ is used so that $\underline{x}_S(s)$ captures all of the school programs c such that the student is indifferent between c and the least preferred school program to which students of type s are matched. In our application, preferences are strict and $\underline{x}_S(s)$ is a single-valued function.

The correspondence $\underline{x}_C(c) \in \mathcal{S} \cup \{\emptyset\}$ describes the lowest priority student-type assigned to program c in match x . Define $\underline{x}_C(c) = \max\{s : s \preceq_c s' \text{ for all } s' \text{ such that } x(c, s'; \pi^S) > 0\}$ if $\sum_{s \in \mathcal{S}} x(c, s; \pi^S) = q_c$, where the maximum is taken with respect to \succeq_c . If $\sum_{s \in \mathcal{S}} x(c, s; \pi^S) < q_c$, then $\underline{x}_C(c) = \emptyset$. If \succeq_c admits indifferences between student types, then \underline{x}_C is a correspondence.

I use the following lemma to simplify the stability conditions in my optimization problems.

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

The description of the stable set provided by equations 3.3 - 3.6 is concise, but it presents practical challenges since the stability constraint is nonconvex and, in general, solving nonconvex optimization problems is NP-hard. 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 , the associated stable subset satisfies equations 3.3-3.5 and

$$(3.7) \quad \text{For all } s \in \mathcal{S}, \quad \sum_{\{c: c \succeq_s \underline{x}_S(s)\}} x(c, s; \pi^S) = \pi^S(s)$$

$$(3.8) \quad \text{For all } c \in \mathcal{C}, \quad \sum_{\{s: s \succeq_c \underline{x}_C(c)\}} x(c, s; \pi^S) = q_c$$

Computing the realizations of \underline{x}_S and \underline{x}_C that are compatible with stability is simple. Consider a potential choice $\underline{x}_C(c) \in \mathcal{S}$. For all $s \in \underline{x}_C(c)$, the set of values of $\underline{x}_S(s)$ compatible with stability is $\{c' : c' \succeq_s c\}$. Having specified each $\underline{x}_C(c)$, it is straightforward to construct the set of $\underline{x}_S(s)$ compatible with stability. There is no guarantee that equations 3.3-3.5 together with 3.7 - 3.8 admit a feasible point for every choice of \underline{x}_C and \underline{x}_S , but checking feasibility is fast and eliminates most of the candidate $(\underline{x}_S, \underline{x}_C)$.

There are four kinds of stability constraints I need to impose. First, some student-types are insured a seat at school program c , and I denote this set of student-types as $\mathcal{I}(c)$. As 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 $s \in \mathcal{I}(c)$. For each $c \in \mathcal{C}$ and each $s \in \mathcal{I}(c)$ I have a constraint of the form

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

which requires $s \in \mathcal{I}(c)$ be assigned to a program that is weakly preferred to program c .

Second, since the overdemanded school programs are, by definition, overdemanded, stability requires that all of the seats at overdemanded programs are assigned to some student-type:

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

If one instead imposes the capacity conditions as inequalities, then an optimal match might leave 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 stability constraints apply to walk-zone school programs. Let $\overline{\mathcal{W}}$ be the set of walk-zone school programs, and for $c \in \overline{\mathcal{W}}$ 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 takes the form

$$(3.9) \quad \text{If } c \in \overline{\mathcal{W}} \text{ and } \underline{x}_C(c) = \mathcal{W}(c), \quad \sum_{s \in \mathcal{W}(c) \cup \mathcal{I}(c)} x(c, s; \pi^S) = q_c$$

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

$$(3.10) \quad \text{If } c \in \overline{\mathcal{W}}, \underline{x}_C(c) \neq \mathcal{W}(c), \text{ and } s \in \mathcal{W}(c), \quad \sum_{\{c' \in \mathcal{C} \cup \{\emptyset\} : c' \succeq_s c\}} x(c', s; \pi^S) = \pi^S(s)$$

Equation 3.10 requires that any student-type s with walk-zone priority at school program c must be assigned to a program that she weakly prefers to c . Otherwise (c, s) forms a blocking pair.

Summarizing, I impose the following stability constraints

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

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

$$(3.13) \quad \text{If } c \in \overline{\mathcal{W}} \text{ and } \underline{x}_C(c) = \mathcal{W}(c), \quad \sum_{s \in \mathcal{W}(c) \cup \mathcal{I}(c)} x(c, s; \pi^S) = q_c$$

$$(3.14) \quad \text{If } c \in \overline{\mathcal{W}}, \underline{x}_C(c) \neq \mathcal{W}(c), \text{ and } s \in \mathcal{W}(c), \quad \sum_{\{c' \in \mathcal{C} \cup \{\emptyset\} : c' \succeq_s c\}} x(c', s; \pi^S) = \pi^S(s)$$

Given 10 overdemanded school programs, there are 1,024 possible combinations of \underline{x}_C and \underline{x}_S to consider, and I find that 24 of the 1,024 sets of constraints satisfy equations 3.3-3.5 and 3.7 - 3.8. The feasibility of each of the 1,024 sets of stability constraints can be tested in polynomial time. Given that the full optimization problem (i.e., including the ODIC constraints) can be solved in polynomial time for each of the 24 subproblems, finding the global optimum is tractable even with limited computational resources.

3.3. Incentive Compatibility Constraints. I use **ordinal dominance incentive compatibility (ODIC)** as my definition of incentive compatibility, which is based on the ordinal ranking of the students' preferences.²⁵ The ODIC conditions are defined as follows:²⁶

²⁵One might have assumed students have cardinal utilities for each school program and then impose Bayesian incentive compatibility (BIC) conditions that require that the expected utility from a truthful report exceeds that of any misreport. The ODIC conditions are equivalent to requiring that truthful reporting of the ordinal ranking of school programs be optimal for any cardinal utility that generates the ordinal ranking, and thus are stronger than BIC.

²⁶The ODIC constraints in the continuum model are equivalent to a requirement of envy-freeness between student-types with the same verifiable trait.

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

$$(3.15) \quad \sum_{\{c': c' \succeq_s c\}} \frac{x(c', s; \pi^S)}{\pi^S(s)} \geq \sum_{\{c': c' \succeq_{s'} c\}} \frac{x(c', s'; \pi^S)}{\pi^S(s')}$$

π^S appears on both sides of equation 3.15, which means that a single agent's deviation has no effect on the aggregate distribution of actions. The ODIC conditions impose up to $|\mathcal{C}|$ constraints for each student-type pair s, s' . The first constraint implies that the probability of matching with student-type s 's most preferred school is not increased by declaring $s' \neq s$. The second constraint implies that the probability of matching with student-type s 's two most preferred schools is not increased by declaring $s' \neq s$. The k^{th} constraint implies that the probability of matching with any of student-type s 's k most preferred schools is not increased by declaring $s' \neq s$. When all of these constraints hold, the match is ODIC for student-type s .

Since the verifiable components of the students' types are known to BPS, I only need to impose constraints that insure it is optimal for students to declare the nonverifiable components of their type (i.e., their preferences) truthfully. When I allow the mechanism to condition on a subset of the verifiable traits, I only impose ODIC constraints between student-types s, s' that share the same realizations of the conditioning traits. The most restrictive set of ODIC constraints I consider applies ODIC constraints between student-types s, s' that share the same priorities. Therefore, I impose ODIC constraints between (for example) students from different zip codes that have the same priorities. The least restrictive set of ODIC constraints I consider only applies ODIC constraints between student-types s, s' that share the same realization across every verifiable attribute. Under these constraints, ODIC constraints do not apply between students living in different zip codes, and so the outcome provided to a student in one zip code is less constrained by the outcome provided a student in another zip code.

The existence of a stable and ODIC match is known for markets with a finite number of agents. When the distribution of student-types and school program capacities admit only rational values, I can treat the continuum of students (program seats) of each student-type (program) as if it were composed of a finite number of agents, where each agent represents the same positive measure of students (school program seats). It is straightforward to compute a stable and ODIC match using a continuum analog of the Gale-Shapley mechanism. When the distributions have irrational values, I use limit analysis techniques to argue that stable matches must exist.

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

I now discuss why the ODIC constraints would make the constrained optimization problem computationally intractable for the analogous problem with a finite number of students. In a

model with N agents, the ODIC constraint would be

$$(3.16) \quad \text{For all } c \in \mathcal{C}, \quad \sum_{\{c':c' \succeq_s c\}} \frac{x(c', s; \pi^S)}{\pi^S(s)} \geq \sum_{\{c':c' \succeq_s c\}} \frac{x(c', s'; \pi^S + \frac{1}{N} [\delta_{s'} - \delta_s])}{\pi^S(s') + \frac{1}{N}},$$

where δ_s is a Dirac measure placing probability 1 on student type s . The right-hand side of Equation 3.16 reflects how the distribution of declared types changes when one of the agents changes her declared preferences. To determine if Equation 3.16 is satisfied, I would need to compute the stable and ODIC match given the altered distribution of declarations. 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

$$(3.17) \quad \text{For all } c \in \mathcal{C}, \quad \sum_{\{c':c' \succeq_s c\}} \frac{x(c', s; \pi^S)}{\pi^S(s)} \geq \sum_{\{c':c' \succeq_s c\}} \frac{x(c', s'; \pi^S + \frac{1}{N} [\delta_{s'} - \delta_s] + \frac{1}{N} [\delta_{s''} - \delta_{s'''}])}{\pi^S(s') + \frac{1}{N}}$$

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 impractical with even a handful of students.

The implementation of the ODIC constraints is detailed in Appendix D. The potential size of the set of ODIC constraints raises the question of whether the constrained optimization problem scales well. There are many school systems that have a small enough number of student-types and schools that one could apply my model. For example, Chicago has 11 selective enrollment high schools, so my methods ought to be tractable if one wished to redesign the enrollment system for this subset of the Chicago school system. Similarly, the San Francisco public school system includes 17 high schools, some of which may not be overdemanded. I conclude that it is likely my techniques are applicable to many, if not all, public school systems.

3.4. Linking the Finite and Continuum Models. This subsection addresses the links between the continuum optimization problems solved in Sections 4 and 5 and the analogous problems with a finite number of agents. I begin by proving Theorem 1, which provides conditions under which the stable set and the set of stable and ODIC matches are generically continuous. Theorem 2 then uses the continuity result to prove that a continuum match that satisfies the ODIC constraints is approximately ODIC in a finite agent market. Theorem 3 then uses the continuity to prove that the maximum of a continuous function over the set of feasible and stable matches in a finite agent market approaches the maximum in the continuum market as the number of agents grows. I close by summarizing Appendix C, which describes how to implement a solution to an optimization problem as a mixture over deterministic school assignments.

Now I turn to the continuity of the sets of matches of interest. Let $S((q_c)_{c \in \mathcal{C}}, \pi^S)$ denote the set of stable matches given school program capacities $(q_c)_{c \in \mathcal{C}}$ and a student-type distribution π^S .

Let $SIC((q_c)_{c \in \mathcal{C}}, \pi^S)$ denote the subset of $S((q_c)_{c \in \mathcal{C}}, \pi^S)$ that also satisfies the ODIC constraints. A property of $S((q_c)_{c \in \mathcal{C}}, \pi^S)$ or $SIC((q_c)_{c \in \mathcal{C}}, \pi^S)$ is **topologically generic** if it holds over a dense, open set of π^S and $(q_c)_{c \in \mathcal{C}}$.²⁷ My later results hinge on the generic continuity of these sets.

Theorem 1. $S((q_c)_{c \in \mathcal{C}}, \pi^S)$ is continuous in $((q_c)_{c \in \mathcal{C}}, \pi^S)$ for topologically generic choices of $((q_c)_{c \in \mathcal{C}}, \pi^S)$ where $q_c > 0$ for all c . $SIC((q_c)_{c \in \mathcal{C}}, \pi^S)$ is continuous in $((q_c)_{c \in \mathcal{C}}, \pi^S)$ for topologically generic choices of $((q_c)_{c \in \mathcal{C}}, \pi^S)$ where $q_c > 0$ for all c and π^S has full support.

Recall that $SIC((q_c)_{c \in \mathcal{C}}, \pi^S)$ denotes the set of stable and ODIC matches in the continuum, let $SIC_N((q_c)_{c \in \mathcal{C}}, \pi^S)$ denote the set of stable matches that satisfy the N -agent ODIC condition (equation 3.17). Define the optimizers of a continuous objective function $F(x)$ as follows

$$x(\circ, \circ; \pi^S) = \underset{x \in SIC((q_c)_{c \in \mathcal{C}}, \pi^S)}{\operatorname{argmax}} F(x) \text{ and } x_N(\circ, \circ; \pi^S) = \underset{x \in SIC_N((q_c)_{c \in \mathcal{C}}, \pi^S)}{\operatorname{argmax}} F(x)$$

From the Theorem of the Maximum, we know that the solution to an optimization problem is upper hemicontinuous under continuous changes in the space of constraints. Theorem 1 therefore implies the correspondence $x(\circ, \circ; \pi^S)$ is upper hemicontinuous in π^S , and continuous if the optimizer is unique. When $F(x)$ is linear, uniqueness must be checked numerically. However, if $F(x)$ is strictly concave at the solution, then uniqueness follows automatically. I do not observe nonunique solutions in the problems I solve in Sections 4 and 5, which motivates the following assumption.

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

My goal is to prove that if truth-telling is an exactly optimal strategy in the mechanism with a continuum of agents, then truth-telling is an approximately optimal strategy for all of the agents. By “approximately optimal” I mean truthfulness satisfies the following ε -ODIC constraint.

Definition 2. $x(c, s; \tilde{\pi}^S)$ satisfies ε -ODIC for $\varepsilon > 0$ under measure $\tilde{\pi}^S$ in the N -agent economy if for all s, s' , and c

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

where $\tilde{\pi}^S = \pi^S + \frac{\delta_{s'} - \delta_s}{N}$ where δ_s is a Dirac measure on s .

I can now state my approximate equilibrium result.

Theorem 2. If Assumption 1 holds, then for any $\varepsilon > 0$ there exists $N^*, \delta > 0$ such that if $N > N^*$ and $\|\pi^S - \tilde{\pi}^S\| < \delta$, then $x(c, s; \tilde{\pi}^S)$ satisfies ε -ODIC under measure $\tilde{\pi}^S$.

²⁷I employ the usual topology over Euclidean spaces.

Let $x_N(c, s; \pi^S)$ denote a stable and ODIC match in the N -agent economy. For my remaining results, I need to make an assumption analogous to Assumption 1 about $x_N(c, s; \pi^S)$.

Assumption 2. Let $\{x_N(c, s; \pi^{S,N})\}_{N=1}^{\infty}$, $\pi^{S,N} \rightarrow \pi^S$, denote a sequence of optimizers of $F(x)$ in $SIC_N((q_c)_{c \in \mathcal{C}}, \pi^{S,N})$. Assume that $x_N(c, s; \pi^S) \rightarrow x_{\infty}(c, s; \pi^S)$ where $x_{\infty}(c, s; \pi^S)$ is continuous in an open neighborhood of π^S for all (c, s) .

Assumption 2 rules out mechanisms that use discontinuities to enforce incentive constraints. On the other hand, using discontinuities in this fashion exposes the participants to a high degree of risk as small changes in the aggregate distribution of types would result in large changes in the mechanism's outcome. Moreover, mechanisms that are used in practice tend to satisfy Assumption 2 (see, for example, Ashlagi and Shi [7] and Azevedo and Budish [9]).

Now I consider whether the continuum model captures all of the possible feasible and stable matches in the finite agent setting. Theorem 3 makes two claims. First, it proves that the optimal value of the problem solved using the discrete model converges to the optimal value of the continuum constrained optimization problem as N grows so long as the limit is continuous.²⁸ Second, Theorem 3 proves that if the optimizer of the continuum model is unique, then optimizers of the N -agent market must approach the optimizer of the limit market as N grows. I use the notation $SIC_{N,\varepsilon}((q_c)_{c \in \mathcal{C}}, \pi^S)$ denote the set of stable matches that satisfy the N -agent ε -ODIC condition (equation 3.18)

Theorem 3. Let Assumptions 1 and 2 hold. If $V_N(\pi^{S,N}) = F(x_N(c, s; \pi^{S,N}))$ and $V(\pi^S) = F(x(c, s; \pi^S))$, then $\limsup_{N \rightarrow \infty} V_N(\pi^{S,N}) = \liminf_{N \rightarrow \infty} V_N(\pi^{S,N}) = V(\pi^S)$. Moreover, if $x(c, s; \pi^S)$ is the unique maximizer, it must be that $x_N(c, s; \pi^S) \rightarrow x(c, s; \pi^S)$.

Appendix C proves that one can implement a match from the continuum game in the finite game using mixtures over deterministic matches that individually satisfy the same stability, capacity, and individual rationality constraints that $x(\circ, \circ; \pi_E^{S,N})$ satisfies. The key is using results from Budish et al. [16] that extend the Birkhoff-von Neumann theorem to settings with constraints. See Appendix C for definitions and other details.

4. WELFARE MAXIMIZATION IN THE BOSTON HIGH SCHOOL MATCH

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. In order to use my optimization approach, I need to define a social welfare function (SWF) to use as an objective function to select between the constrained Pareto optimal matches that satisfy stability and ODIC. I describe my SWF in three steps. First, $|\{c' : c' \succeq_s c\}|$ is the

²⁸One can use the Arzelá-Ascoli Theorem to prove that any sequence $\{x_N(c, s; \pi^{S,N})\}_{N=1}^{\infty}$ has a convergent subsequence, but only under the condition that the members of the sequence are drawn from an equicontinuous family.

	Best Stable, ODIC match	Gale-Shapley Mechanism	Worst Stable, ODIC match
$R(x)$	1.752 [1.721, 1.790]	1.762 [1.729, 1.796]	1.794 [1.758, 1.830]

TABLE 4. Welfare Analysis

number of schools that student s weakly 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 $\sum_{c \in \mathcal{C}} \frac{x(c,s;\pi^s)}{\pi^s(s)} |\{c' : c' \succeq_s c\}|$. My SWF is the average expected rank over the set of students

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

Throughout I refer to $R(x)$ as the average welfare generated by match x . I focus on $R(x)$ due to its salience to policy-makers as a measure of school assignment effectiveness.²⁹ Maximizing welfare involves minimizing $R(x)$ subject to the constraints on capacity and individual rationality (equations 3.3-3.5), stability (equations 3.11-3.14), and ODIC (equation 3.15).

Throughout Section 4.1, each student's verifiable trait v includes only the student's priorities at each school program. This means that the mechanisms I describe provide the same outcome to any two students with the same preferences and priorities.³⁰ I do this to provide as close a comparison as possible to the Gale-Shapley mechanism, which does not treat students differently based on demographics traits such as (for example) ethnicity.

Table 4 describes the average welfare of the best and worst stable and ODIC matches that do not condition on demographics as well as the welfare generated by the Gale-Shapley mechanism.³¹ 95% confidence intervals are provided beneath the statistics. The difference between the best and worst stable and ODIC matches is 0.042, 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.

Table 5 displays the effect of removing the stability constraints and/or ODIC constraints and re-solving the optimization problem. In all cases I continue to insist that the capacity constraints hold with equality, which I impose as a minimal notion of stability. I conclude that it is the ODIC constraints, rather than the stability constraints, that limit my ability to improve student welfare. One can impose stability on top of the ODIC constraint at a welfare cost of 0.004, which is the equivalent of moving 11 students to a school one rank higher in their preference list. Moreover,

²⁹Other common metrics include the percentage of students assigned their first choice, one of their top two choices, etc. The average rank has the benefit of including information about all of the assignment ranks in a single statistic.

³⁰The definition of ODIC (Definition 1) continues to apply given the restricted content of the verifiable traits.

³¹I compute the Gale-Shapley outcome in a limit economy using the techniques from Azevedo and Leshno [11] with uniform tie-breakers as per the algorithm used by BPS.

	Best Match	Best Stable Match	Best ODIC Match	Best Stable, ODIC Match
$R(x)$	1.572 [1.544, 1.608]	1.589 [1.560, 1.625]	1.748 [1.715, 1.784]	1.752 [1.721, 1.790]

TABLE 5. Welfare Analysis

the gap between the best match and the best stable match is 0.017, 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 0.176, a gap that is equivalent to moving 441 students to a school one rank higher in their preference list.

Abdulkadiroğlu et al. [2] uses data from BPS and the New York City school system to assess the cost of strategy-proofness.^{32,33} The authors compare the status quo outcome with the outcome that is generated by applying Erdil and Ergin’s [26] stable improvement cycles to the status quo outcome. Table 3 of Abdulkadiroğlu et al. [2] implies that moving from the Gale-Shapley outcome to a student-optimal stable match results in a change equivalent to moving 9.7 students to a school one rank higher in their preference list, and the authors conclude that the welfare costs of incentive compatibility are small.³⁴

When I repeat this exercise in my data, I find a welfare improvement equivalent to moving 18.2 students to a school one rank higher in their respective preference lists. The large gap between the welfare generated by the best stable match (Table 5) and the match found by the stable improvement cycles algorithm applied to my data is the result of a large set of Pareto optimal, stable matches. Since the best stable match I find and the outcome of the stable improvement cycles algorithm are both Pareto optimal, one can conclude that the best stable match improves the assignment of many students, but there is at least one student that is made worse off in the best stable match.

Finally, one can use the shadow prices of the constraints to evaluate how to allocate resources to maximize welfare. For example, the highest shadow prices for the capacity constraints are associated with the four programs offered by Snowden International. 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.

³²Che and Tercieux [20] conduct a similar exercise using data from the NYC school match.

³³With strict priorities, the student-optimal stable match is obtained by Gale-Shapley. Characterizations of the welfare ranking have been provided in a number of different ways in the literature. A recent contribution in terms of participation is provided by Alva and Manjunath [6].

³⁴The authors also 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. The authors conclude that there is a low cost of stability in the BPS school assignment system.

	Unconditional	Ethnicity	Zip Code	All Demographics
Best ODIC Match	1.748 [1.715, 1.784]	1.710 [1.677, 1.737]	1.689 [1.645, 1.713]	1.634 [1.587, 1.655]
Best Stable, ODIC Match	1.754 [1.721, 1.790]	1.717 [1.679, 1.745]	1.699 [1.654, 1.723]	1.646 [1.600, 1.671]
Worst Stable, ODIC Match	1.794 [1.758, 1.830]	1.825 [1.799, 1.872]	1.862 [1.842, 1.911]	1.924 [1.906, 1.976]
Worst ODIC Match	1.806 [1.772, 1.841]	1.836 [1.814, 1.884]	1.873 [1.858, 1.924]	1.945 [1.930, 1.998]

TABLE 6. Effect of ODIC Conditioning on $R(x)$

4.2. The Benefits of Conditioning on Demographics. Section 4.1 assumed the verifiable trait v only described the student’s priorities at each school. In this section, I include increasing amounts of demographic information (e.g., zip code) into the verifiable trait. Therefore, students with the same preference ranking and priorities can be offered different assignments if the students have different realizations of the demographic variable(s) included in v . For example, if the student’s zip code is included in v , then two students living in different zip codes can be offered different assignments even if their declared preferences and priorities are identical. One can interpret this as a relaxation of an implicit “equal-treatment-of-equals” fairness requirement since the set of “equals” for a given student (i.e., the number of students with the same realization of v) becomes smaller as v includes more demographic information. The relaxation of equal-treatment-of-equals implies that the set of ODIC constraints is relaxed. For example, two student-types with v realizations that include different zip codes do not impose ODIC constraints on each other. Since the ODIC constraints are the primary barrier to improving student welfare (see Table 4), it is perhaps not surprising that weakening these constraints would result in a significant welfare improvement. I also impose capacity and individual rationality constraints (equations 3.3-3.5) and stability constraints (equations 3.11-3.14). Note that although the student’s demographics affect her assignment, her preferences are also taken into account.

I consider four classes of mechanisms in Table 6. 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 or the zip code of the student’s home residence. The final column (*All Demo.*) considers mechanisms that conditions on the students’ ethnicity, zip code, and free school lunch status.

Table 6 presents the best and worst ODIC matches as well as the best and worst stable and ODIC matches for each form of conditioning. The welfare improvement from implementing the best stable, ODIC match that conditions on all the demographics (relative to the status quo) is

equivalent to moving 291 students (out of 2,513) to a school one rank higher in their preference list. In other words, a mechanism that conditions on the full set of demographic information finds a stable, ODIC match that recovers over 70% of the gap between the status quo assignment and the welfare possible if one ignores the ODIC conditions. The primary takeaway is that conditioning on the student demographics gives the mechanism significantly greater freedom to improve the assignment.

Now let us turn to a comparison of the different mechanisms in Table 6.³⁵ Moving from a mechanism that does not condition on any demographics to a mechanism that conditions on either ethnicity or zip code provides a welfare improvement equivalent to moving between 93 and 138 students (out of 2,513) to a school one rank higher in their preference list. A direct reading of Table 6 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. However, this ignores the fact that the outcomes across mechanisms are correlated, which means that these differences are significant. 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 variables. Similarly, since the final column describes the outcome of an optimal mechanism that conditions on more information than any other mechanism in Table 6, the optimal mechanism from this group must improve on all of the others I consider.

There is a fairness issue that arises when considering conditioning on demographic data. Making the school assignment depend on immutable and exogenous characteristics of a student (e.g., ethnicity) for reasons that have a tenuous intrinsic connection to the goal (e.g., welfare) violates equal-treatment-of-equals as noted at the beginning of this section.³⁶ 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. Unlike the cost of imposing the fairness criteria of stability, which is low (see Table 5), the welfare cost of equal-treatment-of-equals is high.

Certain fairness violations may be 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

³⁵Table 10 in Appendix E displays the number of students that are assigned to schools at a particular rank of their preference relation under the status-quo Gale-Shapley mechanism, the optimal mechanism that does not condition on demographics, and the optimal mechanism that conditions on all demographics.

³⁶In 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.

treats students from different parts of Boston differently. Since zip-code based assignment may be more acceptable, both on fairness and legal grounds, 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 7 and Figure 1b). Using zip code or geographic data as a proxy for other demographic traits might very well be legal. For example, the Chicago Public School system uses the demography of the census tract in which a student lives as an input to the school assignment system, although the formal system uses proxies for ethnicity (e.g., fraction of the population that speaks a language other than English) rather than ethnicity directly (Chicago Public Schools [21]).

5. DISTRIBUTIONAL GOALS

In this section, I 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 status-quo Gale-Shapley mechanism on the remaining objectives. Section 5.3 provides a graphical analysis through 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. Unless otherwise noted, the matches are allowed to condition on the full set of student demographics so that I can explore the set of stable and ODIC matches as completely as possible.

5.1. Metrics. School c is diverse if the demographics of the student body at school c closely resemble the demographics of the student population of BPS as a whole. My objective function is the average diversity across all of the overdemanded BPS schools.³⁷

Let \bar{p}_g denote the fraction of the BPS population in demographic group g , where g is one of the six ethnic groups listed 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 . $e(s)$ denotes the ethnic group that student-type s belongs to, and $\mathbb{1}\{e(s) = g\}$ is an indicator that takes the value 1 if $e(s) = g$ and 0 otherwise. With this notation in hand, I can write

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

I then define the diversity of school c , denoted $d_c(x)$, as the average across groups g of the squared difference between the program's demographics, $p_{g,c}(x)$, and BPS as a whole, \bar{p}_g , as

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

³⁷Enforcing stability restricts my ability to change the student bodies of underdemanded schools since any student can claim a seat at an underdemanded school in lieu of her current assignment. Therefore, I focus on overdemanded schools.

My metric of aggregate school diversity, $D(x)$, is the average of d_c over the school programs:

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

In my tables and figures I have, for expositional clarity, used the following average absolute deviation metric for diversity³⁸

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

D_{Abs} provides a linear scale for the diversity metric in-line with the welfare and neighborhood school metrics.

The metric I use for 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 under school assignment x . I use the notation $l(s) = 1$ to denote the event that students of type s are eligible for a free lunch, so I can write

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

$\bar{f} = 0.665$ 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

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

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

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

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

$$S_{Abs}(x) = \frac{1}{\sum_c q_c} \sum_c q_c \left\| f_c(x) - \bar{f} \right\|.$$

The final distributional goal I consider is the encouragement of neighborhood schools. Let $n_c(x)$ denote the fraction of school c 's students drawn from the school's walk-zone. My metric for the encouragement of neighborhood schools is the average of $n_c(x)$ across 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? I now compute the mechanisms that are optimal with respect to either student welfare, socioeconomic or ethnic diversity, or encouraging

³⁸While I would like to have used the absolute difference criterion in the optimization problem, it is computationally easier to solve a quadratic program.

Metric	Gale-Shapley Mechanism	Constrained Optimal Mechanism	Column 1 - Column 2
Welfare, $R(x)$	1.763 [1.730, 1.796]	1.650 [1.603, 1.673]	0.113 [0.113, 0.136]
Socioeconomic Diversity, $S_{Abs}(x)$	0.120 [0.104, 0.140]	0.008 [0, 0.035]	0.112 [0.124, 0.086]
Encouraging Neighborhood Schools, $N(x)$	0.221 [0.199, 0.241]	0.358 [0.323, 0.391]	-0.137 [-0.111, -0.1582]
Ethnic Diversity, $D_{Abs}(x)$	0.0994 [0.092, 0.113]	0.047 [0.037, 0.063]	0.052 [0.063, 0.043]
Ethnic Diversity Using Zip Code, $D_{Abs}(x)$	0.0994 [0.092, 0.113]	0.048 [0.038, 0.063]	0.051 [0.063, 0.043]

TABLE 7. Optimal Mechanisms For Each Metric

neighborhood schools while performing no worse than the status quo Gale-Shapley mechanism on the remaining goals. 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 four metrics. Table 7 presents the outcomes for each constrained optimal mechanism—note that each metric has a distinct constrained optimal mechanism. Recall that lower values represent an improvement for all of my metrics except for encouraging neighborhood schools.

What my analysis reveals is that it is possible to make a significant difference on any one of the metrics while continuing to do no worse than the Gale-Shapley mechanism on the remaining three. For example, Table 6 shows that the optimal outcome from a welfare maximizing, ODIC, and stable mechanism is an average rank of 1.646. Table 7 shows that I can achieve an average rank of 1.650 while doing no worse on the remaining goals. The gap of 0.004 is equivalent to moving 11.8 students to a school one rank higher in their preference list.

The gains along other desiderata are also sizable. In the status quo assignment 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 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.³⁹

³⁹The 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.

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 8. % of Students From Each Ethnic Group at Overdemanded Schools Under the Constrained Optimal Mechanism

Finally, consider the encouragement of neighborhood schools. Even after imposing the BPS walk-zone priorities, the Gale-Shapley mechanism only fills 22% of the seats at these schools with students from the walk-zones.⁴⁰ The optimal constrained mechanism manages to fill 36% of the seats with students from the walk-zones without sacrificing welfare or diversity.

Table 8 presents statistics on the demographics at the overdemanded schools under the constrained optimal mechanism for maximizing ethnic diversity and under the status quo Gale-Shapley mechanism. The first set of columns refer to the constrained optimal mechanism, while the second set of columns refers to the Gale-Shapley mechanism.^{41,42}

The constrained optimal mechanism for maximizing ethnic diversity does so through two channels. First, the mechanism minimizes the **gap in means**, which is the difference between the overall demographics of the overdemanded schools and the student population of BPS. For example, the 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. Thus, the gap in means has dropped from 11%(= 48.3% – 37.3%) to 5.9%(= 43.2% – 37.3%).⁴³

⁴⁰The 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.

⁴¹I computed the numbers for Table 8 using the continuum analog of the discrete Gale-Shapley mechanism.

⁴²Tables 11 and 12 in the appendix provides school-by-school demographic statistics under the status-quo and constrained optimal ethnic diversity mechanism.

⁴³Since fewer African-American students are enrolled into the overdemanded schools as the gap in means closes, these students are (presumably) worse off. However, there is no Pareto ranking since the students moved into the school to close the gap have improved assignments.

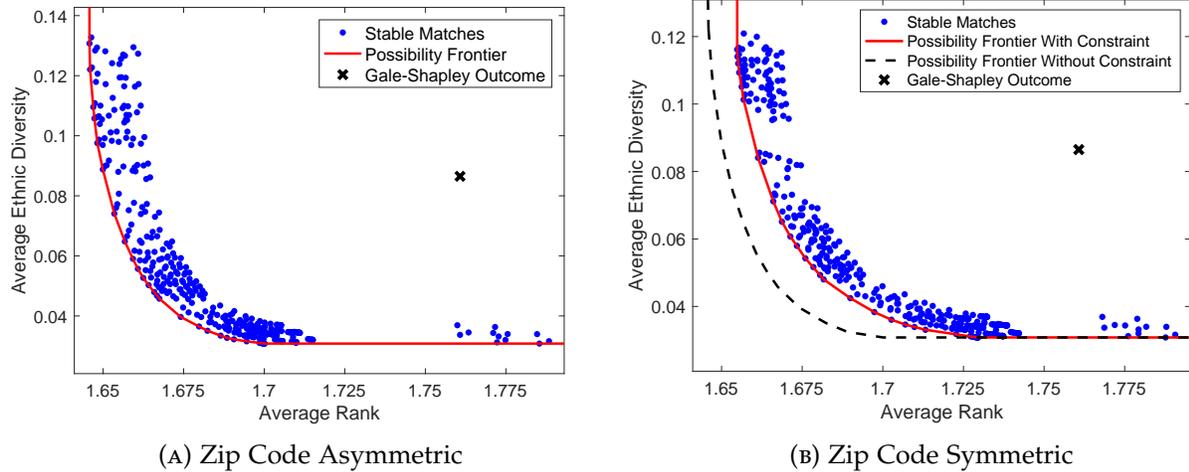


FIGURE 1. Diversity vs Welfare

The gap in means for the other three large ethnic groups are reduced by a similar fraction. The mechanism also limits the **across school gap**, which is the difference in demographics amongst the overdemanded schools. For example, the solution to the optimization problem results in schools with student bodies that range from 26.9% to 37.2% hispanic. This range is three times larger (20.4% to 51.7%) in the status quo assignment. The fraction of white students ranges from 7.2% to 18.3% under the constrained optimal mechanism, which is a 50% decrease of the range generated by the Gale-Shapley mechanism. A similar effect is found for the Asian and African-American populations at these schools.⁴⁴

One might be concerned that achieving these diversity gains requires conditioning the assignment on the students' ethnicities, which could violate the U.S. District court ruling in *McLaughlin v. Boston School Committee* (1996). In order to assuage this concern, the row labeled "Ethnic Diversity Using Zip Code" in Table 7 maximizes racial diversity under the requirement that the assignment mechanism treat all students from the same zip code symmetrically. Requiring symmetric treatment of students within a zip code reduces my ability to achieve racially diverse schools, but the reduction is small. Using zip code as a proxy for ethnic background is an effective, if indirect, tool for encouraging diversity.

5.3. Diversity Possibility Frontiers. The goal of this section is to explore the trade-offs between ethnic diversity and student welfare by creating a possibility frontier that describes the realizations of these goals that are compatible with stability and ODIC. To this end, I minimize the 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 $\alpha \in [0, 1]$ allows me to trace out the welfare-diversity possibility frontier.

⁴⁴The across school gap can be closed by moving students between the overdemanded schools (unlike the gap in means), which means it is not obvious that closing this gap necessarily makes any students worse off.

Figure 1a presents the results of my analysis with $D_{Abs}(x)$ 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.⁴⁵ Matches with higher student welfare are 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 takeaways from my analysis. First, almost all of the gains in terms of diversity can be reaped at a small welfare cost.⁴⁶ Second, the Gale-Shapley outcome is inside the frontier and lacking in both welfare and diversity.

Figure 1b replicates the diversity analysis under the requirement that the assignment mechanism treat all students from the same zip code symmetrically, which implicitly means zip code serves as a proxy for ethnic background. I have plotted the stable matches for the zip-code-symmetric solution as well as the welfare-diversity possibility frontier for the less constrained model (i.e., as shown in Figure 1a). I am still able to improve on the Gale-Shapley outcome on both metrics even given the zip-code-symmetry constraint.

An analysis of the tension between ethnic and socioeconomic diversity is contained in Appendix G. I find that there is essentially no tension between these goals.

5.4. Encouraging Neighborhood Schools Possibility Frontiers. The goal of the walk-zone priority system was to encourage students to enroll in their neighborhood schools. 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.⁴⁷

Figure 2a reports the result of solving my problem with the walk-zone priority system in place for various weightings of the diversity and walk-zone metrics, $D(x)$ and $N(x)$ respectively. Figure 2a uses the absolute deviation metric, $D_{Abs}(x)$, 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. The outlying clusters of solutions in the lower-left and upper-right of the figure, one denoting solutions 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

⁴⁵Recall that to define the stability constraints one must decide which overdemanding schools admit students from outside the walk-zone into the walk-zone seats. For each of α there are 24 sets of stability constraints to consider.

⁴⁶Recall that a 0.01 decrease in $R(x)$ is equivalent to moving 25.1 students to a school one rank higher in their preference list.

⁴⁷In Appendix H 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.

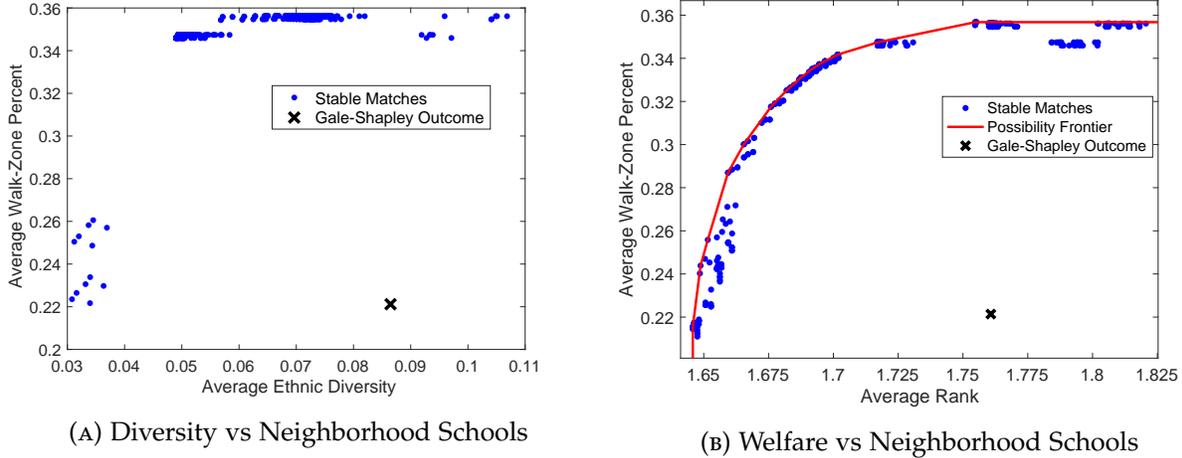


FIGURE 2. Diversity vs Welfare

assignment that both encourages diversity and creates strong neighborhood schools despite the racial segregation of many of Boston’s zip codes. 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.

Figure 2b describes the tension between encouraging neighborhood schools and student welfare. 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 desiderata.

6. CONCLUSION

In this paper, I introduce a novel analysis tool for school choice settings based on constrained optimization problems that assign a continuum of students to a continuum of school program seats. By assuming that I have a continuum of agents on each side of the market, I am able to describe the stability and ODIC constraints in a tractable form. The constrained optimization approach allows the market designer to encode his or her goals into an objective function, and the solution to the resulting constrained optimization problem achieves the global optimum subject to the stability and ODIC constraints. Two major advantages of this approach are that I need not make parametric assumptions on student preferences and the objective function can include distributional goals such as school diversity alongside goals such as student welfare.

In the context of the BPS high school match, I find that the welfare gap between the status quo outcome and the best stable, ODIC match that does not condition on demographics is equivalent to moving 26 students to a school one rank higher in their preference list. 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.

I then maximize student body diversity with respect to ethnicity and socioeconomic status. I show that there are stable matches that yield significantly more diverse outcomes than the Gale-Shapley algorithm while at the same time achieving the same level of student welfare. The increased diversity makes the demographics of the student bodies at the schools more similar to the broader BPS population and shrinks the demographic differences between the schools. Moreover, the tension between student welfare and diversity appears small.

For my third analysis, I study the extent to which one can encourage neighborhood schools while at the same time increasing student welfare or school diversity. I again find that one can do significantly better than the Gale-Shapley outcome on all metrics simultaneously. 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 increases the fraction of the students drawn from the schools' walk-zones and achieve high levels of diversity.

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 a 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 incorporating distributional objectives in a school choice problem. Second, I find that one can make significant improvements relative to the Gale-Shapley outcome with respect to welfare and my distributional goals without sacrificing either stability or incentive compatibility. Third, the application of my framework 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 framework is tractable for many moderately large school systems such as BPS.

One avenue for future work is to explore other goals and constraints of school systems. For example, busing costs are a significant budgetary concern for BPS (Ashlagi and Shi [7]). One might also hope to optimize educational attainment, although this would require a detailed model of the educational attainment production function that is difficult to credibly identify. I have focused on student welfare, school diversity, and encouraging neighborhood schools due to their prominence as public policy objectives and because of the potential tensions between these goals. Although there are obviously other potential goals of interest, I leave these to future work.

REFERENCES

- [1] Abdulkadiroğlu, A.; N. Agarwal; and P. Pathak (2017) "The Welfare Effects of Coordinated Assignment: Evidence from the New York City High School Match," *American Economic Review*, 107, pp. 3635 - 3689.
- [2] 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.
- [3] Abdulkadiroğlu, A. and T. Sönmez (2003) "School Choice: A mechanism design approach," *American Economic Review*, 93 (3), pp. 1954 - 1978.
- [4] N. Agrawal and P. Somaini (2014) "Demand Analysis Using Strategic Reports: An Application to a School Choice Mechanism," *mimeo*.
- [5] Ajayi, K. (2012) "School Choice and Educational Mobility," *mimeo*.
- [6] Alva, S. and V. Manjunath (2017) "Strategy-proof Pareto-improvement under voluntary participation," *mimeo*.
- [7] Ashlagi, I. and P. Shi (2014) "Optimal Allocation without Money: an Engineering Approach," *Management Science*, 62 (4), pp. 1078 - 1097.
- [8] Ashlagi, I. and P. Shi (2014) "Improving Community Cohesion in School Choice via Correlated-Lottery Implementation," *Operations Research*, 62 (6), pp. 1247-1268.
- [9] Azevedo, E. and E. Budish (2018) "Strategy-proofness in the Large," *Review of Economic Studies*, forthcoming.
- [10] Azevedo, E. and J. Hatfield (2012) "Complementarity and Multidimensional Heterogeneity in Matching Markets," *mimeo*.
- [11] Azevedo, E. and J. Leshno (Forthcoming) "A Supply and Demand Framework for Two-Sided Matching Markets," *Journal of Political Economy*.
- [12] Azevedo, E. and J. Hatfield (2012) "Complementarity and Multidimensional Heterogeneity in Matching Markets," *mimeo*.
- [13] Baïou, M. and M. Balinski (2000) "The stable admissions polytope," *Mathematical Programming*, 87, pp. 427 - 439.
- [14] Bodoh-Creed, A. (2013) "Efficiency and information aggregation in large uniform-price auctions," *Journal of Economic Theory*, 148 (6), pp. 2436 - 2466.
- [15] Bonnans, J. and A. Shapiro (2000) *Perturbation analysis of optimization problems*. New York: Springer-Verlag.
- [16] 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.
- [17] Burgess, S.; E. Greaves; A. Vignoles; and D. Wilson (2014) "What Parents Want: School Preferences and School Choice," *The Economic Journal*, 125, pp. 1262-1289.
- [18] Calsamiglia, C.; C. Fu; and M. Güell, "Structural Estimation of a Model of School Choices: the Boston Mechanism vs. Its Alternatives," *mimeo*.
- [19] Che, Y-K; J. Kim; and F. Kojima (2014) "Stable Matching in Large Economies," *mimeo*.
- [20] Che, Y-K and O. Tercieux (2017) "Efficiency and Stability in Large Matching Markets," *mimeo*.
- [21] Chicago Public Schools (2017) *Admissions Policy for Magnet, Selective Enrollment and Other Options for Knowledge Schools and Programs*. Downloaded from <http://policy.cps.edu/download.aspx?ID=82> on 4 December 2017.
- [22] Curtis, C. (1984) *Linear Algebra*, Springer-Verlag: New York.
- [23] Dur, U.; S. Kominers; P. Pathak; and T. Sönmez (2014) "The Demise of Walk Zones in Boston: Priorities vs. Precedence in School Choice," *mimeo*.
- [24] Echenique, F and M.B. Yenmez (2014) "How to Control Controlled School Choice," *The American Economic Review*, 105 (8), pp. 2679 - 2694.
- [25] 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, 648-683.

- [26] 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.
- [27] Erdil, A. and T. Kumano (2013), "Prioritizing diversity in school choice," *mimeo*.
- [28] Fack, G.; J. Grenet; and Y. He "Beyond Truth-Telling: Preference Estimation with Centralized School Choice," *mimeo*.
- [29] Featherstone, C. (2006) "What's the Matter with Tie-Breaking? Improving Efficiency in School Choice," *American Economic Review*, 98, pp. 669-689.
- [30] Fragiadakis, D. and P. Troyan (2017) "Improving matching under hard distributional constraints," *Theoretical Economics*, 12, pp. 863 - 908.
- [31] Gale, D. and L.S. Shapley (1962) "College Admissions and the Stability of Marriage," *The American Mathematical Monthly*, 69, pp. 9-15.
- [32] Hafalir, I.; M. Yenmez; and M. Yildirim (2013) "Effective affirmative action in school choice," *Theoretical Economics*, 8, pp. 325-363.
- [33] Hastings, J.; T. Kane; and D. Staiger (2009) "Heterogeneous Preferences and the Efficacy of Public School Choice," *mimeo*.
- [34] Y. He "Gaming the Boston School Choice Mechanism in Beijing", *mimeo*.
- [35] Immorlica, N. and M. Mahdian (2005) "Marriage, Honesty, and Stability," *SODA 2005*, pp. 53-62.
- [36] Kamada, Y. and F. Kojima (2015), "Efficient matching under distributional concerns: Theory and application," *American Economic Review*, 105, 67-99.
- [37] A. Kapor; C. Neilson; S. Zimmerman (2017) "Heterogeneous Beliefs and School Choice Mechanisms," *mimeo*.
- [38] Kesten, O. (2010) "School Choice with Consent," *The Quarterly Journal of Economics*, 125 (3), pp. 1297 - 1348.
- [39] Kesten, O. and U. Unver (2015) "A theory of school-choice lotteries," *Theoretical Economics*, 10, pp. 543 - 595.
- [40] Kojima, F. and M. Manea (2006) "Strategy-Proofness of the Probabilistic Serial Mechanism in Large Random Assignment Problems," *mimeo*.
- [41] Kojima, F. and P. Pathak (2009) "Incentives and Stability in Large Two Sided Matching Markets," *American Economic Review*, 99 (3), pp. 608 - 627.
- [42] Kominers, S. and T. Sönmez (2016) "Matching with Slot-Specific Priorities: Theory," *Theoretical Economics*, 11 (2), pp. 683-710.
- [43] Lee, S. (2012) "Incentive Compatibility of Large Centralized Matching Markets," *mimeo*.
- [44] Q. Liu and M. Pycia (2016) "Ordinal Efficiency, Fairness, and Incentives in Large Markets," *mimeo*.
- [45] Manea, M. (2009) "Asymptotic ordinal inefficiency of random serial dictatorship," *Theoretical Economics*, 4, pp. 165 -197.
- [46] Pathak, P. and P. Shi (2017) "How Well Do Structural Demand Models Work? Counterfactual Predictions in School Choice," *mimeo*.
- [47] Robinson, S. (1977) "A Characterization of Stability in Linear Programming," *Operations Research*, 25 (3), pp. 435-447.
- [48] 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.
- [49] Roth, A.; U. Rothblum; and J. Vande Vate (1993) "Stable Matchings, Optimal Assignments, and Linear Programming," *Mathematics of Operations Research*, 18, pp. 803-828.
- [50] Roth, A. and M. Sotomayor (1990) *Two Sided Matching: a Study in Game-Theoretic Modeling and Analysis*, Cambridge University Press: Cambridge, England.
- [51] Shi, P. (2016) "Assortment Planning in School Choice," *mimeo*.

APPENDIX A. FOR ONLINE PUBLICATION: PROOFS

The first subsection contains most of the results from the body of the paper. Due to their complexity, I reserve a second subsection for the proofs of the continuity result (Theorem 1).

A.1. Results other than Theorem 1.

Proposition 1. *Equations 3.3, 3.4, 3.6, and 3.5 are necessary and sufficient for a match $x(\circ, \circ; \pi^S)$ to be stable.*

Proof. Equations 3.3 and 3.4 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 3.5 holds. Equation 3.6 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 school program 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 school program capacities take irrational values.

Lemma 2. *When the school program 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 $m\pi^S(s)$ ($m q_c$) 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 school program 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 $\{x_n(\circ, \circ; \pi^S)\}_{n=1}^\infty$, where $x_n(\circ, \circ; \pi^S)$ is a feasible and stable match in the $(\pi_n^S, (q_C^n)_{c \in \mathcal{C}})$ economy. Recall that $x(\circ, \circ; \pi^S) \in [0, 1]^{(|\mathcal{C}|+1)(|\mathcal{S}|+1)}$ and note that this space is compact. From the compactness of the space, there exists a convergent subsequence of $\{x_n(\circ, \circ; \pi^S)\}_{n=N}^\infty$ with a limit $x_\infty(\circ, \circ; \pi^S)$. Since the feasibility and stability equations are

equalities, $x_\infty(\circ, \circ; \pi^S)$ must satisfy these equations for the $(\pi^S, (q_C)_{C \in \mathcal{C}})$ economy. Therefore $x_\infty(\circ, \circ; \pi^S)$ is a feasible and stable match for $(\pi^S, (q_C)_{C \in \mathcal{C}})$. \square

Lemma 1. *Equation 3.2 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 school programs to which c is strictly preferred by s . Therefore

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

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

Theorem 2. *If Assumption 1 holds, then for any $\varepsilon > 0$ there exists $N^*, \delta > 0$ such that if $N > N^*$ and $\|\pi^{S,N} - \pi^S\| < \delta$, then $x(c, s; \pi^{S,N})$ satisfies ε -ODIC under measure $\pi^{S,N}$.*

Proof. Given the continuity of $x(c, s; \circ)$ provided by Assumption 1, our result follows immediately. \square

Theorem 3. *Let Assumptions 1 and 2 hold. If $V_N(\pi^{S,N}) = F(x_N(c, s; \pi^{S,N}))$ and $V(\pi^S) = F(x(c, s; \pi^S))$, then $\limsup_{N \rightarrow \infty} V_N(\pi^{S,N}) = \liminf_{N \rightarrow \infty} V_N(\pi^{S,N}) = V(\pi^S)$. Moreover, if $x(c, s; \pi^S)$ is the unique maximizer, it must be that $x_N(c, s; \pi^{S,N}) \rightarrow x(c, s; \pi^S)$.*

Proof. First note that since $x_\infty(c, s; \pi^S)$ is continuous (Assumption 2), $x_\infty(c, s; \pi^S)$ satisfies the continuum model ODIC constraints. Since the stability constraints do not depend on the size of the market, these are also satisfied by $x_\infty(c, s; \pi^S)$. Therefore $x_\infty(c, s; \pi^S) \in \text{SIC}((q_C)_{C \in \mathcal{C}}, \pi^S)$, which implies $\lim_{N \rightarrow \infty} V_N(\pi^{S,N}) \leq V(\pi^S)$.

From Theorem 2 and Assumption 2 (so $\pi^{S,N} \rightarrow \pi^S$), there exists a sequence $\{\varepsilon_N\}_{N=1}^\infty$ where $\varepsilon_N > 0$ and $\varepsilon_N \rightarrow 0$ such that $x(c, s; \pi^{S,N})$ is an ε -ODIC under measure $\pi^{S,N}$ for large enough N . Therefore, $x(c, s; \pi^{S,N}) \in \text{SIC}_{N, \varepsilon_N}$ and so $F(x(c, s; \pi^{S,N})) \geq V_N(\pi^{S,N})$. Since F is continuous, as $\pi^{S,N} \rightarrow \pi^S$ we have $F(x(c, s; \pi^{S,N})) \rightarrow F(x(c, s; \pi^S)) = V(\pi^S) \geq V_N(\pi^{S,N})$. Therefore, $\lim_{N \rightarrow \infty} V_N(\pi^{S,N}) = V(\pi^S)$. \square

A.2. Proof of Theorem 1. 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

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

I now define a matrix A that represents the capacity and stability constraints of my problem. There must be one row for each school program and student-type, for total of $|\mathcal{C}| + |\mathcal{S}|$ rows. The matrix must have one column for each (school program, 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 school program. 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 also denote a match in column vector notation $\mathbf{x} \in \mathbb{R}^{(|\mathcal{C}|+1)(|\mathcal{S}|+1)}$ where $x_{(c,s)} = x(c, s; \pi^S)$ and the rows of \mathbf{x} have the same order as the columns of A . For $\mathbf{x} \geq 0$, a feasible and stable match solves

$$(A.1) \quad A\mathbf{x} = \mathbf{q}$$

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

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

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.

Denote the set of solutions to equation A.1 as $\mathbf{X}(\mathbf{q}; \underline{x}_C, \underline{x}_S)$. Since \mathbf{q} is a function of $(q_c)_{c \in \mathcal{C}}$ and π^S , \mathbf{X} is implicitly a function of the underlying distributions of student-types and school program 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 $\mathbf{X}(\mathbf{q}; \underline{x}_C, \underline{x}_S)$ is a convex valued, nonempty, continuous correspondence with respect to \mathbf{q} .*

Proof. Claim 1 follows from theorem 8.8 of Curtis [22], which states that equation A.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 [22] implies the existence of a solution. Since the system is linear, the set of solutions must be convex. If A does not have full rank, then it the rank of A is strictly less than the rank of $[A \ \mathbf{q}]$ (i.e., no solution exists) for a topologically generic set of \mathbf{q}

Now I prove continuity of the correspondence of solutions. First note that upper hemicontinuity follows from my definition of equation A.1 as a system of linear equalities. What remains

is to show lower hemicontinuity, which requires that if $Ax = \mathbf{q}$, then for any sequence $(\mathbf{q}_i)_{i=1}^{\infty}$ where $\mathbf{q}_i \rightarrow \mathbf{q}$ I can choose \mathbf{x}_i such that $Ax_i = \mathbf{q}_i$ and $\mathbf{x}_i \rightarrow \mathbf{x}$. 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 $\|\mathbf{x}_i - \mathbf{x}\| < \gamma$ where $Ax_i = \mathbf{q}_i$.

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 A.1. Now I am faced with a perturbation $\varepsilon \in \mathbb{R}^{|\mathcal{C}|+|\mathcal{S}|}$ of \mathbf{q} and wish to solve

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

which is equivalent to solving

$$(A.3) \quad A\delta = \varepsilon$$

I solve the equivalent problem

$$Bz = \varepsilon$$

which admits a solution since B has full rank. The vector \mathbf{z} can be converted into a vector δ that solves equation A.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

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

Therefore, there exists \mathbf{y}_i such that

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

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

I am still one step removed from making claims regarding sets of stable, feasible matches because $\mathbf{X}(\mathbf{q}; \underline{x}_C, \underline{x}_S)$ may include solutions that do not satisfy the non-negativity constraints of the match. Let $\mathcal{S}(\mathbf{q}; \underline{x}_C, \underline{x}_S)$ 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 A.1 does not constrain $x(c, s; \pi^S)$ when (c, s) is a blocking pair with respect to \underline{x}_C and \underline{x}_S . To this end let $\mathbb{R}_{SC} \subset \mathbb{R}^{(|\mathcal{C}|+1)(|\mathcal{S}|+1)}$ be a product set whose i^{th} dimension is equal to \mathbb{R}_+ if the i^{th} column of A denotes a (c, s) where (1) $\pi^S(s) > 0$ and (2) (c, s) 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.⁴⁸ Given this definition

$$(A.4) \quad \mathbf{S}(\mathbf{q}; \underline{x}_C, \underline{x}_S) = \mathbf{X}(\mathbf{q}; \underline{x}_C, \underline{x}_S) \cap \mathbb{R}_{SC}$$

which insures that $\mathbf{S}(\mathbf{q}; \underline{x}_C, \underline{x}_S)$ contains only the solutions of equation A.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} where $q_c > 0$ for all c , 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. 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 $\mathbf{X}(\mathbf{q}; \underline{x}_C, \underline{x}_S)$ 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 $\mathbf{X}(\mathbf{q}; \underline{x}_C, \underline{x}_S)$ lies in the open set $\mathbb{R}^{(|C|+1)(|S|+1)} - \mathbb{R}_{SC}$. Since I know that $\mathbf{X}(\mathbf{q}; \underline{x}_C, \underline{x}_S)$ is continuous in \mathbf{q} , for any \mathbf{q}' sufficiently close to \mathbf{q} , I must have that $\mathbf{X}(\mathbf{q}'; \underline{x}_C, \underline{x}_S) \subset \mathbb{R}^{(|C|+1)(|S|+1)} - \mathbb{R}_{SC}$. This last fact implies $\mathbf{S}(\mathbf{q}'; \underline{x}_C, \underline{x}_S)$ is empty.

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 $\mathbf{x} \in \mathbf{X}(\mathbf{q}; \underline{x}_C, \underline{x}_S)$ are on the non-trivial interior of \mathbb{R}_{SC} in the sense that, for each such \mathbf{x} there exists $\mathbf{x}' \in \mathbb{R}_{SC}$ and a dimension j such that $x_j = 0 < x'_j$. Define $\delta \in \mathbb{R}_{SC}$ where $\delta_{(c,s)} = \gamma, \gamma > 0$ if that dimension of \mathbb{R}_{SC} is not equal to $\{0\}$. Let $\varepsilon = A\delta < \mathbf{0}$, and note that as $\gamma \rightarrow 0$ I have $\varepsilon \rightarrow \mathbf{0}$.

Consider the following system:

$$(A.5) \quad A\mathbf{z} = A(\mathbf{x} - \delta) = \mathbf{q} - \varepsilon$$

From the construction of A , for γ sufficiently small $\mathbf{q} - \varepsilon$ is a valid capacity vector, so the set of solutions \mathbf{z} to this system that lie in \mathbb{R}_{SC} define $\mathbf{S}(\mathbf{q} - \varepsilon; \underline{x}_C, \underline{x}_S)$.⁴⁹ But also note that \mathbf{z} solves equation A.5 if-and-only-if there is some $\mathbf{x} = \mathbf{z} + \delta \in \mathbf{Y}(\mathbf{q}; \underline{x}_C, \underline{x}_S)$ such that \mathbf{x} solves equation A.1. Since \mathbf{x} is on the non-trivial exterior of \mathbb{R}_{SC} and I chose $\delta \in \mathbb{R}_{SC}$ to be strictly positive in all dimensions where possible, it must be the case that $\mathbf{z} = \mathbf{x} - \delta \notin \mathbb{R}_{SC}$. Since this holds for all

⁴⁸In other words, the i^{th} dimension of \mathbb{R}_{SC} is $\{0\}$ if $\mathbf{x} \in \mathbf{X}(\mathbf{q}; \underline{x}_C, \underline{x}_S)$ implies $x_{(c,s)} = 0$.

⁴⁹One might have worried that $\mathbf{q} - \varepsilon$ for elements corresponding to s where $\pi^S(s) = 0$. Our construction of δ insures that this is not the case because $\delta_{(c,s)} = 0$ for these student-types.

solutions to equation A.5, it must be the case that

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

In other words, I can choose \mathbf{q}' arbitrarily close to \mathbf{q} such that claim 1 holds, implying the set of \mathbf{q} where claim (1) or claim (2) holds is dense. \square

Lemma 4 leads almost immediately to Theorem 2.

Theorem 2. $S((q_c)_{c \in \mathcal{C}}, \pi^S)$ is continuous in $((q_c)_{c \in \mathcal{C}}, \pi^S)$ for topologically generic choices of $((q_c)_{c \in \mathcal{C}}, \pi^S)$ where $q_c > 0$ for all c . $SIC((q_c)_{c \in \mathcal{C}}, \pi^S)$ is continuous in $((q_c)_{c \in \mathcal{C}}, \pi^S)$ for topologically generic choices of $((q_c)_{c \in \mathcal{C}}, \pi^S)$ where $q_c > 0$ for all c and π^S has full support.

Proof. First note that $S((q_c)_{c \in \mathcal{C}}, \pi^S) = \bigcup_{(\underline{x}_C, \underline{x}_S)} S(\mathbf{q}; \underline{x}_C, \underline{x}_S)$ 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 the first claim about $S((q_c)_{c \in \mathcal{C}}, \pi^S)$.

To prove my claim about $SIC((q_c)_{c \in \mathcal{C}}, \pi^S)$, first note that if we express the match as $\tilde{x}(c, s; \pi^S) = x(c, s; \pi^S) / \pi^S$, then the ODIC constraints can be described as a matrix of integers that is independent of \mathbf{q} . This implies that the set of \tilde{x} that satisfy the ODIC constraints is fixed. Since the mapping from \tilde{x} to x is continuous in π^S if π^S has full support, then the set of x that satisfy ODIC is continuous in π^S . Therefore, the intersection of $S((q_c)_{c \in \mathcal{C}}, \pi^S)$ and the set of x that satisfy ODIC, which is $SIC((q_c)_{c \in \mathcal{C}}, \pi^S)$, is continuous in $((q_c)_{c \in \mathcal{C}}, \pi^S)$ for topologically generic choices of $((q_c)_{c \in \mathcal{C}}, \pi^S)$ where $q_c > 0$ for all c and π^S has full support. \square

APPENDIX B. DESCRIPTION OF THE GALE-SHAPLEY ALGORITHM

The algorithm uses two inputs. First, the algorithm requires an ordinal preference ranking of available school programs from each student. Second, the school programs must provide a strict priority ordering of the students. Since BPS used coarse priority classes (e.g., all students within a walk-zone without a sibling at the school are in a single class), ties were broken by a single random tie breaker that was assigned to each student. The resulting strict priorities are used by the algorithm. Given the set of student preference profiles, the student-proposing 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.

The Gale-Shapley mechanism has the following timing:

- (1) Each student submits a rank ordered list of up to 10 school programs to the mechanism and begins the algorithm without an assignment.
- (2) Each student without an assignment applies to her most favored school program that has not yet rejected her application.
- (3) Each school program tentatively admits students from highest to lowest priority out of the pool of applicants at the current step and the set of students tentatively assigned to that program previously. Once the program has tentatively admitted a number of students equal to the program's capacity, all students that remain in the pool are rejected.
- (4) If all students have a school assignment, then the algorithm halts and the assignments are finalized. If there is a student without an assignment, return to step 2.

It is possible for a student to submit a preference ranking consisting entirely of overdemanding 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, 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.

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.⁵⁰

⁵⁰The 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.

APPENDIX C. FOR ONLINE PUBLICATION: IMPLEMENTATION THEORY FOR OUTCOMES OF THE
CONSTRAINED OPTIMIZATION PROBLEMS

A question left unaddressed in Section 3.4 is whether the match can be implemented. The output of the constrained optimization problem is a probability that each student of type s is assigned to each school program 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 $|\mathcal{C}|$ school programs (the first $|\mathcal{C}|$ columns) or is assigned to her most preferred underdemanded school program (the final column). To define the assignment associated with a match x , all of the cells corresponding to students of type s and school program 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 school programs must be a *pure assignment* that places each student in a single school program. In other words, the final assignment must contain only 0 or 1 values - a student is either matched to a school program (1) or not matched with that school program (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. [16], 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. [16]), I use Lemma 1 to rewrite the constraints in a way that fits within the framework of Budish et al. [16].

Proposition 3. *Any assignment satisfying equations 3.3 through 3.6 can be implemented.*

Proof. I rely on theorem 1 of Budish et al. [16]. From this result, it suffices to argue that the constraints defining my match are a bihierarchy. The capacity constraints on students and school programs 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. [16] 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:

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

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

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 school program. So equations 3.3 and C.1 form a hierarchy, while equations 3.4 and C.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. [16] implies \mathbf{X} can be implemented. \square

One might imagine directly soliciting preferences from the agents, solving the constrained optimization problem of interest to generate a school assignment, and then implement the assignment as per Proposition 3. There are several issues with employing a mechanism based directly on software that solves a constrained optimization problem. First, while it is possible to assess the size of the ε term in an ε -ordinal dominance equilibrium, it is reasonable to worry how the size of the ε relates to the student's perception of the incentive. This latter question can only be answered by estimating a structural model of student utility, which carries all of the usual worries about parametric assumptions, etc. In addition, one needs to be careful how one handles "rare" types in the finite setting since Theorem 1 only insures continuity when the support of π^S is fixed. A variety of ad hoc solutions exist, but further work will be needed to flesh out the theoretical and practical issues attached to these options.

The second issue 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. Since it is difficult even for academic economists to understand the exact algorithm that generates a solution to a convex program, reoptimizing with respect to $\pi_E^{S,N}$ each year would not satisfy validation or procedural transparency. For this reason alone, one would be right to worry that the mechanism might not generate truthful preference declarations in practice.

APPENDIX D. FOR ONLINE PUBLICATION: IMPLEMENTING THE ODIC CONSTRAINTS

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 $x(c, s'; \pi^S)$, $s' = (v', \succeq_{s'}) \notin \mathcal{S}_{\exists}$, 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' . Let $\tilde{x}(c, s; \pi^S) = x(c, s; \pi^S)$ if $c \succeq_s \emptyset$ and $\tilde{x}(c, s; \pi^S) = 0$ otherwise. The definition of $\tilde{x}(c, s; \pi^S)$ is equivalent to giving the student the assignment $x(c, s; \pi^S)$ and allowing her to transfer to an underdemanded school if she is assigned to a school program c where $\emptyset \succeq_s c$. If $\mathcal{S}_{\exists}(v')$ is nonempty, then

$$(D.1) \quad x(c, s'; \pi^S) = \arg \max_s \left\{ \tilde{x}(c, s; \pi^S) : s \in \mathcal{S}_{\exists}(v') \right\}$$

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 D.1 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' if she reports her type truthfully. 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 $x(c, s'; \pi^S)$ for $s' = (v', \succeq)$ if $\mathcal{S}_{\exists}(v')$ is empty. Let $\mathcal{C}(s')$ denote the union of the set of underdemanded school programs; any school programs to which s' is insured a seat; and any school programs at which s' has walk-zone priority and where the school program 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 trait component of s' . This means that s'' can only be hurt by nontruthfully claiming to have type s' .

APPENDIX E. FOR ONLINE PUBLICATION: ADDITIONAL TABLES

School Name	Probability of Enrollment in an Open Seat
Brighton High	0.591
Excel High	0.385
Lyon	0.222
Snowden International, Chinese	0.189
Snowden International, French	0.070
Snowden International, Japanese	0.142
Snowden International, Spanish	0.071
Another Course College	0.107
Urban Science Academy	0.886
Tech Boston Academy	0.128

TABLE 9. Probability of Enrollment in an Open Seat Under the Status Quo Gale-Shapley Algorithm

\succsim Rank	Gale-Shapley Mechanism	No Demographic Conditioning	Condition on All Demographics	Column 2 - Column 1	Column 3 - Column 1
1 st	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]
2 nd	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]
3 rd	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]
4 th	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]
5 th	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]
6 th	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]
7 th	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]
8 th	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]
9 th	0.1 [0, 0.4]	0.1 [0, 0.4]	0 [0, 0]	0.01 [0, 0.03]	-0.11 [-0.4, 0]

TABLE 10. Number of Students Receiving Assignments at Each Rank in their Preference Lists Under Different Mechanisms

The mechanism that conditions on all of the student demographics generates 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.

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 11. Enrollment Under the Gale-Shapley Mechanism in Terms of the Number of Students

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 12. Enrollment Under the Constrained Optimal Mechanism for Diversity 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 13. Demographics at Overdemanded Schools In Convex Combination Solution in Terms of the Number of Students

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 14. Empty Seats Under the Menus-And-Reserves Mechanism

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 15. Demographics at Overdemanded Schools Under Iterative Menus-and-Reserves Mechanism in Number of Students

Metric	Optimum of Convex Combination of Metrics	Gale-Shapley
Welfare , $R(x)$	1.697 [1.659, 1.725]	1.763 [1.730, 1.796]
Ethnic Diversity , $D_{Abs}(x)$	0.063 [0.052, 0.075]	0.0994 [0.092, 0.113]
Encouraging Neighborhood Schools , $N(x)$	0.319 [0.274, 0.342]	0.221 [0.199, 0.241]

TABLE 16. Maximizing a Convex Combination of the Metrics

APPENDIX F. FOR ONLINE PUBLICATION: MAXIMIZING A CONVEX COMBINATION OF WELFARE,
DIVERSITY, AND ENCOURAGING NEIGHBORHOOD SCHOOLS

Now I would like to consider how well one can do on all of my metrics at once. I assess this by choosing an arbitrary convex combination of the student welfare, ethnic diversity, and neighborhood school 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 to decide. 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 desiderata relative to the status quo.

The values across the three metrics of interest as well as bootstrap standard errors are presented in Table 16. There is some tension between the metrics, which is the reason why I cannot quite achieve the performance of the constrained optimal mechanisms described in Table 7. That being said, the solution to the convex combination generates values much closer to those achieved by the constrained optimal mechanisms than the values of the metrics realized at the Gale-Shapley assignment. Because of the relatively small tension between increasing student welfare and fostering the distributional goals, there are school assignments that “have it all”—an outcome that realizes high welfare and diversity while at the same time encouraging neighborhood schools.

Table 17 presents statistics about the demographics across the overdemanded schools at the solution to the convex combination, and Table 13 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 8 and 17, 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. The maximal and minimal fractions of each school composed of each demographic group, representing the across-school-gap, are also close to those that obtain under the constrained optimal mechanism.

Ethnicity	BPS Avg.	Convex Combination of Metrics		
		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 17. % of Students From Each Ethnic Group at Overdemanded Schools In Convex Combination Solution

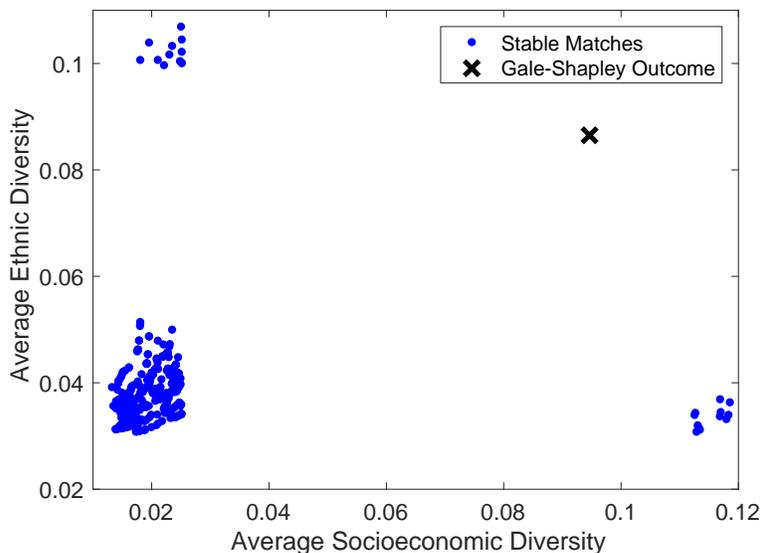


FIGURE 3. Ethnic vs Socioeconomic Diversity

APPENDIX G. FOR ONLINE PUBLICATION: ANALYSIS OF THE TENSION BETWEEN 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 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.

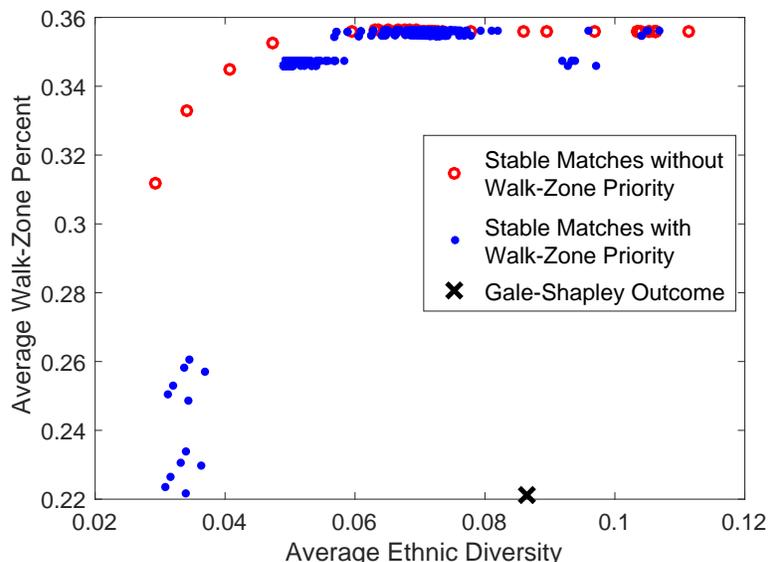


FIGURE 4. Diversity vs Neighborhood Schools without Walk-Zone Priorities

APPENDIX H. FOR ONLINE PUBLICATION: ENCOURAGING NEIGHBORHOOD SCHOOLS AND DIVERSITY WITHOUT WALK-ZONE PRIORITY CONSTRAINTS

Figure 4 presents my results on the tension between encouraging neighborhood schools and school diversity without imposing walk-zone priority constraints using the same format as Figure 2a. I have plotted the same set of stable matches displayed in Figure 2a along with the stable matches I compute without imposing the walk-zone priority constraints. If I choose a match 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 5 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 2b. 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 5 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 restrict my ability to encourage neighborhood schools. These welfare-optimal matches are towards the left side of Figure 5 where there is a significant gap between the possibility frontiers.

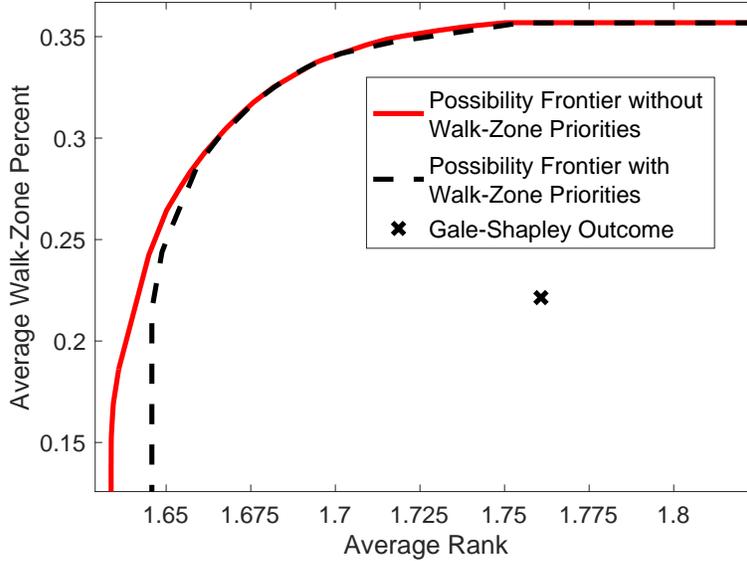


FIGURE 5. Welfare vs Neighborhood Schools without Walk-Zone Priorities

APPENDIX I. IMPLEMENTATION THROUGH MENUS-AND-RESERVES MECHANISMS

Suppose that the distribution π^S has been observed and a solution to one of the constrained optimization problems solved in Sections 4 and 5, $x(\circ, \circ; \pi^S)$, has been computed. In this section I describe mechanisms that implement $x(\circ, \circ; \pi^S)$ in three steps. First I define the benchmark menus-and-reserves mechanism, which helps illuminate how the menus and reserves work together as well as the connections between the mechanism I propose and other mechanisms that use reserves. Second, I analyze a mechanism that uses only the reserves, which helps illustrate the importance of the menu component of the mechanism. Finally, I define the iterative menus-and-reserves mechanism, which is the form of the menus-and-reserves mechanism that has the most practical appeal. The three mechanisms vary according to how closely they realize the desired solution to the constrained optimization problem. I also show that all three mechanisms can be run using the Gale-Shapley algorithm, which means the mechanisms are feasible and ODIC.

All three of my menus-and-reserves mechanisms are based on two ideas. First, I restrict the schools that each student can rank to a **menu** that is a function of the student's verifiable trait. Second, I allocate a **reserve** of seats at each school for students with each possible verifiable trait realization, and the reserves insure that the number of students of each type enrolled in each program approximately matches the solution to the optimization problem. The timing of the static menus-and-reserves mechanism is as follows:

- (1) Students are provided menus by the mechanism.
- (2) Each student submits a rank ordered list of schools from 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 a priority system defined by the the reserves.

A menu for a student with verifiable trait $v \in \mathcal{V}$, denoted $\mathcal{M}(v)$, includes all of the underdemanded schools as all of the overdemanded schools to which students with trait v are assigned with positive probability by the solution to the optimization problem.⁵¹ The set of overdemanded schools in $\mathcal{M}(v)$ is:

$$(I.1) \quad \cup_{\succeq} \{c : x(c, s; \pi^s) > 0 \text{ where } s = (\succeq, v)\}$$

The second component of the mechanism is the reserves, which are implemented through a priority structure. We use the term “reserve” because the mechanism gives students with each verifiable trait higher priority for some of the seats at each school program—it is as if the mechanism has reserved a set of seats for students with each trait. In the game with a finite number of students, the reserve allocates a number of seats at school program c for students with trait v equal to:

$$(I.2) \quad \mathcal{R}(v, c) = N \sum_{\succeq} x(c, (\succeq, v); \pi^s)$$

One needs to round $\mathcal{R}(v, c)$ to the nearest integer since $x(c, s; \pi^s)$ could generate a fractional assignment. This rounding will result in the total number of seats in the reserves to be above the capacity at some schools and below the capacity at others.⁵² When the number of seats assigned via reserves exceeds a school’s capacity, I randomly remove seats from the reserves at that school. When a school has fewer seats assigned via reserves than the capacity, I randomly add seats to the reserves at that school. In the end, every seat at every school is part of a reserve for some $v \in \mathcal{V}$.

Now I describe how to implement the menus and reserves through a priority structure. Taking inspiration from how the status quo mechanism treats walk-zone and open seats within each school program as distinct schools, the menus-and-reserves mechanism treats each set of reserved seats within each program as a distinct school. Students that are insured a seat at c have the highest priority for any seat at that school, and I assume that students that are insured seats are assigned seats from the set of $\mathcal{R}(v, c)$ seats reserved for their type before occupying any other seats.⁵³ Students with verifiable trait v have the second highest priority for seats reserved for students with that trait. Any student-types $s = (\succeq, v)$ where $c \in \mathcal{M}(v)$ have the third highest

⁵¹Menus are used in the current BPS mechanism, implying they are politically feasible.

⁵²The gap between the total number of seats in the reserves and the capacity of the schools ranged from -2 to $+3$, which represents less than 2% of the seats at each school.

⁵³Recall that being insured a seat at a school is a component of v .

priority for any seat in c that is not explicitly reserved for students with trait v .⁵⁴ All other student-types are considered unacceptable by school c . This last group consists of the student-types $s = (\succeq, v)$ such that $c \notin \mathcal{M}(v)$, which means that the menus are directly encoded in the priority structure. Ties in priority are broken using randomly distributed tie-breakers that are assigned to each student.⁵⁵

Because the menus-and-reserves system can be implemented using the Gale-Shapley algorithm, it inherits all of the desirable properties of the Gale-Shapley mechanism. These properties are summarized in the following remark.

Remark 1. *The match generated by a menus-and-reserves mechanism has the following properties:*

- (1) *School program capacities are not violated.*
- (2) *Truthful revelation of ordinal rankings is incentive compatible for the students.*
- (3) *The match is stable with respect to the priorities defined by the menus-and-reserves mechanism.*

It is worth taking a moment to consider the differences between the menus-and-reserves mechanisms and previous mechanisms that used quotas and/or reserves. Real-world school choice settings often use reserves to set aside seats for students with particular demographic traits. As noted above, the status quo mechanism used by BPS does this for students within each school's walk-zone. As a result of the practical interest in reserves, several previous papers have provided mechanisms and discussed practical issues related to how a reserve/quota is implemented (e.g., Agrawal and Somaini [4], Erdil and Kumano [27], Hafalir et al. [32], Ehlers et al. [25], Kamada and Kojima [36], Fragiadakis and Troyan [30]).⁵⁶ In order to satisfy lower bound constraints, many of these mechanisms consider artificial capacities that are weakly below the true capacity of the school program. The idea is that the artificially low capacity will prevent students from enrolling in a popular school, and the rejected students may help satisfy lower bound constraints by enrolling elsewhere. In a menus-and-reserves mechanism, the menus control where students can enroll and are designed to guide the students towards enrolling in school programs where their enrollment helps satisfy a distributional goal. The menus obviate the need to impose artificial capacity constraints that result in empty seats at popular schools. It is not clear to me how one could have designed effective menus without solving some underlying optimization

⁵⁴One can consider an even stronger notion of a reserve 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.

⁵⁵To fully replicate the solution of the continuum optimization problem, I would have to make each student's priorities potentially conditional on her preference declaration. Although conditioning priorities on preference declarations often violates incentive compatibility (e.g., the 1999 Boston mechanism), in this case the conditioning would remain approximately incentive compatible in line with a cardinal version of Theorem 2. I simulated a menus-and-reserves mechanism with preference dependent priorities, but the outcomes were negligibly different from those generated by the menus-and-reserves mechanism without preference dependent priorities.

⁵⁶Some papers differentiate between hard reserves that must be satisfied for a feasible match and soft reserves that merely serve as goals. My reserves take the form of soft reserves.

Metric	Menus-and-Reserves	Convex Combination of Metrics, Limit Model	Gale-Shapley Mechanism, Limit Model
Welfare, $R(x)$	1.725 [1.692, 1.754]	1.697 [1.659, 1.725]	1.763 [1.730, 1.796]
Ethnic Diversity, $D_{Abs}(x)$	0.082 [0.068, 0.094]	0.063 [0.052, 0.075]	0.0994 [0.092, 0.113]
Encouraging Neighborhood Schools, $N(x)$	0.427 [0.403, 0.453]	0.319 [0.274, 0.342]	0.221 [0.199, 0.241]

TABLE 18. Maximizing a Convex Combination of the Metrics

problem, which none of the prior works in the quotas and reserves literature has done to my knowledge.

The first column of Table 18 presents the realizations of my three metrics for the menus-and-reserves implementation of the solution to the convex combination (Section 5.2, Tables 16 and 17). I chose the menus and reserves based on the status quo distribution of types and did not adjust the menus or reserves across bootstrap runs. The second column of results refers to the solution of the same optimization problem 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-reserves 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-reserves mechanism in the finite setting as one might naively predict from the limit model. One reason that the performance drops relative to that realized in the limit model is that I am not adapting the menus or reserves 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 reserve or so many students with the trait that the reserve cannot accommodate all of them.⁵⁷ Both of these effects will distort the school assignment of the menus-and-reserves mechanism away from the exact solution of the limit model. A second reason the menus-and-reserves mechanism will differ from the limit model solutions is that one needs to round the fractional matchings, which will cause larger discrepancies at smaller schools.⁵⁸

⁵⁷In the later case, the excess students will compete for seats from other reserves that do not have enough students with the corresponding verifiable trait v to fill them.

⁵⁸If I run the menus-and-reserves 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 1.700, 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-reserves mechanism to the bootstrap

Based on the preceding discussion, one might assume that the menus-and-reserves system must also do worse than the limit model solution on the metric for encouraging neighborhood schools, but in fact the opposite occurs. The menus-and-reserves mechanism does 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.

By design, none of the schools can have an enrollment over capacity in a menus-and-reserves scheme. However, it is possible that a school will have empty seats because not enough students have that school in their menus. To get a sense for the scale of this problem, Table 14 in the appendix 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 at least half of the time. In addition, with the exception of the smallest schools (Snowden International and Lyon), all of the schools are at 90% of capacity or more at least 85% of the time.

The failure of the static menus-and reserves mechanism to fill the overdemanded schools is driven by the restrictions imposed by the menus on the schools that the students can rank. I now consider **reserves-without-menus** mechanisms that allow any student to rank any school, which is equivalent to setting $\mathcal{M}(v) = \mathcal{C}$. Examining the outcomes generated by this mechanism also help illustrate the role that the menus play in realizing the solution to the constrained optimization problem, $x(c, s; \pi^S)$. 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. The performance of the reserves-without-menus mechanism for the three metrics of interest is displayed in Table 19 and will be discussed shortly.

The third mechanism I wish to consider is the **iterative menus-and-reserves** mechanism, which adapts the menus-and-reserves mechanism to insure overdemanded schools are at or near full capacity. If a school falls below a capacity threshold after an initial execution, then the mechanism includes that school in the menus of all of the students and reruns the menus-and-reserves mechanism, hence the description of the mechanism as “iterative.” The iterative menus-and-reserves mechanism proceeds as follows:

- (1) Each student submits a complete preference ranking of the school programs to the mechanism.
- (2) The designer chooses a lower bound on the percentage of the seats at each overdemanded school program that must be assigned.
- (3) Build the menus $\mathcal{M}(v)$ as per Equation I.1 and set $\mathcal{A} = \emptyset$.

samples, I conclude that most of the failure of the finite mechanism to replicate the outcomes of the continuum solution is due to the choice to not adapt the mechanism to variation in π^S .

- (4) For each student, remove any school programs from her preference ranking that are not included in her menu.
- (5) Run the menus-and-reserves mechanism.
- (6) Identify any school programs that do not satisfy the lower bound on the fraction of its seats that are assigned.
 - If there is such a school program, add one such school program to the menu of all of the students, let $\mathcal{A} = \mathcal{A} \cup \{c\}$, and go to step 4.⁵⁹
 - If there is no such school program, finalize the assignment.

Note that despite the iterative nature of the mechanism, the students need only submit their preference ordering once. The mechanism automatically applies the menu relevant for the current iteration of the mechanism. The iterative menus-and-reserves mechanism will eventually terminate with all of the schools satisfying the constraint on the percentage of assigned seats. To see this, 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. 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, which would alter the menu of schools offered to her in future steps of the mechanism.⁶⁰ However, the probability of this event vanishes as the number of seats in each school program grows.

In order to attain exact incentive compatibility, the menus offered each student must be made independent of the student's preference declaration. If the student cannot manipulate her menu by declaring her preferences nontruthfully, then she can only be hurt by doing so. To formalize this, let \mathcal{A}_{-i} denote the value of \mathcal{A} realized after running the iterative menus-and-reserves mechanism with all of the students except for student i . Since the iterative menus-and-reserves mechanism can be solved quickly, this process can be completed for all students rapidly. Finally, run the static menus-and-reserves mechanism where the reserves are unchanged and each student is assigned a menu equal to $\mathcal{M}_i = \mathcal{A}_{-i} \cup_{\succeq} \{c : x(c, s; \pi^S) > 0 \text{ where } s = (\succeq, v)\}$. Note that the first component of \mathcal{M}_i are the schools at risk of being underdemanded and the second component of \mathcal{M}_i is the menu that would be offered the student in the static menus-and-reserves mechanism.

⁵⁹In my implementation, I chose the school program at random from the set of programs violating the lower bound on the seats allocated.

⁶⁰In the continuum version of the iterative menus-and-reserves mechanism, a single student's deviation cannot alter the menu she is offered at future stages, which means the mechanism is incentive compatible. It is easy to prove that truthfulness is approximately ODIC in a finite setting using an argument similar to that provided in the proof of Theorem 2.

Metric	Iterative Menus- and-Reserves	Reserves w/o Menus	Static Menus- and-Reserves
Welfare, $R(x)$	1.716 [1.683, 1.749]	1.730 [1.697, 1.768]	1.725 [1.692, 1.754]
Ethnic Diversity, $D_{Abs}(x)$	0.084 [0.071, 0.098]	0.097 [0.085, 0.108]	0.082 [0.068, 0.094]
Encouraging Neighborhood Schools, $N(x)$	0.403 [0.370, 0.436]	0.349 [0.326, 0.372]	0.427 [0.403, 0.453]

TABLE 19. Maximizing a Convex Combination of the Metrics

Table 19 shows the distribution of student welfare, diversity, and encouragement of neighborhood schools achieved by the iterative menus-and-reserves mechanism, the reserves-without-menus mechanism, and the static menus-and-reserves mechanism. I require that the schools be at least 95% filled in the iterative menus-and-reserves mechanism. Menus are important for encouraging racial diversity within the schools, which can be seen from the fact that the reserves-without-menus mechanism does roughly as well on the ethnic diversity metric as the status quo Gale-Shapley mechanism (see Table 7). 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.⁶¹

I now address three practical difficulties with implementing a menus-and-reserves mechanism. The first difficulty 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. If a new school is introduced, the designer could allow all of the students to rank the new school program. Once information about the distribution of student preferences has been collected, the optimization problem can be re-solved and the menus-and-reserves mechanism can be rebuilt. Unplanned changes include aggregate shocks to the distribution of student preferences caused by an unexpectedly poor outcome of a school evaluation or an expansion of a school's offerings of college-level courses. If there is a shock to the distribution of preferences, then the iterative menu-and-reserves 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-reserves system around the new solution.

⁶¹It is difficult to compare the results of Table 19 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-reserves system and (2) I allow the capacity constraints to be satisfied weakly in the menus-and-reserves system.

The second 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-reserves mechanism may make it impossible for some students to enroll in certain schools because those schools are not on the students' menus. My preferred solution is to modify the optimization problem underlying the menus-and-reserves system by (1) identifying the student-types that rank each overdemanded school under their true preferences 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.

The final concern is one of robustness. The above mechanism makes the most sense if the priority structure and menus remain fixed from year to year. In this case, the mechanism is transparent in that the menu and priorities are known in advance of any interaction with the mechanism. To test how well this process would work, I experimented with using the implementation based on the 2011-2012 school year data for the 2012-2013 school assignment problem. If it is effective using data from a year later, then this would have given confidence that the mechanism was robust.

A few changes to the market needed to be accounted for. First, a new school program was introduced, which I chose to add to each student's menu. Second, the capacities of the overdemanded schools changed. I chose to randomly assign (across the reserves) any seats associated with a capacity increase, and randomly eliminate any seats associated with a capacity reduction. I experimented with other techniques for handling these changes, but these did not yield qualitative changes in my results.

I found that the performance of the iterative menus-and-reserves mechanism was roughly equivalent to that of the menus without reserves mechanism. The reason is that there were rarely enough students with each overdemanded school on his or her menu to fill each overdemanded school to capacity. In the 2011-2012 school year, the number of students with each overdemanded school on his or her menu was on average triple the capacity of the school. Fixing the menus, the same statistic for the 2012-2013 school year is 1.22 times the capacity.

This level of performance suggests that the menus-and-reserves mechanism is not robust across years. Of course, one could collect preference information each year and redesign the menus-and-reserves mechanism. However, this would inherit all of the problems (e.g., transparency) of directly implementing a solution of a constrained optimization problem as a randomization over deterministic matches (see Appendix C) since the construction of the menus is directly based on the solution to the optimization problem.