

The Efficiency of Race-Neutral Alternatives to Race-Based Affirmative Action: Evidence from Chicago’s Exam Schools*

Glenn Ellison Parag A. Pathak[†]

August 2016

Abstract

Several public K-12 and university systems have recently shifted from race-based affirmative action plans to race-neutral alternatives. This paper explores the degree to which race-neutral alternatives are effective substitutes for racial quotas using data from the Chicago Public Schools (CPS), where a race-neutral, place-based affirmative action system is used for admissions at highly competitive exam high schools. We develop a theoretical framework that motivates quantifying the efficiency cost of race-neutral policies by the extent admissions decisions are distorted more than needed to achieve a given level of diversity. According to our metric, CPS’s race-neutral system is 24% and 20% efficient as a tool for increasing minority representation at the top two exam schools, i.e. about three-fourths of the reduction in composite scores could have been avoided by explicitly considering race. Even though CPS’s system is based on socioeconomic disadvantage, it is actually less effective than racial quotas at increasing the number of low-income students. We examine several alternative race-neutral policies and find some to be more efficient than the CPS policy. What is feasible varies with the school’s surrounding neighborhood characteristics and the targeted level of minority representation. However, no race-neutral policy restores minority representation to prior levels without substantial inefficiency, implying significant efficiency costs from prohibitions on affirmative action policies that explicitly consider race.

*We thank Katie Ellis, Susan Ryan and the staff at Chicago Public Schools for their expertise and help with the data. Vivek Bhattacharya, Alex Olssen, Rahul Singh provided excellent research assistance. We’ve also benefitted from the assistance of undergraduates Jackie Bredenberg and Brandon Enriquez. Ellison acknowledges support from the Toulouse Network for Information Technology. Pathak acknowledges support of National Science Foundation grant SES-1056325 and SES-1426566. Pathak is a member of the scientific advisory board of the Institute for Innovation in Public School Choice.

[†]Department of Economics, Massachusetts Institute of Technology, Cambridge MA 02142 and NBER, e-mail: g Ellison@mit.edu and ppathak@mit.edu

1 Introduction

Affirmative action is one of the most contentious issues in American social policy, nowhere more so than in the context of admission to selective educational institutions. Originally a 1965 executive order requiring federal contracts to “take affirmative action” to ensure that minorities and women are employed where available, affirmative action led to increasingly widespread efforts to boost minority participation in public K-12 school systems and universities in the 1970s and 1980s. More recently, such plans have come under attack on two fronts. Some states have banned race-based affirmative action.¹ And in 2003 the US Supreme Court ruled in *Grutter vs. Bollinger* that “strict scrutiny” must be applied to race-based plans: they must serve a compelling government interest that cannot be effectively achieved in a race-neutral manner. A June 2007 Supreme Court decision applied the earlier decision to strike down race-based admissions plans for public schools in Seattle and Jefferson County, Kentucky.² Against this backdrop, a number of public K-12 and university systems have shifted from race-based affirmative action plans to race-neutral alternatives.³

This paper explores the degree to which race-neutral alternatives are an effective substitute for racial quotas using data from Chicago Public Schools (CPS). CPS has adopted a new race-neutral admission policy for their selective high schools, including the district’s flagship Northside and Walter Payton high schools. From 1980-2007, a consent decree mandated race-based admissions for all Chicago magnet and selective enrollment schools. Admissions decisions were and are based on a composite score that combines middle school grades, standardized ISAT test scores, and a special entrance exam. Through 2009, the admissions procedure employed racial quotas that placed a cap on the number of white students in each school. In 2010, CPS replaced this system with a place-based affirmative action system where disadvantage is measured by a neighborhood proxy of socioeconomic status. Each census tract in the city was assigned to one of four tiers using an index combining five variables: median family income; a measure of adult educational attainment; home ownership rates; and the prevalence of single-parent households and non-native English speakers. Schools first filled 40% of their slots with the applicants having the highest composite scores. The remaining 60% of slots were then filled by

¹This includes constitutional amendments passed by ballot initiative in California in 1996 and Michigan in 2006.

²Department of Education guidelines suggest that a school district could draw attendance zones based on the racial composition of particular neighborhoods, as well as on race-neutral factors such as the average household income and average parental education level of particular neighborhoods within the school district, provided that all students within those zones would be treated the same regardless of their race (OCR 2011).

³Kahlenberg (2003) describes socioeconomic admissions criteria across a number of districts, while Kahlenberg (2008) lists 60 US school districts using socioeconomic status as a factor in student assignment.

dividing the slots equally across the four tiers and filling the slots with the highest-scoring remaining students living in census tracts belonging to each tier. The system was modified in 2012: the number of tier-reserved slots was increased from 60% to 70% and a sixth variable (test scores in the local elementary school) was added to the index.

Our focus on Chicago is motivated by several reasons. First, Chicago’s new plan is held up as a national model for achieving racial and ethnic diversity in selective public schools (e.g., Kahlenberg (2014)).⁴ Following the 2007 US Supreme Court decision, federal Department of Justice and Education Guidelines (OCR 2011) describe acceptable alternatives for achieving diversity at selective schools, which mirror Chicago’s plan:

For students who meet the basic admissions criteria, a school district could give greater weight to the applications of students based on their socioeconomic status, whether they attend underperforming feeder schools, their parents’ level of education, or the average income level of the neighborhood from which the student comes...

Second, Chicago employs an affirmative action scheme with explicit rules, which can be analyzed quantitatively. In particular, Chicago’s assignment mechanism is a variant of the student-proposing deferred acceptance algorithm, which is strategy-proof for participants (Dubins and Freedman 1981, Roth 1982). This property motivates simulating alternative assignment rules ignoring how these alternative policies might change how applicants rank schools. Third, the fact that Chicago’s schools are differently situated (with respect to neighborhood segregation) provides an opportunity to explore the effectiveness of race-neutral plans across settings. Finally, Chicago’s exam schools are highly visible and are frequently included in lists of the best U.S. public high schools, which makes them of independent interest.

Legal arguments about the effectiveness of race-neutral policies often solely focus on how minority representation would change when race-based policies are replaced. This emphasis is misguided, because it is always possible to eliminate racial gaps in admissions in a race-neutral manner by admitting students completely at random. However, elite schools do not randomly admit students because it conflicts with the main reason for their existence: they provide advanced students access to programs tailored to their knowledge and abilities. Therefore, assessing the “effectiveness” of race-neutral policies inherently requires measuring two components: whether the policy provides desired diversity (or

⁴Similar place-based schemes also exist in other affirmative actions systems for colleges and universities, including in Israel (Alon and Malamud 2014). For U.S. college admissions, Cashin (2014) argues for prioritizing applicants from disadvantaged neighborhoods rather than by race.

other) benefits and how race-neutral restrictions prevent admitting the students schools would most like to admit.

To motivate our analysis, Table 1 lists admissions outcomes for selected pairs of students. The top student in each pair is a higher-scoring applicant who is denied admission under the CPS tier plan, while the bottom student in the pair is a lower-scoring student who is admitted. In the first two pairs, underrepresented minority students are denied admission despite having perfect grades in 7th grade, perfect scores on the specialized school entrance exam, and very high ISAT scores. Meanwhile, the CPS procedure accepted students with much less impressive records: the comparison Hispanic student has a much lower score on the entrance exam and the comparison black admittee had worse grades and a somewhat lower entrance exam score. In these cases, the admitted students were admitted because they live in lower SES census tracts. The next two pairs provide examples in which the student denied admission qualifies for free/reduced lunch, while the lower-achieving student does not. The high-achieving students in these examples would be admitted if one eliminated affirmative action. They would also be admitted under a race-based policy: when minorities are added more efficiently one does not need to displace as many white and Asian-american students. Race-based policies are also better at accepting minority students with scores just below the cutoffs. The top students in the fifth and sixth pairs are such examples. In each case, the high-achieving students denied admission are free/reduced lunch eligible, yet are missed by the CPS race-neutral policy. The comparison students are admitted because they live in lower-SES census tracts, but are not themselves poor enough to qualify for a subsidized lunch. While these examples are extreme, they clearly illustrate inefficiencies with race-neutral admissions. The paper proposes a methodology to characterize inefficiency and uses it to systematically examine race-neutral admissions in Chicago.

Section 2 develops a simple model that motivates our measure of the effectiveness of an affirmative action plan. Students are assumed to benefit both from a curriculum tailored to their knowledge and ability, and from learning within a diverse student body. The optimal admissions policy is race-conscious and can be implemented either as a racial quota or a bonus scheme that gives some number of points to members of the underrepresented group. If race cannot be considered, the effectiveness of policies based on proxies correlated with race depends on the relationship between applicant attributes and race. We show, via an example, that it is possible that even a highly correlated proxy may be completely ineffective. Our main result shows that the social welfare produced by any policy, no matter how complex, is a function of two simple properties of the allocation: the average baseline achievement

Student	Race	Free Lunch	Score Components			Comp. Score	Application	
			GPA	ISAT	Exam		School	Admit?
A	Hispanic	No	4.0	96.3	100	98.8	Payton	No
B	Hispanic	No	4.0	94.0	77.7	90.6	Payton	Yes
C	Black	No	4.0	96.3	100	98.8	Payton	No
D	Black	No	3.25	100	93	89.3	Payton	Yes
E	Asian	Yes	4.0	98.0	99	99.0	Payton	No
F	Asian	No	4.0	100	85	95.0	Payton	Yes
G	White	Yes	4.0	98.0	99	99.0	Payton	No
H	White	No	3.75	93.0	97	93.9	Payton	Yes
I	Hispanic	Yes	4.0	96.0	95	97.0	Payton	No
J	Hispanic	No	4.0	94.0	77.7	90.6	Payton	Yes
K	Black	Yes	4.0	99.3	89	96.1	Northside	No
L	Black	No	4.0	94.3	76.7	90.3	Northside	Yes

Table 1: Examples of the efficiency costs of race-neutral affirmative action

of admitted students and minority representation. There is a Pareto frontier of efficient policies in this two-dimensional space. Race-neutral policies lie inside the frontier. We define a welfare-motivated notion of relative efficiency that represents the policies’ distance to the efficient frontier.

Section 3 uses data from CPS to estimate the shape of the Pareto frontier, which represents the fundamental constraints for affirmative action policies. We report on CPS’s two most selective selective-entry high schools: Walter Payton College Preparatory High School (Payton) and Northside College Preparatory High School (Northside). Focusing on the two most selective schools allows us to examine where affirmative action policies might have the largest impact on assignments. The schools also attract different applicant pools due to their geographic location, allowing us to examine whether race-neutral policies perform better in some situations. Payton’s central location is appealing to students living in a number of predominantly low-income black and Hispanic neighborhoods, whereas Northside is far from most predominantly black neighborhoods and would naturally draw more of its poor and minority students from surrounding middle-class neighborhoods.

We first note the percentage of underrepresented minority students that each school would have if admissions were based solely on the “composite score” that CPS currently uses. At both schools, minority representation can be increased substantially from its level under a purely score-based admis-

sions scheme with only a slight decrease in average composite scores, but score declines become much steeper as minority representation is further increased. Since the new policy is based on socioeconomic factors, we also report how policies affect the number of low-income students qualifying for free- and reduced-price lunch at each school and consider Pareto frontiers relevant when both minority- and low-income representation are concerns.

Section 4 examines Chicago’s socioeconomic-based affirmative action plan and variants that reserve more or fewer slots for students from low-SES areas. Our primary focus is on the relative efficiency of the CPS plan as a race-neutral method for increasing minority representation. Reserving a small number of slots for students from lower-socioeconomic tiers is nearly as efficient as racial quotas if one only wants to *slightly* boost minority enrollment. But the CPS policy is much less efficient when one tries to use it to keep minority representation anywhere close to its former level. There is also a limit to what can be accomplished with these policies ignoring efficiency concerns: minority representation at Payton would decrease even if one were to allocate 100% of seats to the tier-specific quotas. The CPS policy is more effective at Payton than Northside, because there are relatively fewer high-scoring applicants in the lowest SES areas who apply to Northside.

SES-based affirmative action plans are intended to also provide integration on other dimensions including low-income representation. In theory, such a benefit may offset these plans’ decreased efficiency as a tool for racial integration. But in practice we find that Chicago’s plan is actually less efficient than racial quotas as a means for increasing the number of low-income students at Payton and Northside.

We also examine within-school heterogeneity in scores and within-school racial gaps. While the majority of admitted students have 99th percentile scores under either plan, there are more lower-scoring students and their scores are lower under the CPS plan. The shift to race-neutral affirmative action could increase or decrease within-school racial gaps. One effect that tends to narrow the gap is that some of the lower-tier slots go to whites and Asians, which brings down their average scores. But a second opposite effect is that some of the highest scoring blacks and Hispanics (including some low-income students qualifying for free lunch) fail to gain admission which also brings down the minority average. We find that the second effect is stronger than the first and that the SES-based plans increase the within-school racial gap in composite scores.

Section 5 explores how much of the inefficiency of Chicago’s SES-based system is driven by a suboptimal implementation of the general idea as opposed to being inherent to any race-neutral policy. The current CPS socioeconomic index is an unweighted average of six tract-level characteristics.

Intuitively, an index-based rule can only be nearly as efficient as a race-conscious rule if the index is highly correlated with the minority status among students near the acceptance/rejection margin. The CPS index does not have this property. We identify alternative policies (involving SES-related “bonus points”) that would be somewhat more efficient. Using additional census tract characteristics to predict minority status, chosen by a cross-validated LASSO procedure, does not result in a significant improvement.

We also examine approaches inspired by rules like Texas’s top 10% rule, implemented based on Chicago’s 77 neighborhood areas or census tracts. We conceptualize these rules as a much more granular version of Chicago’s tier-based admissions. Instead of allocating slots to four tiers, the policies assign slots to much smaller entities like the census tract or neighborhood area. A policy using census tracts would be much less efficient than the CPS policy. Intuitively, the policy is highly inefficient when used for programs serving extreme high-achievers because they are outliers: there will be many cases in which geographic unit A has two candidates who are much stronger than the top candidate from geographic unit B, even when units A and B serve students of similar socioeconomic status. An implementation based on neighborhood areas is much better than the tract-based implementation and provides another way to improve the current CPS policy. Our biggest-picture conclusion, however, is that all feasible place-based policies appear to be substantially less efficient than race-based policies.

This paper is most closely related to the innovative work of Fryer, Loury, and Yuret (2008), which developed a model of “color-blind” vs. “sighted” affirmative action to show that the reduction in the quality of the student body is an efficiency cost of color-blindness. For university admissions, they consider a thought experiment in the College and Beyond dataset where each of several colleges admits a student body half as large as their current one from an applicant pool consisting of their current students. We build on their analysis and expand on their empirical work in several dimensions: we examine examine more realistic changes to admissions policies, an approach made possible by having data on the full applicant pool; we compare a variety of race-neutral policies; and we examine additional outcomes including within-school achievement gaps and effects on the representation of low-income students. Cestau, Epple, and Sieg (2015) develop an econometric model of the referral process for taking the admissions tests for selective elementary schools. They report that profiling by race and income together with affirmative action based on free lunch status can achieve 80% of level of black enrollment as a race-based affirmative action plan. Corcoran and Baker-Smith (2014) study admissions policies at New York’s exam schools, focusing on a descriptive account of the decision to apply. Though their main interest is not in affirmative action, they simulate top 10% rules based

on 7th grade scores and find that it leads to an increase in black and Hispanic representation at the schools.

Our model is also related to several other papers on affirmative action. Chan and Eyster (2003) show that if an admissions office is prohibited from using affirmative action, they can still increase minority representation by placing less weight on applicant qualifications. In a study motivated by Chicago’s assignment mechanism, Dur, Pathak, and Sönmez (2016) characterize the best precedence, or order in which affirmative-action and non-affirmative action slots at schools are processed, for applicants from particular tiers. In our simulations, we use their results to focus on the best tier-blind precedence for applicants from the most disadvantaged tier. Epple, Romano, and Sieg (2008) develop an equilibrium model of affirmative action and tuition policies to show that a ban on affirmative action leads to a decline in minority students at top-tier colleges. Ray and Sethi (2010) show that any score-maximizing race-neutral affirmative action plan must be a non-monotone function of past performance.

2 A Model of Elite Schools and Affirmative Action

2.1 Rationalizing elite schools

We begin with a simple model in which it is efficient for a school system to create an elite school when students benefit from a curriculum tailored to their ability or preparation. Formally, suppose that a school system serves a continuum of students with types θ continuously distributed on \mathbb{R} . The type θ can be thought of as a composite of the student’s ability and/or preparation. The school system operates two schools with curricula $c_1, c_2 \in \mathbb{R}$. Suppose that the expected outcome for a type θ -student if assigned to school i depends on the student’s type and on the match between the students’ type and the curriculum as follows:

$$V_i(\theta) = h(\theta) - k(\theta - c_i)^2, \tag{1}$$

where $h(\theta)$ governs the direct relation between student type and expected outcome, the curriculum c_i is described by the type of the student it serves best, and k indexes the relative importance of student-curriculum matching. An assignment policy specifies which student types are assigned to each school and the curriculum taught in each school.⁵ The optimal assignment policy maximizes the

⁵While curricula are a choice variable in this model, an alternative interpretation is that only school assignments are a choice variable and after school assignments are made, classroom dynamics force teachers to adopt the curriculum best suited to the mean student in the school.

sum of students' expected outcomes.

Given this objective, aggregate outcomes are highest if students are grouped according to θ and each school's curriculum is matched to the set of students it serves. The following proposition formalizes the rationale for elite schools.

Proposition 1 *If students' outcomes are given by (1), the optimal assignment policy is defined by a cutoff $\hat{\theta}$, where students with $\theta > \hat{\theta}$ are assigned to school 1 and students with $\theta < \hat{\theta}$ are assigned to school 2. The curricula are $c_1 = \mathbb{E}(\theta|\theta > \hat{\theta})$ and $c_2 = \mathbb{E}(\theta|\theta < \hat{\theta})$. Furthermore, the optimal $\hat{\theta}$ is the solution to*

$$\hat{\theta} = \frac{\mathbb{E}(\theta|\theta < \hat{\theta}) + \mathbb{E}(\theta|\theta > \hat{\theta})}{2}.$$

In any assignment scheme, the curricula c_1 and c_2 are the mean θ 's of the students assigned to each school. Given any curricula with $c_1 > c_2$, optimal student assignment implies that each student with $\theta > \frac{c_1+c_2}{2}$ is assigned to school 1, which corresponds to the condition on $\hat{\theta}$ in the proposition.

The top panel of Figure 1 provides a simple illustration of a situation in which an elite school is optimal. The curve represents the density $f(\theta)$ of the type distribution, which continues substantially further to the right.⁶ The vertical line indicates the cutoff $\hat{\theta}$ with students to the right of this line being assigned to school 1. The optimal curricula c_2 and c_1 are the conditional means of θ in each school. The cutoff $\hat{\theta}$ is equidistant between c_2 and c_1 .

2.2 Optimal affirmative action with diversity benefits

Continuing our model development, suppose that student outcomes depend on both curriculum matching and classroom diversity. Admissions arrangements at elite schools often reflect both of these concerns. For example, a Blue Ribbon Commission evaluating Chicago's policy felt that it was important for the "programs to maintain their academic strength and excellent record of achievement, but also believes that diversity is an important part of [their] historical success" (BRC 2011).

Suppose that the school system serves two populations indexed by $j = \{o, u\}$, where o denotes overrepresented and u denote underrepresented at the elite school absent affirmative action. Let m be the fraction of students who belong to the population u . Let F_o and F_u be the distributions of types within each population. Due to some source of disadvantage, we assume that the type distribution is

⁶The figure was generated by solving a model with $k = 1$, $d = 1$, and θ having a standard exponential distribution.

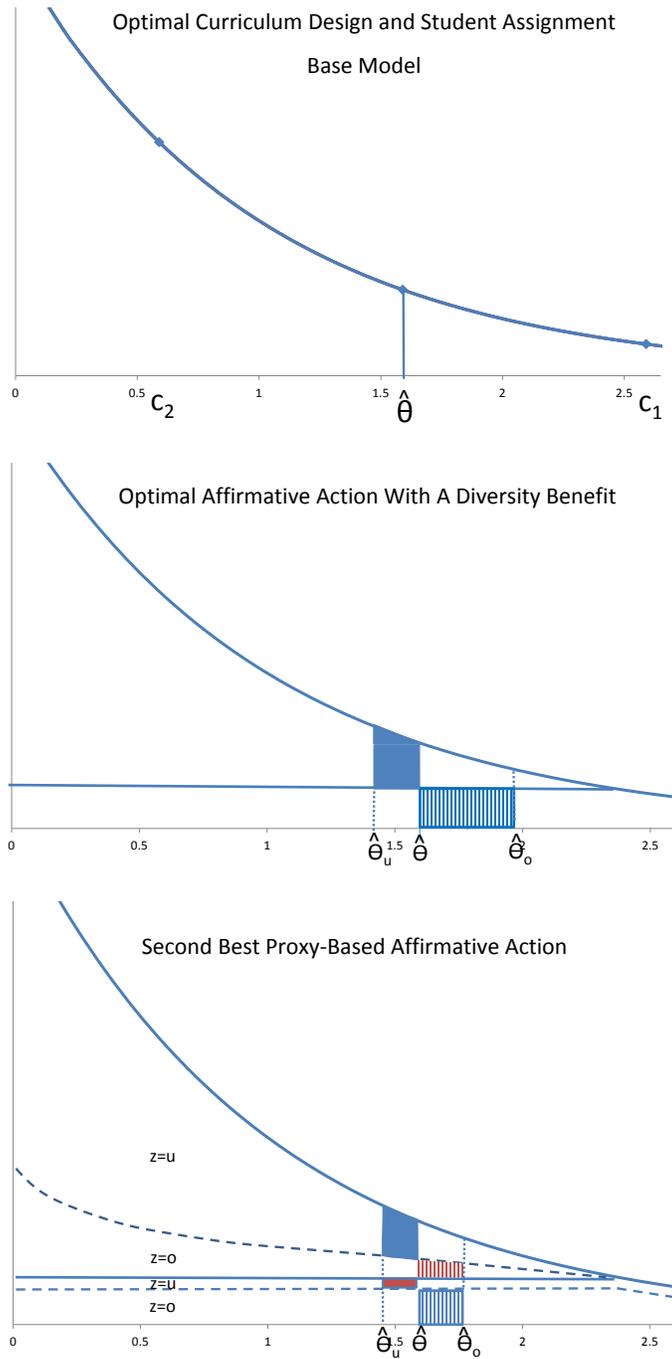


Figure 1: Optimal admissions policies (a) without a diversity benefit; (b) with a diversity benefit; and (c) when only an imperfect proxy is available.

higher in population o :

$$1 - F_o(\theta) \geq 1 - F_u(\theta), \text{ for all } \theta.$$

We add a diversity benefit to the model by assuming that outcomes are highest when each school is representative of the population:

$$V_i(\theta) = h(\theta) - k(\theta - c_i)^2 - d|m_i - m|, \quad (2)$$

where m_i is the fraction of students in school i from population u and d indexes the importance of diversity. We assume that diversity is valued because it improves all students' educational outcomes, e.g. by exposing them to different viewpoints. Similar results could be derived in models in which affirmative action is motivated by other concerns, e.g. if younger generations benefit from minority students' serving as role models as in Chung (2000).

The next proposition shows how the model provides a rationale for affirmative action.

Proposition 2 *If student outcomes are given by (2), the optimal assignment policy is defined by a pair of cutoffs $\hat{\theta}_o$ and $\hat{\theta}_u$ where students from population j are assigned to school 1 if and only if $\theta > \hat{\theta}_j$ for $j \in \{o, u\}$. The optimal cutoffs satisfy*

$$\hat{\theta}_o > \hat{\theta}_u \quad \text{and} \quad \hat{\theta}_o - \hat{\theta}_u = \frac{d}{k(c_1 - c_2)}.$$

This proposition states that within each population, students are perfectly sorted on θ , but the elite-school cutoff is lower for students from population u . The gap between the two cutoffs is larger when diversity is more important (d is larger), when educational outcomes are less sensitive to student-curriculum mismatch (k is smaller), and when curricula are closer together ($c_1 - c_2$ is smaller).⁷ The proof of this proposition is in the appendix.

Affirmative action could be implemented either with a reservation or a bonus scheme. The school system could use quotas and reserve $m(1 - F_u(\hat{\theta}_u))$ seats at school 1 for members of population u . Alternatively, affirmative action could be implemented under a race-as-a-factor scheme in which students from population u are admitted over students of population o if their scores are “close enough,” with close enough being defined as being within $\frac{d}{k(c_1 - c_2)}$.

⁷The cutoffs are related to the curricula by $\hat{\theta}_u = \frac{c_1 + c_2}{2} - \frac{d(1-m)}{k(c_1 - c_2)}$ and $\hat{\theta}_o = \frac{c_1 + c_2}{2} + \frac{dm}{k(c_1 - c_2)}$.

The second panel of Figure 1 provides an illustration. We divide the student population from the first panel into two subpopulations: the lower part (below the horizontal line) reflects the distribution of types within population o and the upper portion reflects the distribution within population u .⁸ If the same cutoff $\hat{\theta}$ were applied to all students, then population u will be substantially underrepresented at school 1. Starting from this point, increasing $\hat{\theta}_o$ and decreasing $\hat{\theta}_u$ improves diversity. Initially, there is no first-order loss in match quality: the added underrepresented students have θ 's that are nearly identical to those of the displaced overrepresented students. Accordingly, a small change of this type always improves welfare. The cutoffs in the figure are those that are optimal given the assumed parameters. The solid shaded area is the set of added underrepresented students. The striped region shows the displaced students from the overrepresented population. At the cutoffs $\hat{\theta}_u$ and $\hat{\theta}_o$, it is no longer welfare-improving to add population u students to and displace population o students from the elite school. While this would improve diversity, the mismatch costs would exactly offset the benefit.

2.3 Affirmative action with imperfect proxies

We next turn our attention to second-best affirmative action policies when membership in a particular population cannot be explicitly considered, but correlated proxies can. Consider the two-population model as above in which a school system cares about the fraction m_i of students from population u at each school i with outcomes given by equation (2). Suppose now in addition that each student has an observable characteristic $z \in \{o, u\}$ that is an imperfect proxy for the population to which they belong. Assume that z is equal to the population j to which the student belongs with probability $p > \frac{1}{2}$ independently of θ . Suppose that the school system is prohibited from considering the population j to which the student belongs in admissions, but can consider the proxy z .

The posterior probability that a student of observed type (θ, z) belongs to population u is

$$m(\theta, z) = \begin{cases} \frac{f_u(\theta)p}{f_u(\theta)p + f_o(\theta)(1-p)} & \text{if } z = u \\ \frac{f_u(\theta)(1-p)}{f_u(\theta)(1-p) + f_o(\theta)p} & \text{if } z = o. \end{cases}$$

Given only the assumptions we have imposed so far, the optimal affirmative action policy need not involve cutoffs. Intuitively, if students with type θ are unlikely to be from population u even if they are

⁸The figure is constructed by solving a numerical example. We take $k = 1$, $d = 1$, and θ exponentially distributed and assume that the mass of overrepresented students of type θ is $\min(0.09, e^{-\theta})$. The o population is over-represented among high values of θ because the distribution continues to the right. In addition, we have imposed the constraint that the number of students in the elite school be the same as in the top panel of the figure. Otherwise, the optimal policy would assign fewer students to the elite school. The elite school would also be smaller in the top panel if that calculation had assumed that diversity is valuable.

observed to have $z = u$, while students with a slightly lower type θ' are likely to be from population u , it could be optimal to assign (θ, u) students to school 2 and (θ', u) students to school 1. In practice, it seems that it would be politically infeasible to adopt rules that penalized students for having higher scores even if a school system did know about differences in likelihood ratios of this type. Accordingly, we restrict our attention here to cutoff policies $(\hat{\theta}_u, \hat{\theta}_o)$ under which the school system assigns students with observed characteristics (θ, z) to school 1 if and only if $\theta \geq \hat{\theta}_z$.

Proposition 3 *Suppose student outcomes are given by (2) and that the school system is prohibited from considering the population j to which a student belongs, but can adopt admissions policies with cutoffs $(\hat{\theta}_u, \hat{\theta}_o)$ that depend on the proxy z . Then, the optimal cutoffs satisfy*

$$\hat{\theta}_o - \hat{\theta}_u = \frac{d}{k(c_1 - c_2)} \left(m(\hat{\theta}_u, u) - m(\hat{\theta}_o, o) \right).$$

Note that the equation characterizing the difference in the cutoffs is nearly identical to that in Proposition 2: the only difference is that the right-hand side is multiplied by $m(\hat{\theta}_u, u) - m(\hat{\theta}_o, o)$.⁹ This term is the expected increase in the number of minority students in school 1 when we replace a marginal student with observed proxy $z = o$ with a marginal student with $z = u$. It is always less than one. The diminished effect on minority representation of each student that is moved reduces the incentive to split the cutoffs. The formula simplifies in special cases. When $p = 1$, so the proxy is completely informative, the final term on the right-hand side will be equal to 1 giving the same formula as in Proposition 2. And when $p = \frac{1}{2}$, so the proxy is completely uninformative, the final term is equal to zero when we implement a single cutoff.¹⁰

The bottom panel of Figure 1 provides a simple illustration of the optimal policy in a numerical example with $p = \frac{3}{4}$.¹¹ The dashed lines show how students within each population j are divided on z . In the lower region of population o students, three-fourths of the students are in the lower $z = o$

⁹Formulas for the cutoffs are also similar to those for the previous model:

$$\hat{\theta}_u = \frac{c_1 + c_2}{2} - \frac{d(m(\hat{\theta}_u, u) - m)}{k(c_1 - c_2)} \text{ and } \hat{\theta}_o = \frac{c_1 + c_2}{2} + \frac{d(m - m(\hat{\theta}_o, o))}{k(c_1 - c_2)}.$$

¹⁰Note that the model of section 2.1 did not include a diversity benefit. One reason for organizing the exposition in this way was to defer bringing up the option of using something other than a cutoff policy to take advantage of differences in likelihood ratios.

¹¹The numerical example is as before: $k = 1$, $d = 1$, θ is exponential with mean 1 and the mass of overrepresented students of type θ is $\min(0.09, e^{-\theta})$. We also fix the size of the elite school to be the same as in the upper two panels.

bar. In the upper region of population u students, three-fourths are in the upper $z = u$ band. The cutoffs $\hat{\theta}_o$ and $\hat{\theta}_u$ now apply to the signal z rather than to the population. Hence, students added by shifting $\hat{\theta}_u$ down (the solid shaded regions) include both students from population u (in blue) and students from population o (in red). The students displaced by raising $\hat{\theta}_o$ (shaded by vertical bars) also include students from both populations.

The optimal cutoffs in this model are closer together than the cutoffs under unrestricted affirmative action because only a fraction of the students who are added to the elite school by an incremental reduction in $\hat{\theta}_u$ are from population u . Hence, the diversity benefit for each student added is smaller than before. The per-student mismatch costs are the same, so it is optimal to reduce the degree of affirmative action. Representation of the disadvantaged group is reduced by two separate mechanisms: (1) holding fixed the number of students affected by the affirmative action program only a portion of the added students come from the underrepresented population; and (2) because affirmative action is a less effective tool, the school system uses it less aggressively. The red regions in the figure highlight the inefficiency of the policy: some students from the underrepresented population are denied admission despite having θ 's that are high enough to earn them admission under a score-based admission system, and some students from the overrepresented population are admitted despite having low θ s.

2.4 An example in which proxy-based affirmative action is ineffective

Although proxy-based affirmative action may be an effective substitute for the unconstrained optimal policy as in the numerical example above, it need not be. We next show that there are situations in which it is actually counterproductive.

We depart from the previous model in assuming that the distribution of θ is lower in the underrepresented population for two reasons: θ is related to income and the underrepresented population has a lower income distribution; and even after conditioning on income, the distribution of θ is lower in the underrepresented population due to some other disadvantage not captured by a variable that can be used in admissions. Specifically, assume that student ℓ 's type is

$$\theta_\ell = \text{Income}_\ell + 0.1 \text{ Over}_\ell,$$

where Income_ℓ is student ℓ 's income and Over_ℓ indicates membership in the overrepresented population. Suppose also that the income distribution is systematically higher in population o : it is uniformly distributed on $[0, 1]$ in population u and uniform on $[0.1, 1.1]$ in population o . As a result, θ is uniform

on $[0, 1]$ in population u and uniform on $[0.2, 1.2]$ in population o . Assume that the two populations are equal in size. Suppose for this example that the elite school has a capacity of 20% of the total number of students.

If the school system uses a purely score-based admissions policy, then 25% of the seats in the elite school are assigned to the underrepresented population. The admissions cutoff is $\hat{\theta} = 0.9$. Thirty percent of population o and 10% of population u have types above this level. We can think of the underrepresentation as equally attributable to the two sources of disadvantage.¹²

Suppose that the school system knows each student's income and can use this variable in its admissions process. Since low income is correlated with underrepresentation and directly effects θ , it seems natural to use income as a proxy variable.

Example 1 *Suppose that the school system divides students into four equal-sized tiers on the basis of income and then selects the top 20% of students from each tier. Then no students from the underrepresented group gain admission to the elite school under the tier-based affirmative action policy.*

Each tier has mass 0.25. The bottom tier consists of all students with incomes below 0.3: underrepresented students with incomes uniformly distributed on $[0, 0.3]$ and overrepresented students with incomes uniform on $[0.1, 0.3]$. The population u students in this tier have θ 's uniform on $[0, 0.3]$. The population o students have θ 's uniform on $[0.2, 0.4]$. The cutoff that selects 20% of this tier is $\hat{\theta}_1 = 0.3$. All low-income students above this cutoff are from the overrepresented group. Calculations for the other tiers are similar. The three higher tiers consist of students with incomes between 0.30 and 0.55, between 0.55 and 0.80, and above 0.80, respectively. The θ cutoffs that select the top 20% of the tiers are $\hat{\theta}_2 = 0.55$, $\hat{\theta}_3 = 0.80$, and $\hat{\theta}_4 = 1.1$. All students in the applicable income range with θ above these cutoffs are from the overrepresented group.

While this result may seem paradoxical at first, it is related to Chan and Eyster (2003)'s observation that it is possible to increase representation of an underrepresented group by adding noise to the decision process. In this example, income is playing two roles. It is a source of disadvantage for the underrepresented population. But it can also be thought of as a source of noise that helps some underrepresented students overcome their other disadvantage when a score-based procedure is used. The income-based affirmative action procedure removes a source of disadvantage, but it also removes

¹²If both populations had incomes uniform on $[0, 1]$, the cutoff would have been $\hat{\theta} = 0.85$ and 37.5% of the seats would have gone to members of the underrepresented group. The underrepresented group is also assigned 37.5% of the seats if the income differences are present, but the group-related difference is not.

a source of noise. The example is constructed so that the latter effect is more important than the former.

Although we've considered a specific proxy system based on four tiers in this example, a number of other income-based affirmative action policies would work just as poorly in this environment. No underrepresented students would qualify for the elite school under any similar tier-based system using a larger number of tiers. Nor would underrepresented students be assigned any seats under a policy that defined student k 's income-adjusted score as $\theta_k - \alpha \text{Income}_k$ and selected students in order of adjusted scores for any $\alpha > \frac{2}{3}$. On the other hand, an unrealistic feature of this example is that the only source of heterogeneity in θ is income-related. Adding an additional source of heterogeneity within each population would make affirmative action work better. But the noise-reduction effect of the correlated proxy would still be present and there are specifications where the tier-based plans would still be counterproductive as a method to address underrepresentation.

In the model of the previous section, proxy-based affirmative action approximates the optimal unrestricted affirmative action plan when the proxy is highly correlated with group membership. The example shows that even a highly correlated proxy can be of no help if it is separately related to the admissions score. This range of possibilities motivates our empirical investigation of proxy-based affirmative action.

2.5 Welfare, Pareto frontiers, and measuring inefficiency

We now discuss the welfare consequences of affirmative action policies in the context of our basic model with two groups and outcomes given by equation (2). Throughout this section, we fix enrollment in the elite school 1 at some fraction n_1 . We will think of the choice of an optimal affirmative action plan as a two-step process: (1) for each level m_1 of minority representation in school 1, we find the optimal plan that makes the fraction of underrepresented students m_1 ; and (2) we then choose an optimal plan by maximizing over m_1 . Let m_1^s be the fraction of students at school 1 who are from population u under a purely θ -based admissions policy. The optimal policy with diversity concerns will have $m_1 \in (m_1^s, m)$, and we restrict attention to this interval for the remainder in the subsequent discussion.

Conditional on m_1 , the optimal plan simply sends the $n_1 m_1$ underrepresented students with the highest θ 's to school 1 along with the $n_1 - n_1 m_1$ highest scorers from the overrepresented population.

The “first-best” cutoff scores are

$$\hat{\theta}_u^{FB}(m_1) = F_u^{-1}\left(1 - \frac{n_1 m_1}{m}\right) \quad \text{and} \quad \hat{\theta}_o^{FB}(m_1) = F_o^{-1}\left(1 - \frac{n_1 - n_1 m_1}{1 - m}\right).$$

The optimal curricula are

$$\begin{aligned} c_1^{FB}(m_1) &= (1 - m_1)\mathbb{E}_o(\theta|\theta > \hat{\theta}_o^{FB}(m_1)) + m_1\mathbb{E}_u(\theta|\theta > \hat{\theta}_u^{FB}(m_1)) \\ c_2^{FB}(m_1) &= (1 - m_2)\mathbb{E}_o(\theta|\theta < \hat{\theta}_o^{FB}(m_1)) + m_2\mathbb{E}_u(\theta|\theta < \hat{\theta}_u^{FB}(m_1)), \end{aligned}$$

where \mathbb{E}_o and \mathbb{E}_u are expectations of θ s drawn from F_o and F_u , respectively.

In the discussion below, it is useful to have notation for the difference between the mean types of the students at the two schools. Using ℓ to index students and the function $s(\ell)$ to denote the school to which student ℓ is assigned, define the difference between mean types as:

$$\Delta(s(\ell)) \equiv \mathbb{E}(\theta_\ell | s(\ell) = 1) - \mathbb{E}(\theta_\ell | s(\ell) = 2).$$

(We often omit the argument of Δ when the policy under discussion is clear.) Write $\Delta^{FB}(m_1)$ for the difference under the optimal policy given representation level m_1 . Note that because the optimal curricula are the mean types within a school, we also have

$$\Delta^{FB}(m_1) = c_1^{FB}(m_1) - c_2^{FB}(m_1).$$

Consider now any other plan that assigns student ℓ to school $s(\ell)$. Suppose that the plan also results in school 1 having a fraction m_1 of students from population u and that the curricula at the schools are c_1 and c_2 . We can decompose social welfare under this plan into a sum of two terms, $W = W_e + W_d$, where W_e reflects curriculum-matching benefits and W_d reflects diversity benefits:

$$\begin{aligned} W_e &= \mathbb{E}[h(\theta_\ell) - k(\theta_\ell - c_{s(\ell)})^2], \\ W_d &= n_1(-d|m - m_1|) + n_2(-d|m - m_2|). \end{aligned}$$

The proposition below gives simple formulas for these two terms that apply whenever curricula are optimal given the student assignment. The proof is in the appendix.

Proposition 4 *Suppose that a fraction m_1 of the n_1 students assigned to school 1 under $s(\ell)$ are from population u with $m_1 < m$. Suppose that the curricula c_1 and c_2 are optimal given the student*

assignments. Then,

$$\begin{aligned} W_e &= \mathbb{E}(h(\theta)) - k \text{Var}(\theta) + kn_1n_2\Delta^2 \\ W_d &= 2n_1d(m_1 - m), \end{aligned}$$

where $\Delta = \mathbb{E}(\theta_\ell | s(\ell) = 1) - \mathbb{E}(\theta_\ell | s(\ell) = 2)$.

Write W^{sb} and Δ^{sb} for the welfare and cross-school difference in average types under a purely score-based admissions policy. Then the difference between social welfare under policy $s(\ell)$ and the score-based policy is

$$W - W^{sb} = 2n_1d(m_1 - m_1^{sb}) - kn_1n_2 \left((\Delta^{sb})^2 - \Delta^2 \right). \quad (3)$$

W_e is an increasing function of Δ because the benefits from dividing the system into two schools comes from differentiating instruction. Students are better off when this can be done to a greater extent. The formula comparing welfare under an affirmative action plan to welfare under a purely score-based plan brings out the basic trade-off. The first term on the right-hand side of (3) reflects the diversity benefits that are proportional to the degree to which the school composition changes, $m_1 - m_1^{sb}$, and to the importance d of diversity. This trades off against the loss in benefits from curriculum matching reflected in the second term. A purely score-based approach maximizes the differences in the types across the schools so any affirmative-action policy will have a disadvantage on this dimension. The magnitude of the disadvantage depends on the degree to which curriculum differentiation Δ is reduced.

Under unrestricted affirmative action, we can think of school systems as choosing among a set of points in (W_d, W_e) space. For each value of m_1 , we can determine the allocation of students to schools that results from a quota that reserves a fraction m_1 of the seats for underrepresented students and compute the values (W_d, W_e) associated with those allocations. This downward-sloping set of points is the Pareto frontier. The relative importance of curriculum matching k and diversity d will determine which point on the frontier a welfare-maximizing school system would choose. Proposition 4 tells us that W_d is a linear function of m_1 and W_e is a monotone function of Δ , which is in turn a linear function of the average type in the elite school, $\mathbb{E}(\theta_\ell | s(\ell) = 1)$. Accordingly, one can equivalently represent affirmative action plans as points in a two-dimensional set with the following axes: the fraction of students at the elite school who are underrepresented minorities on the x -axis and the average composite score among students at the elite school on the y -axis. A corollary of the above

proposition describes the Pareto frontier in this space.

Corollary 1 *The Pareto frontier of affirmative action policies $\hat{\theta}^{FB}(m_1)$ form a downward sloping curve in $(m_1, \mathbb{E}(\theta_\ell | s(\ell) = 1))$ space for $m_1 \in [m_1^{sb}, m]$. The curve has slope zero at its left endpoint and is concave.*

The argument for this result is simple: the slope of the curve is just the difference in scores between the students who are being added and those being displaced when quotas are expanded. This is $-(\hat{\theta}_1^{FB}(m_1) - \hat{\theta}_2^{FB}(m_1))$. This curve is downward sloping because this difference is negative when $m_1 > m_1^{sb}$. The slope is zero at the left endpoint because $\hat{\theta}_1^{FB}(m_1^{sb})$ and $\hat{\theta}_2^{FB}(m_1^{sb})$ are both equal to the score-based cutoff $\hat{\theta}^{sb}$. The curve is concave because $\hat{\theta}_1^{FB}(m_1)$ is increasing in m_1 and $\hat{\theta}_2^{FB}(m_1)$ is decreasing in m_1 which makes the difference increasing in absolute value.

These observations translate into economically important points. The zero-slope result shows that there is no first-order loss to small-scale affirmative action because the added underrepresented students and the displaced overrepresented students will have nearly identical preparation. The concavity result shows that the curriculum-matching losses become more severe as the number of added students increases because the marginal students being added and the marginal students being displaced are more different.

Any other affirmative action policy, regardless of how it is constructed, produces an outcome corresponding to some point in $(m_1, \mathbb{E}(\theta_\ell | s(\ell) = 1))$ space below the Pareto frontier. The distance between the plan and the frontier can be thought of as measuring the efficiency losses incurred by adopting the proxy-based plan instead of a race-based plan. The corollary below provides a formal justification: the efficiency loss relative to what could have been achieved with the same minority representation is a monotone function of the degree to which that average θ at the elite school has been reduced.

Corollary 2 *Suppose that a fraction m_1 of the n_1 students at school 1 under assignment $s(\ell)$ are from population u with $m_1 < m$. Let $\Delta = \mathbb{E}(\theta_\ell | s(\ell) = 1) - \mathbb{E}(\theta_\ell | s(\ell) = 2)$. Then, the welfare loss from policy $s(\ell)$ rather than the most efficient policy under which a fraction m_1 of the students at school 1 are from population u is*

$$W^{FB}(m_1) - W = kn_1n_2(\Delta^{FB}(m_1)^2 - \Delta^2).$$

Our empirical work reports a measure of “relative efficiency” when discussing plans. To define this concept, consider any assignment $s(\ell)$ under which a fraction m_1 of the students at school 1 are

from population u and the difference in the average θ across the two schools is $\Delta(s(\ell))$. This plan will have reduced the average θ at school 1 relative to its level under purely score-based admissions by $n_2(\Delta^{sb} - \Delta(s(\ell)))$. If the school had instead adopted the optimal plan given m_1 , scores would have been reduced by the smaller amount $n_2(\Delta^{sb} - \Delta^{FB}(m_1))$.

We define the relative efficiency of policy $s(\ell)$ by the ratio of these two terms:

$$\text{Relative Efficiency} \equiv \frac{\Delta^{sb} - \Delta^{FB}(m_1)}{\Delta^{sb} - \Delta(s(\ell))}.$$

Note that this measure is scaled so that a plan has a relative efficiency of 100% if it produces the same social welfare as the best possible plan under which school 1 has a fraction m_1 of its students from population u . Relative efficiency will be lower than 100%, but still positive under any plan other than the optimal race-conscious plan. Figure 2 contains an illustration of the calculation. The thicker blue curve represents the Pareto frontier of feasible affirmative action plans with the axes corresponding to the level of minority representation (on the x -axis) and the average score of students assigned to the elite school (on the y -axis). The horizontal dashed line is the average score under a purely score-based admissions policy. The lower red curve might correspond to a family of race-neutral admissions policies that use some proxy correlated with minority status, e.g. reserving some fraction of seats to be allocated according to the proxy and allowing open competition for the rest. The large filled circle highlights one policy that achieves a 40% minority representation level. We describe it as 25% efficient because only 25% of the reduction in test scores was necessary for that level of minority representation. That is, the highlighted policy reduces test scores at the elite school by 5.2 points (from 99 to 93.8) whereas one could have obtained 40% minority representation while only reducing test scores by 1.3 points (from 99 to 97.7) if one had chosen a policy on the Pareto frontier.

The corollary suggests that welfare losses are quadratic in the reduction in Δ , so that incremental welfare losses from a percentage-point decrease in relative efficiency are larger at lower levels of efficiency, e.g. the loss from moving from a 90% efficient plan to an 89% efficient plan is smaller than the loss from moving from a 70% efficient plan to a 69% efficient plan. We chose to define relative efficiency as we have rather than using a quadratic scale because we feel that it is easier to interpret and because the corollary implies that our measure is monotonically related to social welfare losses.

In the model, student outcomes are affected by a one-dimensional measure of diversity, as given by equation (2). Given that Chicago used to employ racial quotas, it seems most natural to measure

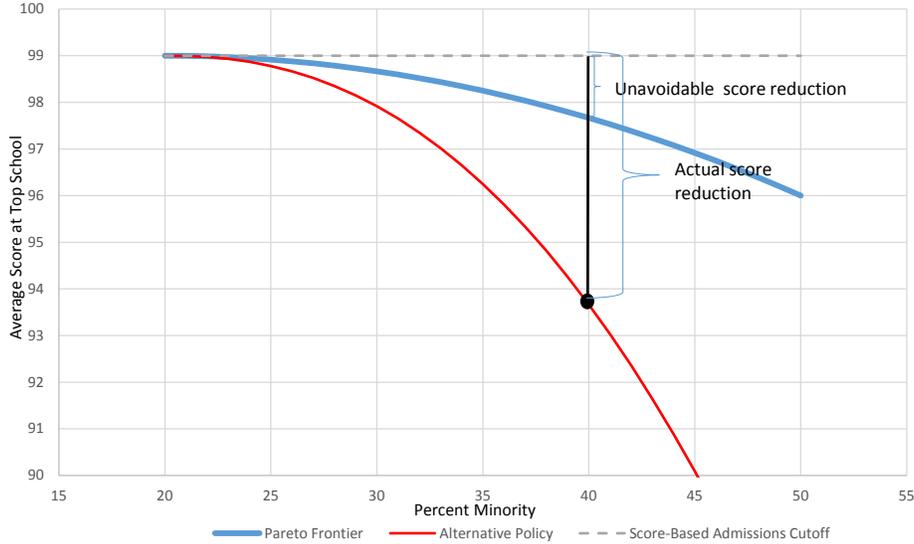


Figure 2: An illustration of the Pareto frontier and a 25% efficient policy

diversity in terms of black and Hispanic representation. However, Chicago’s switch to a socioeconomic-based definition of disadvantage may suggest that diversity should be measured in terms of income. In that case, it seems natural to measure diversity in terms of the fraction of students obtaining subsidized school lunches, a proxy for income. In either case, we can compute an efficiency measure as in Figure 2 with the corresponding one-dimensional measure of diversity. We also will consider situations where diversity is simultaneously measured on two distinct dimensions. In this case, the Pareto frontier becomes three dimensional, with composite scores compared to the two dimensions of diversity.

3 Chicago’s Pareto Frontiers

3.1 Data

Our primary data consists of application files from Chicago Public Schools from the 2013-2014 year and a separate file containing the factors used to compute the tier for each census tract. Each record includes students’ composite scores and their underlying components (7th grade GPA, 7th grade standardized test score, and admissions exam score), school choices, tier, and the final assignment. The application files also indicate whether a student qualifies for special education programming;

because those applicants are assigned through a different process, we exclude them from our analysis sample. Our final analysis sample has 16,818 students and 77,051 student choices, implying that the average applicant ranks about 4.6 schools. We rescale the composite admission scores used by CPS to a 0 to 100 scale that can be roughly thought of as corresponding to a student’s national percentile. For example, an applicant would receive a score of 98.9 if she achieved a score CPS deemed to be at the national 98th percentile on their admissions test, had 97th percentile scores in both English and Math on the ISAT, and had a perfect middle school GPA.¹³

In the 2013-14 school year, CPS used six factors to place neighborhoods into tiers: (1) median family income, (2) percentage of single-parent households, (3) percentage of households where English is not the first language, (4) percentage of homes occupied by the homeowner, (5) adult educational attainment, and (6) average ISAT scores for attendance-area schools. Based on these factors, each census tract is given a socioeconomic score. Tracts are ranked by these scores, and then divided into four tiers, each with approximately the same number of school-age children. Tier 1 tracts are the most disadvantaged, while Tier 4 tracts have the highest scores. The leftmost map in Figure 3 illustrates the distribution of tiers throughout the city. Many of the highest-tier census tracts are at the northern or western edges of the city or along the lake north of downtown. CPS first allocates 30% of the seats at each school solely via composite score and then allocates the remaining 70% of seats with an equal split among each of the four tiers using composite score within tier. For all of our analysis, we follow Chicago and implement the student-proposing deferred acceptance algorithm, assuming that open seats precede tier seats and within seat type, students are ordered based on their composite score.¹⁴ The first two columns of Table 2 show that exam applicants are positively selected compared to CPS 8th graders: their standardized test scores are well above average and they are less likely to qualify for a subsidized lunch. They are also more likely to be female and a little less likely to be underrepresented minorities.

In 2013-14, Payton and Northside had the most competitive admissions criteria, and these two schools are the focus of our empirical investigation.¹⁵ The right two maps in Figure 3 show the

¹³The percentiles are not calibrated to the same population, and a full 100 GPA points are given to any student with straight As in 7th grade so the interpretation of scores as percentiles is not precise.

¹⁴Dur, Pathak, and Sönmez (2016) show that among tier-blind precedence policies, which specify the order in which seats are processed, Chicago’s policy is most favorable to applicants from tier 1. We use this implementation in all exercises in this paper.

¹⁵Throughout the paper, statistics refer to students admitted to Payton and Northside though the main admissions process. They will not exactly match other measures of the schools’ demographics because they do not include a small number of special education and principal’s discretion students who are admitted separately. To compute the allocations produced by alternative admissions policies, we assume that all schools employ those policies when we run the student-

schools' locations and where their students live. The schools are somewhat differently situated, so they provide an opportunity to examine how race-neutral affirmative action plans fare in different environments. Payton is closer to the center of Chicago and enrolls students from across the city. Northside is a relatively attractive location only for students from the northern parts of the city, which are primarily tier 3 and tier 4 neighborhoods. For this reason, there may be greater efficiency costs if an affirmative action plan fails to offer admission to the low-income and/or minority students living in these neighborhoods.

The third and sixth columns of Table 2 show that applicants to Payton and Northside are somewhat positively selected compared to the broader pool of exam school applicants, since they have higher composite scores. There are also fewer black and more Hispanic applicants for Northside compared to the overall applicant pool. The model in Section 2 showed that proxy-based plans must lower minority admittance or lower composite scores compared to racial quotas. At Payton, the tier plan cuts the fraction of black admits nearly in half, but leaves composite scores unchanged compared to the racial quota. At Northside, the tier plan increases the black and Hispanic share, but results in lower composite scores than the racial quota. Across both schools, 124 different students, or roughly one quarter of admitted students, are admitted when the affirmative action system changes.¹⁶

3.2 Unrestricted Pareto frontiers

We start by presenting Pareto frontiers at Payton and Northside when there are no restrictions on using race and/or free-lunch status in the admissions process. These graphs illustrate the extent to which average tests scores must be reduced in order to boost minority or subsidized-lunch representation to various levels.

3.2.1 Trade-offs between composite scores and minority representation

We first examine the Pareto frontier in the dimension of average entrance scores vs. minority representation. The left panel of Figure 4 graphs the highest average entrance score that can be obtained proposing deferred acceptance algorithm.

¹⁶We implemented racial quotas at Payton and Northside by setting separate black, Hispanic, and white/Asian quotas to be equal to the average share of black, Hispanic, and white/Asian students from the three years prior to the adoption of race-neutral admissions. The share of underrepresented minorities was much lower at Northside than at Payton. The 1980 consent decree required that no more than 35 percent of a school be white. Northside had a substantially larger Asian enrollment than Payton.

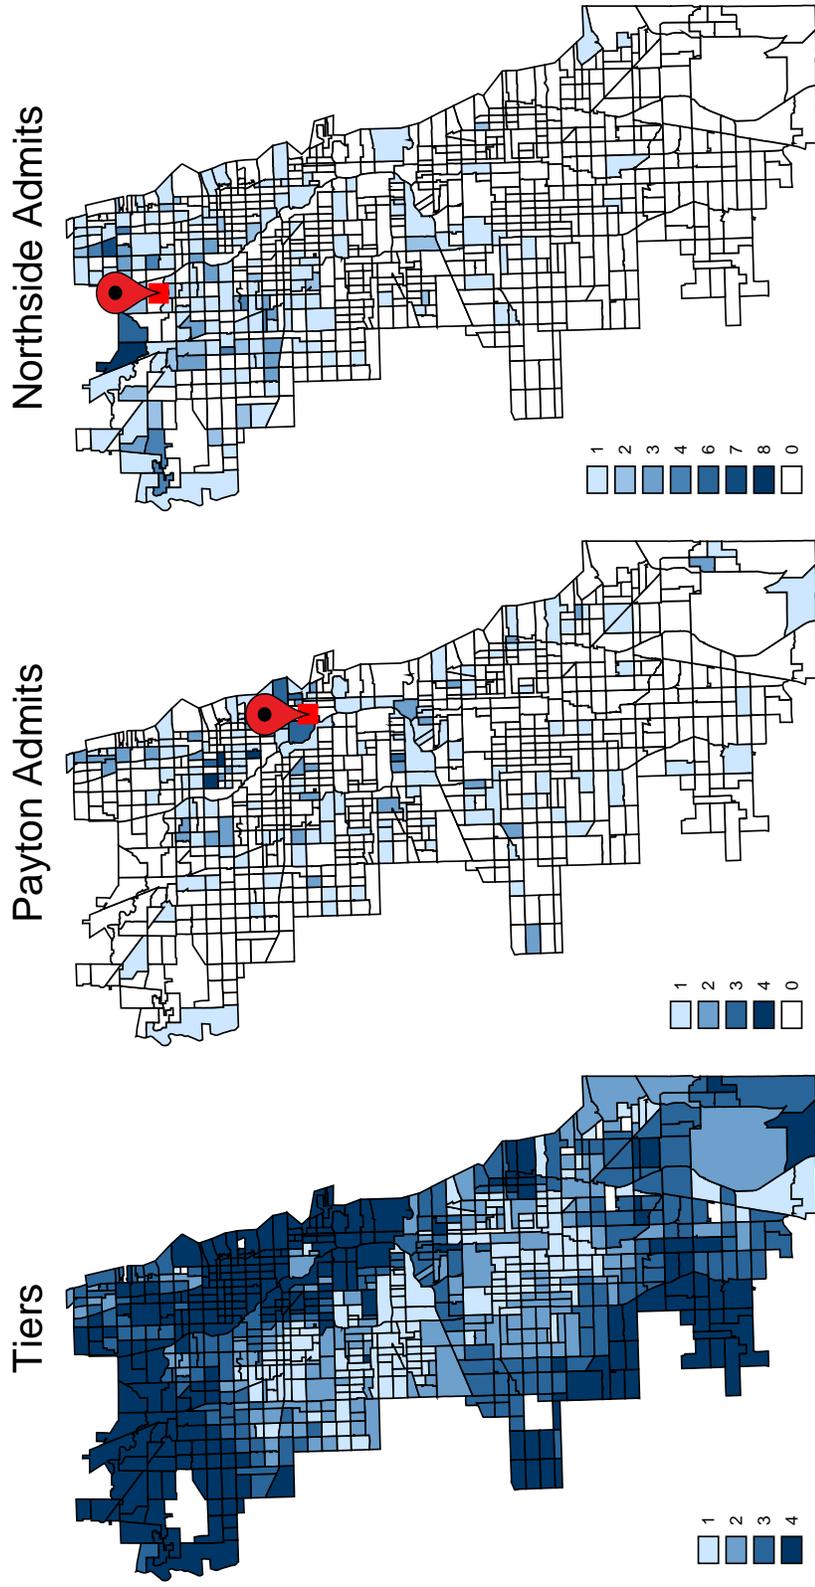


Figure 3: Chicago Tiers and Admitted Students at Payton and Northside

	CPS 8th Graders (1)	Exam Appli- cants (2)	Appli- cants to Payton (3)	Offered Payton 70% tier (4)	Racial quota (5)	Appli- cants to Northside (6)	Offered Northside 70% tier (7)	Racial quota (8)
Black	0.42	0.39	0.37	0.15	0.28	0.26	0.06	0.05
Hispanic	0.45	0.42	0.44	0.22	0.24	0.50	0.29	0.23
Female	0.50	0.55	0.55	0.65	0.67	0.57	0.65	0.65
Subsidized lunch	0.86	0.74	0.71	0.25	0.26	0.71	0.34	0.24
Composite Score		62.2	65.9	98.0	98.0	66.9	97.7	98.9
GPA		61.8	65.9	99.4	99.3	67.4	99.2	99.9
Standardized Test		75.9	78.3	97.6	97.5	78.7	97.5	98.2
Admissions Test		48.8	53.7	97.0	97.1	54.6	96.6	98.7
Observations	27,944	16,818	10,549	220	220	8,274	259	259

Table 2: Descriptive Statistics on Applicants

within each school for each level of minority representation on the x-axis. The blue circles are for Payton and the red squares are for Northside. The hollow markers correspond to the allocation under the current CPS 70% tier-based plan.

The leftmost points on the two graphs indicate that if the schools admit students purely according to composite scores, the average composite score would be 99.1 at Payton and almost 99.0 at Northside. About 20% of the students at each school would be underrepresented minorities. The flatness of the graphs just to the right of these points shows that initially there are only small differences between the added minority students and the displaced majority students. For example, when using a minority quota to increase the minority representation from 20% to 30%, a typical displaced students might have a 96th percentile scores on the entrance exam, 97th percentile scores and on the English and Math ISAT and a perfect middle school GPA, whereas the newly admitted minority students would be students with test scores like 95-95-95 and a perfect GPA or test scores like 99-99-99 with a single B in 7th grade. The remainder of the curve illustrates how much steeper the frontiers become beyond this point. Average composite scores at Payton decline to 99.0 at 30% minority, 98.7 at 40%, 98.2 at 50%, and 97.4 at 60%. The Northside frontier has a virtually identical shape with average scores 0.1 to 0.2 lower at any given level of minority representation.

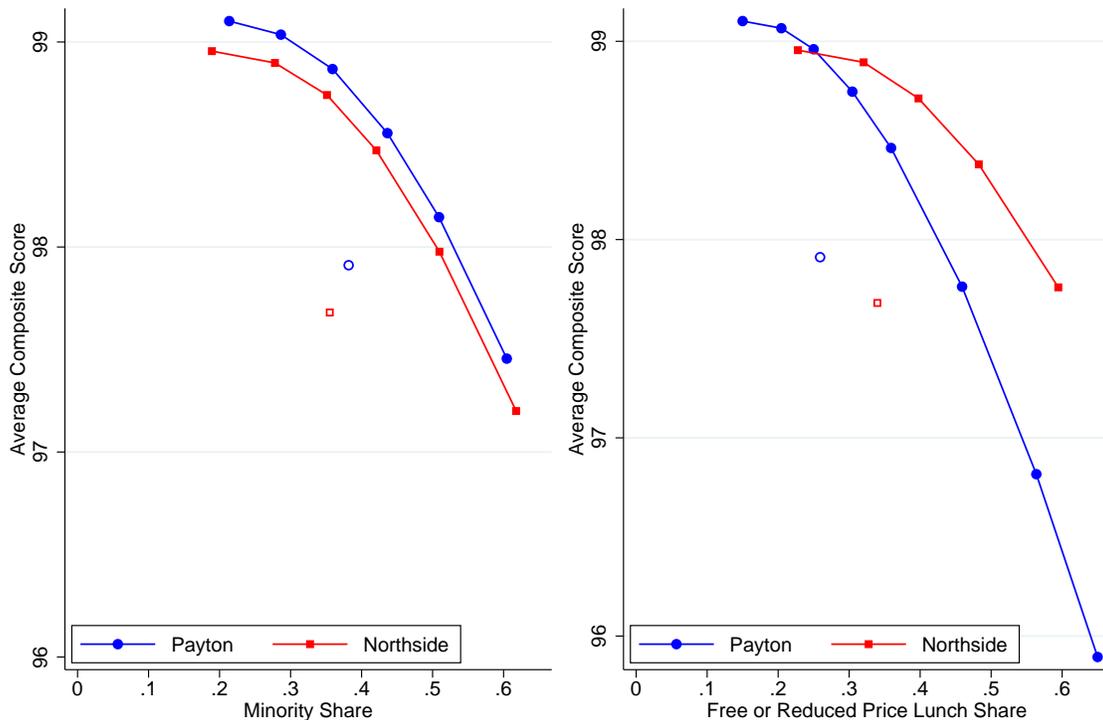


Figure 4: Composite scores vs. minority and free-lunch representation under unrestricted affirmative action

3.2.2 Trade-offs between composite scores and low-income representation

We next turn to trade-offs between average composite scores and the representation of low-income students, as measured by free or reduced-price lunch qualification. The right panel of Figure 4 graphs the maximum average composite score as we vary the number of low-income students. As with the minority affirmative action system, the score-maximizing outcome can be implemented either by setting a subsidized lunch quota or by adding the appropriate number of bonus points to the composite scores of students with subsidized lunches. The shapes of the frontiers are generally similar to those for increasing minority representation. Low-income representation can be increased by about 10 percentage points from its level under a purely score-based policy with only a slight decrease in average composite scores, but the reduction in composite score is more pronounced at higher levels of low-income representation.

One visible difference between the minority and low-income graphs is the differences across schools.

For minority representation, the two curves are very similar, whereas for low-income representation, they are not. The leftmost points on the low-income graphs describe outcomes under a purely score-based admissions policy. Northside has many more students qualifying for free lunch under such a policy (23% vs. 15%). It also seems to be somewhat easier to increase low-income representation at Northside in the sense that the reduction in average composite scores that must be incurred to increase free-lunch representation by any given amount is smaller.

3.2.3 Trade-offs between low-income and minority representation with scores held fixed

School systems may want to maintain both minority and low-income representation at elite schools. We next report on the degree to which these goals are aligned (or in conflict) by presenting slices of three dimensional Pareto frontiers.

In theory, it is easy to present a full three-dimensional Pareto frontier relating minority representation, low-income representation, and average composite scores: given any desired levels m and ℓ of minority and low-income representation, let $f(m, \ell)$ be the highest average composite score that is feasible and graph $f(m, \ell)$. Graphs of surfaces can be hard to visualize, however, so we report a few slices from this surface. Specifically, we suppose that each school wishes to maintain an average composite score close to 98.5, 98.0, or 97.5, and graph the sets of ordered pairs (m, ℓ) corresponding to highest levels of minority and low-income representation that can be *simultaneously* achieved while keeping average composite scores at each level.

The left panel of Figure 5 presents the trade-off between minority and low-income representation for Payton.¹⁷ The small triangle at the bottom-left corner reflects the class composition under a purely score-based admissions policy: the average composite score is 99.1; just over 20% of the students are minorities and 16% are eligible for a subsidized lunch. The three curves then describe the highest levels of minority and low-income representation that can be simultaneously maintained for average composite scores of 98.5, 98.0, or 97.5. The distance between the lower-left triangle and the innermost green curve illustrates that even a modest reduction in the average composite score to 98.5 has a substantial impact on class composition. The rightmost point indicates that minority representation can be increased to 43%, and the leftmost point indicates that one could admit a class in which 35% of

¹⁷Due to the discreteness of the number of applicants and capacities, it is not always possible to exactly have an average score of 98.5, 98.0, or 97.5. We therefore find the minority and free lunch share that most closely approximates these levels of the composite score. Figure 5 reports the average composite score corresponding to each point from our procedure.

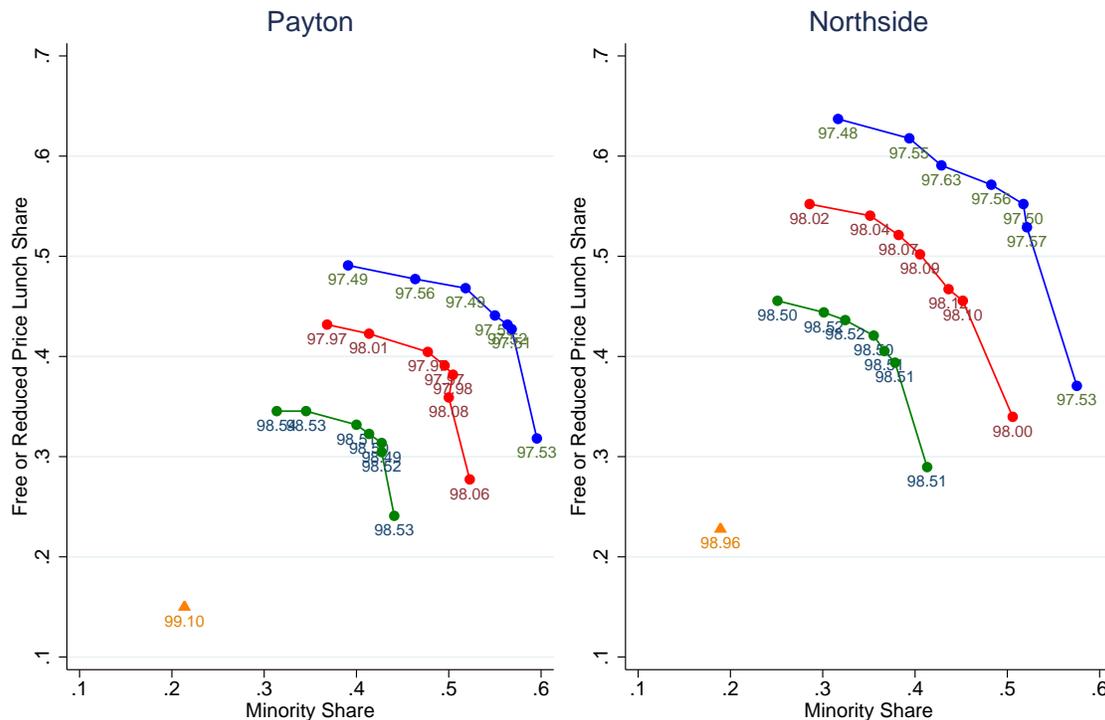


Figure 5: Levels of minority and low-income representation that can be simultaneously achieved at various score levels

students qualified for subsidized lunch. The rightmost points on the three curves correspond to racial preferences of increasing strength. The fact that these three points form an upward sloping curve indicates that pure racial-preference policies have the side effect of increasing low-income representation. But the effect is much less than one-for-one: the policy on the outer curve that achieves nearly 60% minority representation only has 32% of students eligible for subsidized lunch. The leftmost points on the curves correspond to pure low-income preferences. Such policies will similarly have the side-effect of increasing minority enrollment. But again, the effect is much less than one-for-one: a low income preference that reduces the average composite score to 97.5 will increase low-income representation to 49%, but the minority representation is below 40%.

Because the frontiers are extremely curved, a modification of a pure racial-preference that bases a portion of the preference on income rather than race can substantially increase low-income representation while only slightly reducing minority representation. For example, the near verticality of the

outermost curve near its right endpoint means that at average composite scores of 97.5, it is possible to increase low-income representation from 32% to 43% while only reducing minority representation from 59% to 56%. Similarly, the fact that curves are nearly horizontal at the left endpoints indicates that a modification of a pure income-preference that replaces a portion of the income preference with an explicit consideration of race can substantially increase minority representation while having no effect on average scores and almost no effect on low-income representation. For example, one can raise minority representation at Payton from 39% to 46% while holding the average score at Payton fixed at 97.5 and reducing subsidized lunch representation just from 49% to 47.5%. Together these two features suggest that policymakers who value both racial and income diversity would presumably have a strong preference for policies that explicitly considered both race and income.

The right panel of Figure 5 presents a comparable graph for Northside, with the orange triangle indicating the class composition under a purely score-based admissions policy where the average composite score is 99.0: 19% of students are minorities and 23% receive subsidized lunch. The three curves again show low income-minority combinations that are feasible if one wants to keep the average composite score at 98.5, 98.0, or 97.5. The distance of the innermost curve from the orange point is another illustration of the scope for altering class composition while reducing average composite scores just from 99.0 to 98.5. The fact that the entire inner curve is above and to the right of the triangle corresponding to the current class composition reflects that once again racial preferences will have the side effect of raising low-income representation and vice versa. As with Payton, the fact that curves are flat near the left endpoints and quite steep near the right endpoints suggests that policy-makers would presumably prefer an interior point. One clear contrast with Payton is that the Northside curves are located at higher y -values. This again reflects that (contrary to what one might have guessed given Northside's more suburban location) there are actually more high-scoring, low-income students who want to attend Northside.

Implementing a policy that achieves an interior point requires making explicit use of both free-lunch eligibility and minority status. Assuming this is possible, policies can be implemented using simple bonus schemes. Students have x bonus points added to their score if they qualify for free/reduced lunch and y bonus points added if they are a minority. Assignments are then made simply by accepting students in the order of their bonus-adjusted scores. The ratio x/y as determines where we end up between the fewer minorities/more low income or the more minorities/fewer low-income end of each curve. The magnitudes of the bonuses will determine whether we end up on a curve with a higher or lower average composite score. For example, in the case of Payton, the rightmost (52% minority, 28%

subsidized lunch) point on the average-score-98 curve is obtained by giving 6.9 points for minority status and no points for being low-income. The leftmost (37% minority, 43% subsidized lunch) point on the same curve is obtained by giving 7.4 points for low-income status and no points for being a minority. The point two points to the right on the same curve where minority representation is substantially higher (48%) and subsidized lunch representation only a little lower (40%) is obtained by giving 2.8 points for minority status and 5.3 points for subsidized lunch eligibility.

4 How Efficient Are Chicago’s Policies?

4.1 The policy as a tool for racial integration

Chicago’s affirmative action system may be the most sophisticated race-neutral K-12 assignment policy that has replaced an explicit race-based policy. It is therefore of interest to explore whether the CPS policy is an effective substitute for a race-based policy. We measure “effective substitute” by comparing the admitted classes at Payton and Northside with classes that could have been admitted if there were no restrictions on the form of affirmative action.

We begin by examining the effectiveness of CPS’s policy as a means of increasing minority representation. The left panel of Figure 6 reports on Payton. The red circles are the same score-minority Pareto frontier as in Figure 4. The blue triangles show the average scores and levels of minority representation that are realized when one implements policies that reserve various fractions of the seats at Payton to be allocated evenly across the CPS tiers. The current CPS policy corresponds to the larger black square. It involves a 70% reservation and produces a class that is 37% minority and has an average composite score of 98.0.

The CPS policy is nearly efficient at low minority shares. The upper left points on the two curves are by construction identical: reserving zero seats for each SES tier is equivalent to reserving zero seats for minorities. A feature that is not “by construction” is that the portions of the curves close to this point are nearly coincident. For example, the CPS policy can increase minority representation from 20% to 25% while maintaining about the same average composite score as under the race-based policy. At these levels, it is hard to see the gap between the two curves.

Policies based on the CPS tiers are far from efficient, however, when they are used to increase minority representation to anything close to its former level. Most notably, we calculate that the current CPS policy is 24% efficient at Payton. To understand this number, recall, that we defined

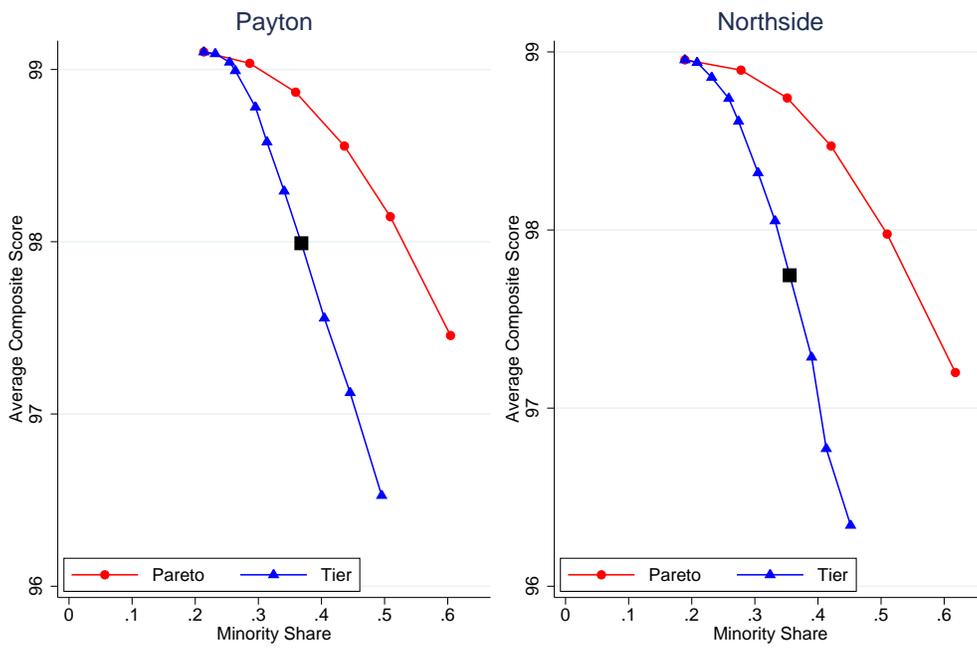


Figure 6: Feasible composite scores at various levels of minority representation: the effect of restricting affirmative action to the CPS race-neutral policy

the efficiency of an affirmative action policy as the fraction of the unavoidable score decline. Looking vertically above the black square, it's worth noting that the policy has reduced average composite scores relative to a pure score-based policy by about 1.1 points (from 99.1 to 98.0) whereas a racial quota policy could have achieved the same level of minority representation while only reducing scores by slightly less than 0.3 points (from 99.1 to a little above 98.8). Hence, only 24% ($\approx 0.3/1.1$) of the decrease in scores that the policy entails was necessary to achieve that level of minority representation. An alternate way to think about the efficiency costs is to look horizontally in the figure. The policy produces a class that is 37% minority and has an average composite score of 98.0. A race-based policy that reduced the average composite score to 98.0 could have yielded a class that is 52% minority.

The right panel of Figure 6 presents a similar graph for Northside. The CPS policy is only 20% efficient as a method for increasing minority representation. The reduced efficiency reflects two features of the curve showing the effects of allocating more seats by CPS tier. Near the upper left endpoint of the graph, the initial flat segment is almost nonexistent. Increasing minority representation from 20% to 30% requires a drop in the average composite score that is more than 50% larger than was needed at Payton. The Northside graph also becomes more concave at high levels of minority representation reflecting that it becomes increasingly difficult to add minority students. Indeed, no policy based on the CPS tiers, not even increasing the tier reservation to 100%, will get minority representation much above 45%. Intuitively, the CPS policy may do worse at Northside because many of the strong minority applicants to Northside live in the middle class neighborhoods that surround the school. A system that uses residence in a low-SES census tract as a proxy for minority status will do relatively well when minority candidates are segregated into low-SES communities and relatively poorly when minority candidates are geographically dispersed.

4.2 The policy as a tool for socioeconomic integration

The CPS affirmative action policy can be motivated as a tool to achieve socioeconomic as well as racial integration. We turn to an investigation of the efficiency with which the policy increases the number of low-income students attending Payton and Northside. The left panel of Figure 7 contains three graphs related to potential class compositions at Payton. The blue circles represent the Pareto frontier. For each desired level of representation of students eligible for free/reduced lunch, they graph the corresponding highest average composite score. The green triangles show CPS's tier policies varying the fraction of seats for students in each SES tier.

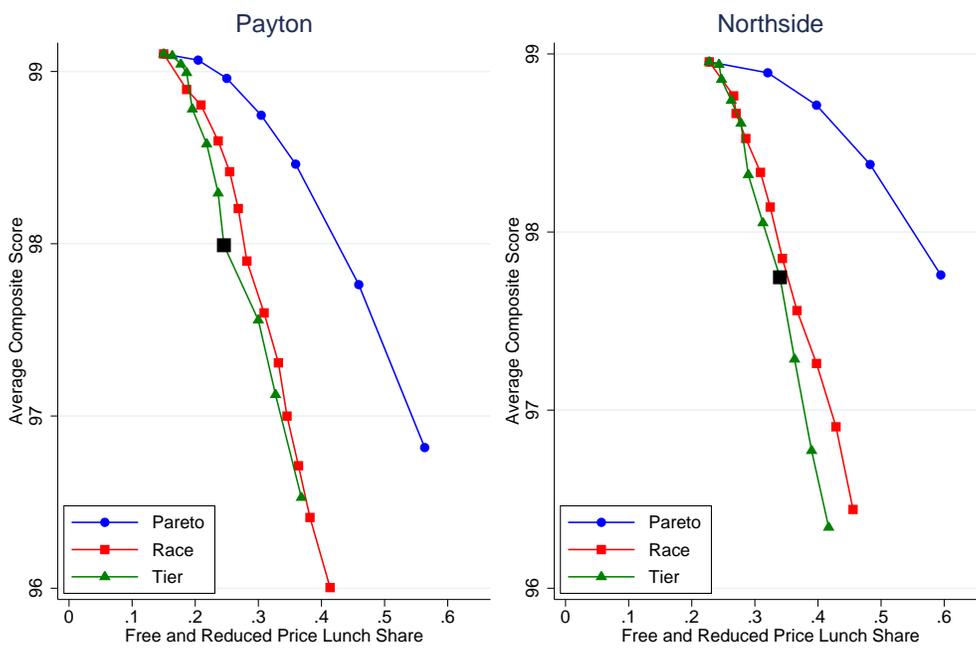


Figure 7: Feasible composite scores at various levels of free/reduced lunch eligibility: CPS and minority quotas vs. the Pareto frontier

Despite its focus on socioeconomic disadvantage, the CPS tier plan is not very efficient as a method for increasing the number of low-income students. At Payton, for example, the current CPS policy increases the proportion of students on free/reduced lunch from 15% to 25% at the cost of reducing average composite scores from 99.1 to 98.0. A simple free/reduced lunch quota could have achieved the same increase in subsidized lunch representation while maintaining an average composite score of 99.0. Comparing the drops with more precision, our efficiency metric says that the CPS policy is only 14% efficient as a method of increasing low-income representation.

Because low-income representation was clearly not CPS's only goal, it is not damning that the CPS policy is not highly efficient on this dimension. One of our more striking observations, however, is that the SES-based CPS policy is actually worse than a racial quota policy as a method for increasing the representation of low-income students. The red squares in the graph illustrate the effect of minority quotas on low-income representation: we again plot the percent free/reduced lunch and the average composite score that results from using minority quotas of between 20% and 70%. For example, we see that a racial quota set so that it produced a class with an average composite score of 98.5 would have yielded a class in which 25% of students received subsidized lunch. The CPS policy is bringing in no more subsidized lunch students despite reducing the average composite score to 98.0.

The right panel contains comparable graphs for Northside. The Northside graph starts further to the right because Northside has more free-lunch students under a purely score-based admissions scheme. The CPS tier scheme is again less effective here than it was at Payton: our efficiency metric says that it is just 10% efficient. The CPS tier vs. race-based comparison also comes out the same way: CPS tier system is admitting fewer subsidized lunch students at each level of the average composite score than would a minority quota.

The fact that CPS's SES-based policy is less efficient than racial quotas as a means to increase low-income representation may at first be surprising: CPS's SES index is based on several variables known to be highly correlated with poverty. However, once we learn that a student from a low SES census tract has an extremely high composite score, it means that the student is very unusual for their census tract. This suggests that their demographic characteristics are different from the census tract mean. In contrast, under the old policy, a high-scoring minority student is less likely to be poor than a randomly selected minority student, but we still know that they are a minority student. Apparently, when one conditions on having high grades and standardized test scores, minority status retains more of its power as a predictor of low-income status than does living in a low-SES census tract.

4.3 Within-School Heterogeneity and Achievement Gaps

A primary motivation for elite public schools is that they make it possible to offer curricula tailored to the knowledge and abilities of a city’s most advanced students. Within-school heterogeneity in preparation/ability reduces this potential benefit. Our focus on how affirmative action reduces average test scores reflects this concern – the average score at an elite school declines as low-scoring students are admitted. In this section, we present more direct analyses of within-school heterogeneity.

4.3.1 Within-school score heterogeneity

Figure 8 reports on within-school heterogeneity across admissions policies. The top three histograms illustrate how the distribution of composite scores within Payton changes as admission shifts from purely score-based to a 20% minority set-aside (which produces a class that is 35% minority) and then to a 40% minority set-aside (producing a 50% minority class). In a purely score-based system, there would be little heterogeneity in composite scores at Payton: over 60% percent of students would have scores of at least 99, most of the rest would have scores of at least 98, and everyone would have scored at least 97. With a 20% set-aside, students with scores between 97.1 and 98.5 are replaced with a students with scores between 95.4 and 97.0. Increasing the set-aside to 40% as in the bottom panel causes an even more notable increase in heterogeneity. On the high end, most students with scores below 99 have been displaced and the bottom tail extends below 93.

The bottom histogram on the left-side of Figure 8 illustrates the distribution of composite scores at Payton under the current CPS tier policy. The most striking feature is that the left tail of low-scoring students now extends down to 89. Concretely, this means that teaching and curricula most accommodate substantially more heterogeneous student preparation. Note also that this occurs despite the fact that the class is only 37% minority – a racial composition much closer to that with the 20% set aside than to that with the 40% set aside. The difference also helps explain why the curve of class compositions feasible under the CPS tier approach is so steeply sloped – replacing an extra unrestricted slot with a tier 1 slot replaces a student with a score of around 99 with a student with a score of around 89.

Northside shows a roughly similar pattern as Payton across admissions policies. At Northside, the lower tail extends down to 95.3 (vs. 95.4 at Payton) with a 20% minority set-aside and to 93.1 (vs. 92.9) with a 40% set-aside. The CPS policy is substantially increasing within-school heterogeneity compared to both quota levels. Indeed, the fact that the CPS policy is less efficient at Northside than

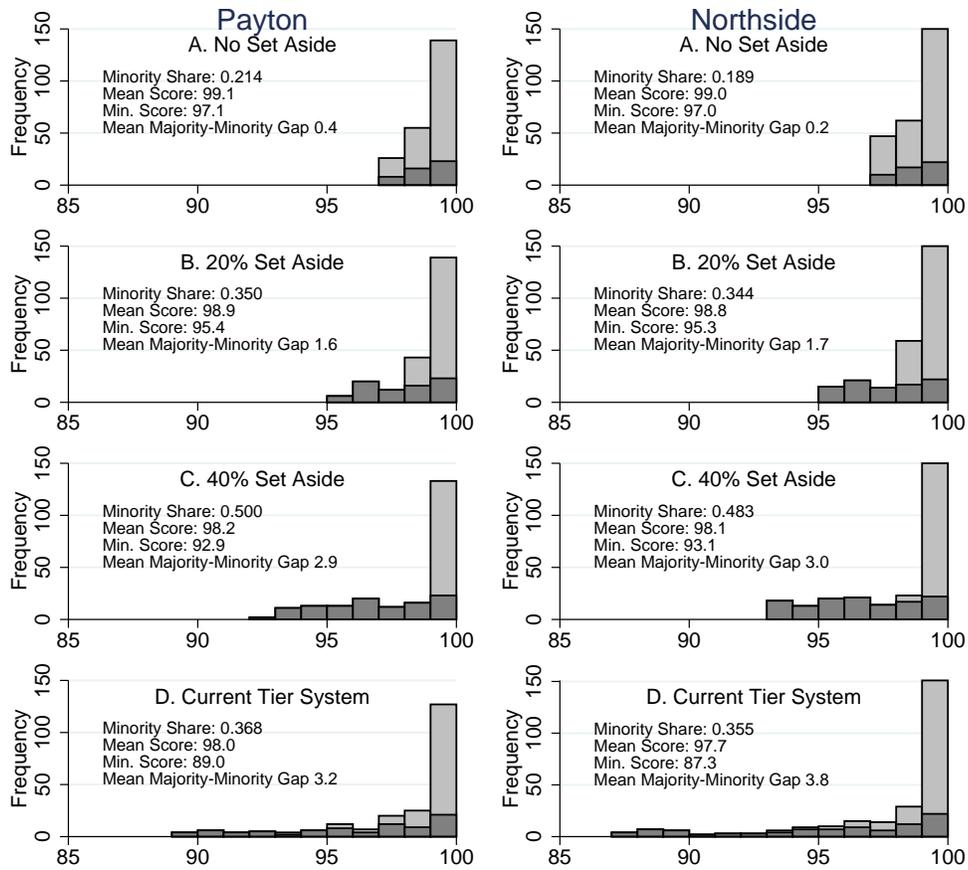


Figure 8: Within-school score distributions under alternate admissions policies

at Payton is visible in the greater extent of the lower tail. It extends down to 87.3 at Northside versus 89.0 at Payton.

4.3.2 Within-school achievement gaps by race

A separate concern about affirmative action policies is that they may contribute to within-school achievement gaps that adversely affect minority students for reasons, e.g. stereotype formation, that we did not try to incorporate into our model (see, e.g., Steele and Aronson (1995) and Austen-Smith and Fryer (2005)). In this section, we illustrate how the gap in composite scores is affected as more seats are reserved for minorities under a quota system and with CPS's shift to a race-neutral policy.

The histograms in Figure 8 were designed to also illustrate the gap in composite scores. Each histogram bar is divided into two portions: the lower dark gray portion reflects the number of under-represented minority students with composite scores in the band and the upper lighter portion reflects the number of white or Asian-American students. The upper histograms make clear that within-school achievement gaps would be small with no set-aside. All admitted students would then have scores of at least 97, and there would be minority students in all three bands. The second row of histograms, corresponding to a 20% minority set aside, show a larger score gap. All majority students would have scores above 98 under such a policy, whereas the minority students scores would be more spread out with a median around 97.5. The gaps widens at a 40% minority set-aside. Almost all of the majority students would then have scores of at least 99, whereas the minority scores would be fairly evenly spread over the range from 93 to 100.

Switching from a racial quota to CPS's tier-based affirmative action affects within-school achievement gaps in two opposing ways. One effect that works to reduce racial gaps is that the policy admits some relatively low-scoring white and Asian-American students who live in low SES neighborhoods. Working in the opposite direction, however, is the fact that a number of high-scoring minority students living in medium to high SES neighborhoods are now being denied admission and that almost all of the lowest-scoring admits (all of whom come from tier 1) are minorities. Given the opposing effects, it is not a priori clear how the shift to a race-neutral policy will affect within-school gaps. Empirically, the answer turns out to be that the shift increases the within-school racial gaps. Under a 20% minority set-aside, the majority-minority score gap in average scores is 1.6 points at Payton and 1.7 points at Northside. Under the CPS tier-based plan, it is 3.2 points at Payton and 3.8 points at Northside. (As before, the larger gap at Northside reflects that the CPS plan is less efficient than at Northside.) A comparison of the histograms in the second and fourth rows of Figure 8 brings this out in more detail.

At both schools, the shift to the race-neutral plan reduces the number of majority students with scores in the 98 to 99 range. We do add a visible number of majority students with scores between 95 and 97, but almost all of the added students with scores below 95 are underrepresented minorities.

Within-school majority-minority achievement gaps are larger under the tier-based policy for any fixed level of minority representation. This fact can be seen in Figure 9, which plots the difference between the average composite score of admitted majority and minority students for each level of minority representation. The solid lines illustrate racial quota policies varying the numbers of reserved slots, with blue circles for Payton and red triangles for Northside. At the left endpoint, the curves have a value of 0.2, which reflects the small difference in average scores of admitted majority and minority students under a purely score-based admissions process. Achievement gaps grow by about 1 point for every 10 points of minority share: they are slightly above 1 point when the school is 30% minority, around 2 points when the schools are 40% minority, and around 3 points at 50% minority. The dashed lines give corresponding numbers for policies based on the CPS tiers: for each desired level of minority representation, we find the CPS tier reservation percentage that produces that level of minority representation and report the majority-minority score gap under that policy. At Payton, the majority-minority gaps are small at first, but grow once enough seats are allocated by tier to increase minority representation beyond 30 percent. The gaps grow even more quickly at Northside, again illustrating that the lower efficiency of the CPS plan in the Northside environment is associated with a larger majority-minority score gap. Compared to a racial quota, the CPS policy roughly doubles the magnitude of the majority-minority score gap at Northside.

5 How Efficient Can Race-Neutral Policies Be?

So far we have compared the CPS plan with race-based policies which would probably be infeasible in the current legal/political environment. Hence, the “inefficiency” we have discussed should be thought of as measuring costs CPS incurred when eliminating racial preferences. It does not measure losses relative to alternate policies CPS could have adopted. In this section, we explore what else CPS could have accomplished in a race-neutral manner. We start by examining predictors of minority status among high achievers.

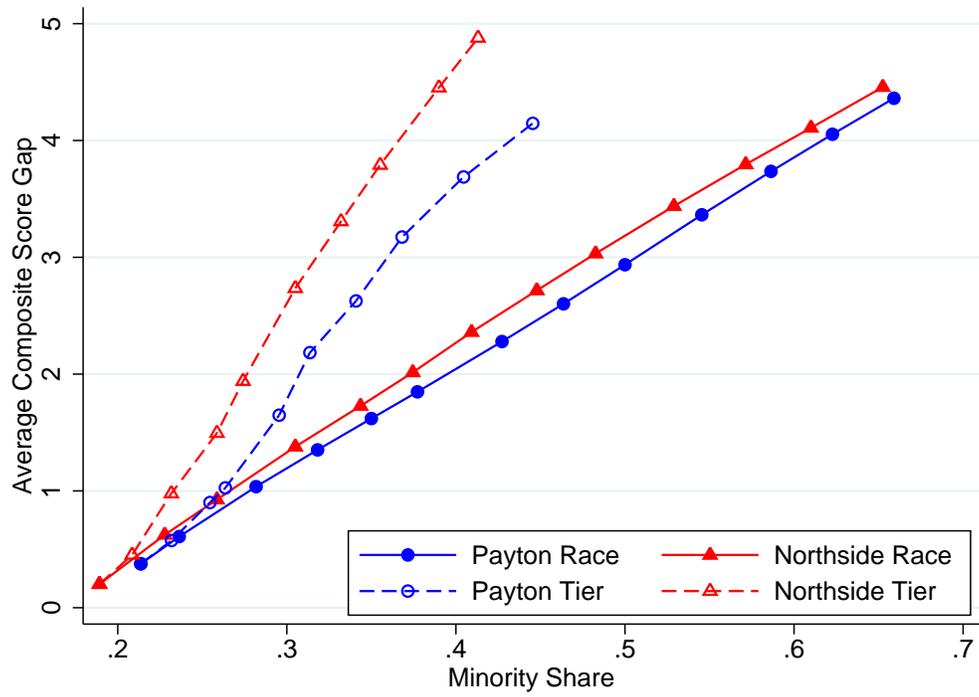


Figure 9: Within school gaps in average composite scores: quotas vs. CPS tier-based selection

5.1 Correlates of underrepresented minority status

An SES-based affirmative action plan increases minority representation if there are more minorities among the students it advantages than among the students it displaces. A complete analysis of any policy is complex, but it is reasonable to expect that proxy-based policies will be effective as means to increase minority representation if the proxies are highly correlated with minority status *within the population of applicants with scores close to the admissions margin*. Although we know a great deal about demographic correlates of minority status, it is unclear whether this knowledge is relevant for our current purposes: when we see a student from a low-SES neighborhood obtain an extremely high score on an admissions test, we learn that the student is very unusual for their neighborhood. Therefore, we need to change our posteriors about the student’s race/ethnicity. In this section, we examine correlates of minority status for high-scoring students.

5.1.1 The CPS demographic variables

Each of CPS’s six tract-level variables is correlated with an applicant’s minority status. The first column of Table 3 presents coefficient estimates from six OLS regressions run on the full dataset of students applying to CPS’s exam schools in 2010-2012. The dependent variable in each regression is an indicator for being an underrepresented minority. Each regression has a single explanatory variable: one of the SES indicators in the current CPS formula. Each of these variables is scaled as a percentile (between 0 and 1) within Chicago’s census tracts.¹⁸ Higher percentiles correspond to what CPS regarded as being of higher SES, e.g. having higher median household income or a higher percentage of two-parent families. Five of the six variables are positively correlated with minority status because the coefficient estimates are negative and significant. The native English speaker variable is not.

The second column of Table 2 restricts the regressions to a subsample more relevant to Payton and Northside admissions: the set of applicants with composite scores of at least 96. Here, all six variables are correlated with underrepresented minority status. Of course, whether variables are individually correlated with minority status is different from whether one would want to weight them positively in an index. The third column presents estimates from an OLS regression of the minority indicator on all six variables. The results suggest CPS’s equally-weighted index may be quite different from the index that would be most aligned with minority status: there are substantial differences

¹⁸For a few census tracts, CPS’s version of Home Ownership Percentile does not exactly correspond to a conversion of their Home Ownership variable to percentiles. The estimates in Table 2 are based on CPS’s Percentile measure. We have re-estimated the regressions in Table 2 by converting Home Ownership to a percentile and found very similar estimates.

in the coefficients on most of the variables. Most strikingly, the coefficient on home ownership is both positive and statistically significant. This suggests that CPS’s inclusion of this variable may disadvantage minority students. Chicago has some predominantly black middle class neighborhoods, such as Washington Heights, in which home ownership is high and some more affluent areas, including parts of Lakeview, the Loop, and the Near West Side, with mostly rental housing. Apparently, such examples are sufficiently common so that including home ownership on top of the other variables can disadvantage minorities.

	Dep. Variable: Applicant is Minority					
	(1)	(2)	(3)	(4)	(5)	(6)
Median Income Percentile	-0.375 (0.006)	-0.657 (0.042)	0.124 (0.076)			
Adult Education Percentile	-0.380 (0.006)	-0.719 (0.040)	-0.486 (0.078)	-0.419 (0.043)		-0.240 (0.047)
Two Parent Percentile	-0.520 (0.006)	-0.919 (0.043)	-0.728 (0.081)	-0.525 (0.069)		-0.523 (0.068)
Native English Percentile	0.094 (0.006)	-0.147 (0.056)	-0.045 (0.077)			
Home Owner Percentile	-0.283 (0.006)	-0.339 (0.043)	0.351 (0.059)			
ISAT Percentile	-0.526 (0.006)	-0.760 (0.041)	-0.324 (0.064)	-0.213 (0.062)		-0.209 (0.061)
Sum of Tract Characteristics					-1.101 (0.057)	
Free or Reduced Price Lunch						0.212 (0.024)
Sample	Exam	Apps	Apps	Apps	Apps	Apps
	Apps	≥ 96	≥ 96	≥ 96	≥ 96	≥ 96
Observations	44,962	1,819	1,819	1,819	1,819	1,819
R-Squared			0.268	0.244	0.171	0.274
Misclassification			21.7%	23.1%	27.2%	22.3%

Table 3: Predictions of Minority Status

It is unclear whether it would be legally or politically viable for CPS to adopt the predicted value from the regression in column 3 as its measure of tract SES: most obviously, the measure could be criticized for explicitly penalizing students for living in neighborhoods with lower median household

incomes. Accordingly, we report in the fourth column a related regression in which first dropped the three variables that had “wrong sign” coefficient estimates in the third column, and then dropped one additional variable that had a wrong sign in this regression. There is some loss in R^2 from making the formula immune to this criticism. The fifth column reports estimates from a regression that uses only the CPS’s equally weighted sum of the six tract characteristics. Note that the R^2 from this regression is yet lower, only 0.17 compared to 0.24 in the previous column. This suggests that modifying the SES index to place more weight on some demographic variables and less on others is worth exploring as a potential means to increase efficiency.

5.1.2 Student-level free lunch data

The current CPS admissions policy affirmative action uses only tract-level variables. Note, for example, that this means that low-income students who qualify for free/reduced price lunch receive no advantage relative to students from their census tract who do not. It seems natural that one might want to include this variable for the direct benefit of increasing the representation of low-income students. It also seems plausible that a subsidized-lunch variable might be highly correlated with minority status, so its inclusion might also further the goal of increasing minority representation. To examine this hypothesis we added an indicator for free/reduced lunch eligibility to the regression in the fourth column of Table 3. The coefficient estimate on the subsidized lunch variable is positive and significant and increases the R^2 of the regression from 0.24 to 0.27. This is not a very large increase, but still makes policies which include individual-level data on subsidized lunch status (and perhaps other variables) a direction that one might explore as a means to also improve minority representation.¹⁹

5.1.3 Alternate demographic variables

The CPS SES index uses a measure of test scores in the local public school and five tract-level variables taken from census data. The five census variables all seem plausibly related to disadvantage, but so are dozens of other variables available from the census: the fraction of households with income below the poverty level, the fraction receiving food stamps, the fraction of adults not in the labor force, the fraction without health insurance, the share of housing units that are vacant, teenage pregnancy rates, etc. In this section, we explore the extent to which SES indexes using alternate variables would

¹⁹We do not pursue this further in this paper. Instead we restrict our comparisons to alternate policies which like the CPS policy use only tract-level variables. We felt that this helps clarify the relative efficiency of policy designs when implemented with comparable data.

be more correlated with minority status in the population of exam school applicants with high test scores.

Our primary analysis follows a two-step approach. First, we manually browsed through the many thousands of data items available from the census at the tract level and used them to define 145 variables that seemed plausibly related to disadvantage and that we thought might turn out to be predictors of minority status.²⁰ The number of variables is still sufficiently large relative to the number of high-scoring minority applicants to make overfitting an important concern, and it also seems implausible that CPS would choose an SES index with so many components. Accordingly, we use a LASSO regression to pick a parsimonious index.²¹ Specifically, we estimate a LASSO regression on the subsample of applicants with composite scores of at least 96 with minority status as the dependent variable and both the CPS variables and our added 145 variables as potential explanatory variables.²² The regularization parameter was chosen by cross validation.²³

The LASSO procedure yielded a model that uses nine of explanatory variables. Table 4 lists the variables that are used in the index in decreasing order of their weight. The model selects three CPS variables: the fraction of single-parent families, average ISAT scores in the local elementary school, and the measure of adult educational attainment. That the LASSO chose these three of out of the 151 available variables is an indicator that the CPS measure predicts minority status. The largest weight is placed on a variable giving the median value of owner-occupied units. The LASSO index also weights two variables that are somewhat less natural as measures of socioeconomic status: the fraction of the foreign-born population who come from Asia and Europe. All nine variables enter with signs matching the most natural interpretation of the variables as indicators of tract SES.

5.2 More efficient race-neutral policies

We now turn to an evaluation of the extent to which alternative race-neutral policies would be more efficient as a means to increase minority representation. We consider three successively larger modifi-

²⁰Not all variables are available for all tracts. We impute the mean value for missing values.

²¹We also investigated the performance of random forest models. They did not fare quite as well in out-of-sample fit and also seem harder to justify as an index of socioeconomic status so we did not pursue this direction further.

²²All variables were scaled to have variance one in our full 2010-2012 sample.

²³We repeatedly estimated the model on half of the sample and assessed its predictive power on the other half. We chose the regularization parameter so that the number of variables included was as small as possible subject to predictive power being within one standard deviation of the best possible. We then estimated the model using this regularization parameter on the full 2010-2012 dataset. All estimation used the `glmnet` package in R. We have also investigated other methods of variable selection and found very similar results for efficiency.

Variable	Weight
Median value for owner occupied units	-0.455
CPS single parent family	-0.264
Fraction of foreign born from Asia	-0.129
CPS ISAT scores in local elementary school	-0.086
Fraction of foreign born from Europe	-0.079
CPS adult educational attainment	-0.046
Fraction of households that are female head, no husband, and Food Stamps/SNAP	0.036
Fraction of children 3–4 in married couple households	-0.007
Fraction of households with children under 18 who received Food Stamps/SNAP	0.002

Table 4: Demographic variables selected by LASSO

cations of the current CPS plan: (1) a modification that uses the current CPS tract-level SES measure but in a continuous manner; (2) a second that also re-weights the CPS variables; and (3) a third that further modifies the current plan by adding variables as suggested by the LASSO regression.

Each of the alternate policies employs a continuous tract-level *TractDisadvantage* variable. In the case of the policy (1), it is the predicted value of a the model in column 5 of Table 3 which can be thought of as using the same SES measure that CPS now uses but without the artificial discretization into tiers; in policy (2), it is the predicted value from the regression model in column 4 of Table 3; and in policy (3), it is the prediction of the LASSO regression. Each student has a multiple of his or her *TractDisadvantage* variable added to their scores as a “bonus” that can be thought of as reflecting disadvantage:

$$AdjustedScore = CompositeScore + wTractDisadvantage.$$

We then run the student-proposing deferred acceptance algorithm using the *AdjustedScore* as each student’s priority ranking. The weight w on SES can be varied continuously. Larger values of w will lead to lower composite scores and higher minority representation at the top schools.

To make as clean a comparison of relative efficiency as possible, we implemented two versions of each of the above plans. One uses a weight that makes the underrepresented minority share at Payton exactly match its value under the current CPS plan. The other exactly matches the current underrepresented minority share at Northside.²⁴ The first column of the upper panel of Table 5 compares the efficiency of the four plans at Payton using the Payton-matching weights. The current

²⁴With a single school-independent w we cannot simultaneously match the minority percentages at both schools. This would be possible if we allowed a broader class of plans: choosing different weights at each school and then running the deferred acceptance algorithm.

CPS plan is 24% efficient at Payton. The second row examines a plan that bases a continuous bonus on the current CPS SES index.²⁵ This plan is 31% efficient at Payton. The improvement can be thought of as coming from eliminating the arbitrary discretization of *TractDisadvantage* inherent in the CPS plan. The third row examines the plan which uses the re-weighted sum of three of the six CPS variables, dropping the other three. It is 34% efficient at Payton. Finally, the fourth row examines a plan which uses a weighted sum of the nine variables selected by the LASSO model as its SES index. It is 38% efficient at Payton. We take the fact that the LASSO model provides only a moderate improvement to suggest that it will be hard to do much better than our fairly simple re-weighted CPS model in a place-based scheme that relies only on tract-level variables.

The third column reports similar figures for Northside for the implementation of each plan that keeps the underrepresented minority share at Northside constant. The results strengthen our earlier observation that race-neutral policies are less effective in the Northside environment. Not only is the current CPS plan less efficient, but improving on the CPS plan is more difficult there. The efficiency gain from re-weighting the CPS variables is just four percentage points (compared to 10 at Payton) and the LASSO-based plan is just three percentage points better than the CPS plan.

Admissions Policy	Efficiency at Payton	Efficiency at Northside	Largest Bonus at Payton	Largest Bonus at Northside
A. Minority Representation				
CPS Tier 70%	24%	20%	10.1	11.6
Bonus – continuous CPS index	31%	19%	9.2	11.5
Bonus – reweighted CPS variables	34%	24%	10.0	12.4
Bonus – LASSO prediction	38%	23%	10.7	15.6
B. Low-Income Representation				
CPS Tier 70%	14%	10%	10.1	11.6
Bonus – continuous CPS index	20%	14%	8.9	9.6
Bonus – reweighted CPS variables	27%	13%	7.4	10.2
Bonus – LASSO prediction	35%	17%	8.3	12.1

Table 5: Efficiency of alternative race-neutral plans

The third and fourth columns of Table 5 report the size of the largest bonus that was used in each plan. Our definition of the *AdjustedScore* gives some bonus points to all students, so we define this

²⁵Note that the CPS tier quota plan could also be regarded as a bonus scheme where students receive bonus points that depend on their tier although the size of the bonus would differ between Payton and Northside.

to be the difference between the largest and smallest number of bonus points that students receive.²⁶ In most cases the largest bonuses at Payton are about 10 points, which means that some students with composite scores around 90 may be accepted while other students with scores of around 99 are rejected. This will only occur, of course, if there are students with scores of around 90 from very low SES census tracts. The plans at Northside use somewhat larger maximum bonuses.

To provide a further illustration of the power of such schemes to increase minority representation when used more or less aggressively, Figure 10 reproduces Figure 6 with the addition of series that describe the outcomes when implementing the re-weighted CPS and LASSO bonus schemes for various weights w . The left panel graphs the average composite score for each level of minority representation at Payton and the right panel gives the corresponding figures for Northside. At Payton the re-weighted CPS bonus scheme (shown in green) noticeably improves on the current CPS tier scheme at minority shares near the current levels. The performance of the LASSO model provides an additional incremental improvement that is smaller but still clearly visible. At Northside, the improvement from re-weighting the CPS variables is visibly smaller. The failure of the LASSO-based model to provide any incremental improvement holds for a broad range of minority representation levels.

We also repeat this exercise predicting whether a student is eligible for a free or reduced price lunch in the bottom panel of Table 5.²⁷ Of course, whether a student receives free or reduced price lunch is directly observed, so it is not necessary to predict this characteristic from census tract information. However, we treat the exercise as providing insight into how well one can predict other measures of disadvantage that one might want reflected in a socioeconomic based affirmative action system, but which the district does not directly observe. We choose the weight to match the fraction of students currently eligible for subsidized lunches at Payton and Northside, respectively. The estimates in the lower panel of Table 5 show that both reweighted CPS indices as well as the LASSO model provide only a moderate improvement over the CPS Tier formula. Together, the two panels of Table 5 suggest that the difficulty of predicting whether a high-scoring student is a minority using tract-level variables extends to predicting other measures of disadvantage.

²⁶In the case of the CPS Tier 70 plan we define the “largest bonus” to be the difference between the realized cutoffs for students from the highest and lowest tiers.

²⁷Appendix Tables A1 and A2 report the comparable regression table and LASSO selected factors corresponding to Table 3 and Table 4 for free lunch. Using the same regularization criteria as in the LASSO minority model to predict subsidized lunch status selects only three variables. We therefore chose the regularization parameter to minimize misclassification. The procedure results in nine variables, two of which are also selected by the LASSO minority model.

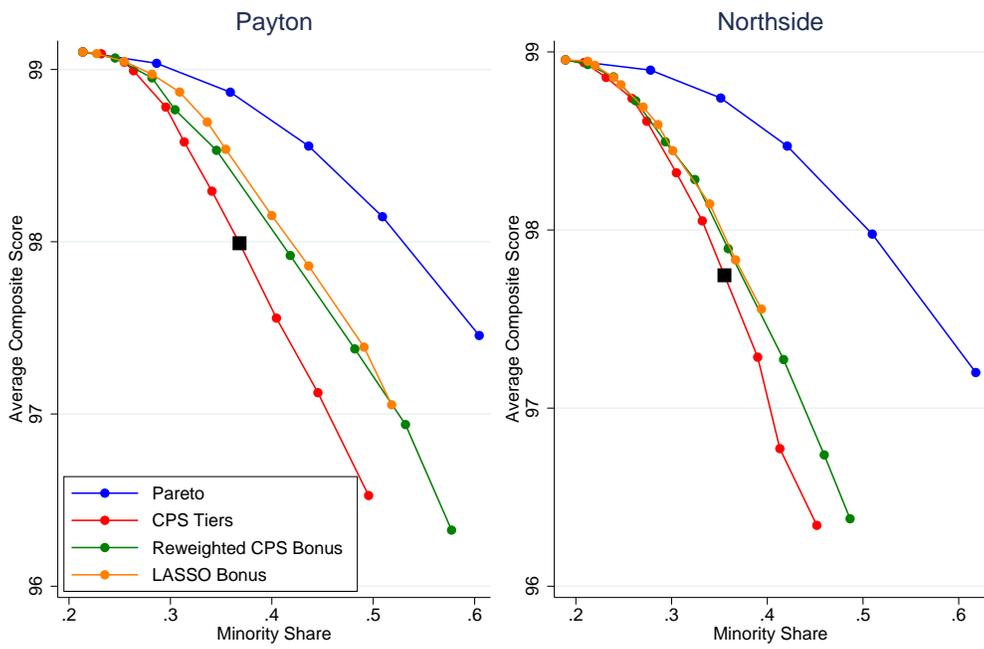


Figure 10: Feasible composite scores at various levels of minority representation for alternate race-neutral plans using bonuses based on tract demographics.

5.3 “Top 10%”-Style Rules

Texas’s “Top 10% Rule” has probably received more attention than any other race-neutral affirmative action policy. In the original 1997 version, all students ranked in a top 10% of a Texas public high school were guaranteed admission to all public universities.²⁸ Similar plans have been enacted or considered in Florida and California. In this section, we examine how such plans might work in the Chicago context.

We consider plans that first fill a fraction Y of the seats at Payton and Northside with the students with the highest composite scores, and then fill the remaining $1 - Y$ fraction of seats using priorities like those in a top $X\%$ rule. A constraint one immediately faces when doing this is that Payton and Northside are much smaller than the University of Texas-Austin: each admits only 220 to 260 students per year. Chicago contains roughly 800 census tracts and CPS operates 478 elementary schools. A plan to admit the top student from each elementary school or from each census tract would therefore be infeasible. It would be particularly infeasible to consider a plan that, like Texas’s, focuses purely on grade-point-averages: each school has many students with perfect GPAs. Motivated by a desire to have a plan that is similar in spirit to a top 10 percent plan, but otherwise as similar as possible to current CPS policies in Chicago, we consider two variants. In each case, the number X needs to be filled in with the (very small) number that exactly fills the $1 - Y$ fraction of seats to be distributed under the top $X\%$ rule.

1. Seats reserved for top $X\%$ students are allocated by census tract. Each student is given a “composite score tract rank” defined to be their rank within the census tract where they live divided by the number of 8th grade students in that census tract.²⁹ Seats are allocated to the students with the lowest composite score tract rank until the reserved seats are filled. Note that when the number of open seats is very small, students from relatively small tracts may not gain admission even if they have a perfect score.³⁰

²⁸More recently the University of Texas at Austin has limited automatic admission roughly to students in the top 8% of their high school class in order to limit the number of students accepted in this way to 75% of incoming in-state students. The remaining students are admitted via a holistic race-conscious policy that is the subject of recent court cases.

²⁹If two students have identical composite scores, which most often occurs with perfect scores, we break the ties randomly. We believe that the rule works better with this tie-breaking than if all students are given an average rank, because under the average-rank system tracts with multiple perfect scorers may get no students admitted to the top schools.

³⁰Some students living in large tracts can also be denied admission despite a perfect score if there are other perfect scorers in their tract.

- Seats reserved for top $X\%$ students are allocated by community area. The city of Chicago is sometimes divided into 77 community areas (which may in turn consist of several named neighborhoods). We assign each student to one community area using the census tract of residence.³¹ We define each student's composite score community rank as the rank within the community divided by the number of 8th graders in the community area and again allocate seats reserved for top $X\%$ students in order of this rank.

Figure 11 illustrates the efficiency of these schemes as a method to increase minority representation. The upper series marked with blue circles is the Pareto frontier. The other two series describe outcomes under the tract-based and community area-based top $X\%$ rules.

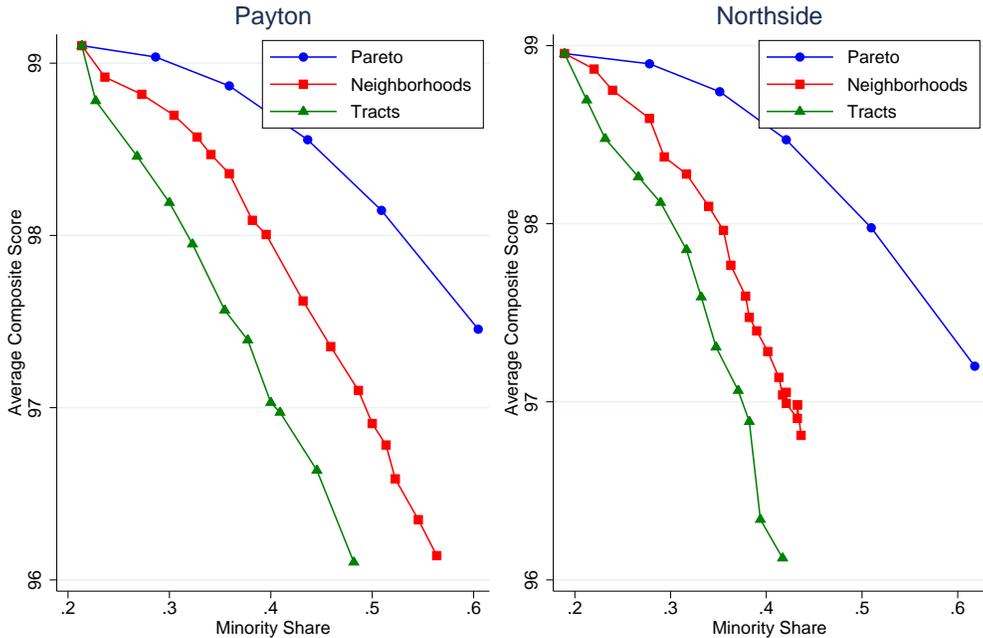


Figure 11: Feasible composite scores at various levels of minority representation: Top 10%-style rules

The lower green series corresponds to the tract-based implementation. It performs very poorly: visibly worse than the community area-based implementation and worse than the current CPS plan, which is not graphed here. In the graph, it is particularly notable that the lines slope steeply downward

³¹Census tract numbers mostly align with community areas so we simply define a community area to be the 5th and 6th digit of the census tract ID.

from the beginning. This reflects that census tracts are small relative to the sizes of the schools. Even in large census tracts, we sometimes find that the single highest-scoring student has a score well below the current cutoffs for Payton and Northside.

The implementation using community areas, represented in the graph by the intermediate series marked in red, works much better. At current levels of minority representation, it is substantially more efficient than the current CPS plan at Payton (33% efficient vs. 24%) and also more efficient at Northside (25% vs. 20%). Hence, it represents another alternative that CPS could consider alongside the bonus schemes we have described above. To maintain current level of minority representation at Payton, one would first allocate 70 of the seats in open competition with the remainder to be allocated via the top X% rule.³²

One potentially troubling feature of the top X% style plans that is not apparent from our average-score based efficiency metric is that they sometimes let in a few students with very low scores. In the community-based implementation, this turns out to not be so severe: the lowest-scoring admitted students have scores of about 87. In the tract-based implementation, the lowest scoring admitted student at Northside has score of 57.7. Presumably, however, one could slightly modify a top X% rule to avoid this by only making those seats available to students above some pre-specified threshold. A switch to a top X% rule would also involve a reduction in low-income representation: the number of students qualifying for subsidized lunch is about about 2 percentage points lower under the top X% plan than under the current CPS policy.

6 Recap and Conclusions

This paper evaluates affirmative action plans as a means to achieve multiple goals: enabling schools to develop curricula tailored to students ability/preparation and letting students learn in a diverse environment. In our framework, the welfare consequences of any policy are a function of a two-dimensional sufficient statistic: the minority representation it achieves and the average scores of the students it assigns to the elite school. Motivated by this result, we propose that affirmative action polices can usefully be evaluated by their position relative to the Pareto frontier of efficient plans in this space.

We have discussed a number of different affirmative action plans in this paper. Table 6 collects

³²Note that we have assumed that open seats are filled first. This number would be different if seats reserved for community areas are filled first.

together some statistics on each of them.³³ Our first main empirical observation is that the CPS tier-based plan is highly inefficient: we estimate that it is just 24% efficient at Payton and 20% efficient at Northside. This means that the reduction in average composite entrance schools that the schools incurred in order to raise minority representation to its current level is four to five times as large as it would have needed to be if CPS were allowed to use race as a factor in admission. Another way to think of inefficiency is in terms of the level of minority representation that can be achieved. No CPS-like plan, not even allocating 100% of the seats by SES tier, would achieve 50% minority representation at Payton and Northside. A race-based plan could do this while simultaneously keeping average composite scores greater than 98.

Table 6 also reports within-school differences in majority and minority composite entrance scores under each plan and the composite score of the lowest-scoring admitted student. As we noted when we presented histograms of scores by race in Figure 8, another cost of the shift to a race-neutral plan is that racial gaps in average scores are about twice as large as they would be under the race-based plans that achieve similar minority representation. Almost all of the students with lower-tail scores are minority students.

Socioeconomic affirmative action has grander goals than just racial integration. Another of our main empirical observations, however, is that these goals are elusive in practice. In particular, we find that the CPS SES-based approach is actually less effective than a minority quota as a means to enroll low-income students. Our investigation of correlates of minority status suggest that this is an inherent problem. Students from low-SES neighborhoods who earn very high grades and standardized tests scores are unusual for their neighborhood and therefore may not be representative of their neighborhood. Minority status is also a weak predictor of low-income status once we know that a student is high achieving. But it works well enough so that we cannot regard SES-based plans using tract-level demographics as having an advantage over racial quotas on the dimension of addressing low-income representation.

We have examined several ways in which the CPS plan could be made more efficient as a means to increase minority representation. Eliminating the arbitrary discretization of tract SES into four tiers

³³Each of the alternate race-neutral policies presented below the current CPS Tier 70 plan in the table involves the choice of a weight or number of open seats that will affect the resulting level of minority representation. As we did in Table 5 we are reporting statistics on two slightly different implementations of each policy: in the top panel when presenting statistics on Payton we implement the policy with a weight chosen to exactly match the current minority share at Payton, and in the bottom panel we study versions that match the minority share at Northside. The measures of lunch efficiency correspond to the policies which match the minority shares.

Admissions Criterion	Average Composite Score	Majority-Minority Score Gap	Fraction Minority	Fraction Free/Red. Lunch	Lowest Scoring Admit	Efficiency Measure: Minority Lunch	
Walter Payton College Prep							
Pure score-based	99.1	0.376	0.214	0.150	97.1	-	-
Minority set-aside 20%	98.9	1.620	0.350	0.186	95.4	100%	7.3%
Minority set-aside 40%	98.2	2.936	0.500	0.268	92.9	100%	25.4%
CPS tiers 60%	98.3	2.626	0.341	0.236	90.2	22.9%	15.7%
CPS tiers 70%	98.0	3.174	0.368	0.245	89.0	24.2%	13.8%
Continuous CPS bonus	98.2	1.969	0.368	0.255	91.3	31.0%	20.9%
Re-weighted CPS bonus	98.3	2.061	0.368	0.236	90.2	33.9%	16.0%
Lasso bonus	98.4	2.351	0.368	0.241	89.7	37.7%	19.6%
Top-10% communities	98.3	2.919	0.368	0.218	87.4	33.6%	10.0%
Top-10% tracts	97.3	5.423	0.368	0.255	67.0	15.2%	10.2%
Northside College Prep							
Pure score-based	99.0	0.202	0.189	0.228	97.0	-	-
Minority set-aside 20%	98.8	1.728	0.344	0.266	95.3	100%	0.3%
Minority set-aside 40%	98.1	3.032	0.483	0.324	93.1	100%	10.4%
CPS tiers 60%	98.1	3.307	0.332	0.313	88.4	18.8%	7.1%
CPS tiers 70%	97.7	3.791	0.355	0.340	87.3	19.6%	9.6%
Continuous CPS bonus	97.7	2.242	0.355	0.382	89.9	19.0%	17.9%
Re-weighted CPS bonus	98.0	2.137	0.355	0.340	90.6	23.9%	11.8%
Lasso bonus	97.9	2.330	0.355	0.382	86.9	22.8%	21.4%
Top-10% communities	98.0	3.219	0.355	0.313	86.9	25.2%	6.8%
Top-10% tracts	97.1	5.542	0.355	0.340	57.7	12.8%	6.3%

Table 6: Comparison of admissions criteria

helps. Re-weighting the CPS variables provides an additional improvement. A Top 10% style plan can also provide an improvement if it is well designed – a plan based on too small units like census tracts (or elementary schools) would be much worse than the current plan and a neighborhood-based plan could also work worse than the one we have shown if some details of the implementation were done differently.

Another of our primary conclusions, however, is that it is hard to do much better in a place-based affirmative action plan. It should therefore be understood that eliminating race-consciousness in admissions comes at a cost. Our comparison of Payton and Northside suggests that this cost will in part depend on the school environment. In Chicago, it is easier to construct relatively efficient race-neutral plans at Payton than at Northside. This may be related to where both schools are situated relative to low-income minority neighborhoods.

We have not tried to address a number of other concerns that naturally arise when cities design affirmative action plans. Top 10% style plans ensure that places will go to students from across the city, which may be politically appealing. Plans also differ in their predictability. Minority shares are more predictable when the fraction of seats that will be reserved for underrepresented minorities is a design variable. Plans that, like the efficiency improvement we suggest, give bonus points to students from lower SES areas, will have more predictable majority-minority score gaps, but may result in minority representation that varies across schools or across years. Another obvious avenue for improvement would be to incorporate individual characteristics in addition to tract characteristics: one could consider whether a student receives free lunch, their parents' education levels, etc. While this is common in college admissions, in the public school context it may raise concerns about whether parents (especially of the disadvantaged) can be counted on to provide the desired information.

References

- ALON, S., AND O. MALAMUD (2014): “The Impact of Class-Based Affirmative Action on Admission and Matriculation Outcomes in Israel,” *Economics of Education Review*, 40, 123–139.
- AUSTEN-SMITH, D., AND R. FRYER (2005): “An Economic Analysis of ‘Acting White’,” *Quarterly Journal of Economics*, 120(2), 551–583.
- BRC (2011): “Blue Ribbon Commission Report on the 2011-2012 Magnet and Selective Enrollment Admissions Policy,” Chicago Public Schools, Available at: http://cpsmagnet.org/ourpages/auto/2011/8/22/47545626/BRC_Report%20Final.pdf and Last accessed: November 18, 2015.
- CASHIN, S. (2014): *Place, Not Race: A New Vision of Opportunity in America*. Beacon Press.
- CESTAU, D., D. EPPLE, AND H. SIEG (2015): “Admitted Students to Selective Education Programs: Merit, Profiling, and Affirmative Action,” *Journal of Political Economy*, forthcoming.
- CHAN, J., AND E. EYSTER (2003): “Does Banning Affirmative Action Lower College Student Quality?,” *American Economic Review*, 93(3), 858–872.
- CHUNG, K.-S. (2000): “Role Models and Arguments for Affirmative Action,” *American Economic Review*, 90(3), 640–648.
- CORCORAN, S., AND C. BAKER-SMITH (2014): “Pathways to an Elite Education: Application, Admission, and Matriculation to New York City’s Specialized High Schools,” NYU, Working Paper.
- DUBINS, L. E., AND D. A. FREEDMAN (1981): “Machiavelli and the Gale-Shapley algorithm,” *American Mathematical Monthly*, 88, 485–494.
- DUR, U., P. A. PATHAK, AND T. SÖNMEZ (2016): “Explicit vs. Statistical Preferential Treatment in Affirmative Action: Theory and Evidence from Chicago’s Exam Schools,” NBER Working Paper 22109.
- EPPLE, D., R. ROMANO, AND SIEG (2008): “Diversity and affirmative action in Higher Education,” *Journal of Public Economic Theory*, 10, 475–501.
- FRYER, R., G. LOURY, AND T. YURET (2008): “An economic analysis of color-blind affirmative action,” *Journal of Law, Economics, and Organization*, 24, 319–355.
- KAHLENBERG, R. (2003): *All Together Now: Creating Middle-Class Schools through Public School Choice*. Brookings Institution Press.
- (2008): “A Report to the Chicago Public Schools on Using Student Socioeconomic Status as a Factor in Student Assignment,” Available at: <http://s3.documentcloud.org/documents/524908/kahlenberg.pdf> and Last accessed: November 15, 2015.

- (2014): “Elite, Separate, and Unequal: New York City’s Top Public Schools Need Diversity,” NY Times, June 22.
- OCR (2011): “Guidance on the Voluntary Use of Race to Achieve Diversity and Avoid Racial Isolation in Elementary and Secondary Schools,” Memorandum of U.S. Department of Justice and U.S. Department of Education, Office of Civil Rights, Available at: <http://www2.ed.gov/about/offices/list/ocr/docs/guidance-ese-201111.pdf> and Last accessed: November 15, 2015.
- RAY, D., AND R. SETHI (2010): “A Remark on Color-Blind Affirmative Action,” *Journal of Public Economic Theory*, 12(3), 399–406.
- ROTH, A. E. (1982): “The Economics of Matching: Stability and Incentives,” *Mathematics of Operations Research*, 7, 617–628.
- STEELE, C. M., AND J. ARONSON (1995): “Stereotype Threat and the Intellectual Test Performance of African Americans,” *Journal of Personality and Social Psychology*, 69(5), 797–811.

Appendix

Proof of Proposition 2. Within each population using some cutoff is optimal for the same reason as before: holding fixed the curricula and the number of students from population j assigned to school 1, it is clearly optimal to send the highest θ students from population j to school 1. Assignments favoring population o ($\theta_o < \theta_u$) are obviously suboptimal: moving the cutoffs closer together improves diversity at both schools and reduces mismatch. A purely θ -based assignment also cannot be optimal: a slight increase in $\hat{\theta}_o$ together with a slight decrease in $\hat{\theta}_u$ that keeps enrollment at school 1 constant improves diversity and has no first-order effect on mismatch.³⁴ Hence, the optimum must have $\hat{\theta}_o > \hat{\theta}_u$.

The formula for the difference between the cutoffs comes directly from the first-order condition for the impact of slightly increasing $\hat{\theta}_o$ and slightly decreasing $\hat{\theta}_u$, keeping the the number of students assigned to school 1 fixed. The change in the mismatch costs for a type $\hat{\theta}_o$ student from being switched to school 1 is $-k(\hat{\theta}_o - c_1)^2 - (-k(\hat{\theta}_o - c_2)^2) = 2k\hat{\theta}_o(c_1 - c_2) + k(c_2^2 - c_1^2)$, and the change for a type θ_u student from the opposite switch is $2k\hat{\theta}_u(c_2 - c_1) + k(c_1^2 - c_2^2)$. Hence the total change in mismatch costs from changes that move an infinitesimal mass δ of students in each direction is $2k(\hat{\theta}_o - \hat{\theta}_u)(c_1 - c_2)\delta$. Such a shift also increases m_1 by δ/n_1 and reduces m_2 by δ/n_2 where n_1 and n_2 are the masses of students at schools 1 and 2. The optimum must have $m_1 < m < m_2$ so both changes hurt diversity.³⁵ The fraction n_1 of students at school 1 have a change of $-\delta\delta/n_1$, and the fraction n_2 students in school 2 have a change of $-\delta\delta/n_2$, so the total change in diversity benefits is $2d\delta$. Hence, the first order condition is $2k(\hat{\theta}_o - \hat{\theta}_u)(c_1 - c_2) + 2d = 0$, which gives the formula in the proposition. \square

Proof of Proposition 3. Suppose that system uses cutoffs $(\hat{\theta}_u, \hat{\theta}_o)$. If a small mass δ of students with types around (θ, z) are shifted from school 2 to school 1 the change in curriculum matching utility is to first order $-k\delta(((\theta - c_1)^2) - (\theta - c_2)^2) = -2k\delta(c_1 - c_2) \left(\frac{c_1+c_2}{2} - \theta\right)$. Again, approximating to first order, the fraction minority at school 1 is increased by $\delta(m(\theta, z) - m_1)/n_1$. The fraction minority at school 2 is increased by $\delta(m_2 - m(\theta, z))/n_2$. The net change in diversity-related utility is the sum of the effects on the n_1 students in school 1, the n_2 students in school 2, and the δ students who are shifted. The three terms add up to $2d\delta(m(\theta, z) - m)$. The first-order condition for the optimality of $\hat{\theta}_u$ is that the first-order approximations of the curriculum-matching and diversity-benefits must sum to zero at $\hat{\theta}_u$. This gives $d(m(\hat{\theta}_u, u) - m) = k(c_1 - c_2) \left(\frac{c_1+c_2}{2} - \hat{\theta}_u\right)$, which implies that $\hat{\theta}_u = \frac{c_1+c_2}{2} - \frac{d(m(\hat{\theta}_u, u) - m)}{k(c_1 - c_2)}$.

The first order condition for $\hat{\theta}_o$ is by the same calculation $\hat{\theta}_o = \frac{c_1+c_2}{2} - \frac{d(m(\hat{\theta}_o, o) - m)}{k(c_1 - c_2)}$. Subtracting these two expressions gives the formula for $\hat{\theta}_o - \hat{\theta}_u$ in the Proposition. \square

³⁴This argument does not rely on the assumption that the diversity benefit is constant. With equal cutoffs there is a discrete gap between the minority composition at school 1 and the optimum, so a similar result would obtain even if the diversity benefit were quadratic in $(m - m_i)$.

³⁵If $m_1 > m > m_2$ the change provides a diversity benefit while reducing mismatch costs so $\hat{\theta}_o$ and $\hat{\theta}_u$ could not have been optimal.

Proof of Proposition 4. The optimal curricula have $c_i = \mathbb{E}(\theta_\ell | s(\ell) = i)$. Hence,

$$W_e = \mathbb{E}(h(\theta)) - k(n_1 \text{Var}(\theta_\ell | s(\ell) = 1) + n_2 \text{Var}(\theta_\ell | s(\ell) = 2)).$$

The sum of the two variances expands as

$$\begin{aligned} & n_1 (\mathbb{E}(\theta_\ell^2 | s(\ell) = 1) - \mathbb{E}(\theta_\ell | s(\ell) = 1)^2) + n_2 (\mathbb{E}(\theta_\ell^2 | s(\ell) = 2) - \mathbb{E}(\theta_\ell | s(\ell) = 2)^2) \\ &= \mathbb{E}(\theta_\ell^2) - n_1 (\mathbb{E}(\theta_\ell | s(\ell) = 1)^2) - n_2 (\mathbb{E}(\theta_\ell | s(\ell) = 2)^2). \end{aligned}$$

Given that the weighted sum of the conditional means is $\mathbb{E}(\theta)$ and their difference is Δ their values are $\mathbb{E}(\theta_\ell | s(\ell) = 1) = \mathbb{E}(\theta) + n_2 \Delta$ and $\mathbb{E}(\theta_\ell | s(\ell) = 2) = \mathbb{E}(\theta) - n_1 \Delta$. Substituting these expressions into the above formula gives

$$\begin{aligned} W_e &= \mathbb{E}(h(\theta)) - k (\mathbb{E}(\theta^2) + n_1 \mathbb{E}(\theta)^2 + 2n_1 n_2 \mathbb{E}(\theta) \Delta + n_1 n_2^2 \Delta^2 + n_2 \mathbb{E}(\theta)^2 - 2n_1 n_2 \mathbb{E}(\theta) \Delta + n_1^2 n_2 \Delta^2) \\ &= \mathbb{E}(h(\theta)) - k \mathbb{E}(\theta^2) + k \mathbb{E}(\theta)^2 + k n_1 n_2 (n_1 + n_2) \Delta^2. \\ &= \mathbb{E}(h(\theta)) - k \text{Var}(\theta) + k n_1 n_2 \Delta^2. \end{aligned}$$

When $m_1 < m$ the fact that $n_1 m_1 + n_2 m_2 = m$ gives $m_2 > m$ and $m_2 - m = \frac{n_1}{n_2} (m - m_1)$. Hence,

$$\begin{aligned} W_d &= -n_1 d |m_1 - m| - n_2 d |m - m_2| = -n_1 d (m - m_1) - n_2 d \frac{n_1}{n_2} (m - m_1) \\ &= -2n_1 d (m - m_1). \end{aligned}$$

The formula for $W - W^s$ is follows immediately from subtracting the formulas for welfare under the two plans. \square

A Appendix

	Dep. Variable: Applicant is Free or Reduced Price Lunch					
	(1)	(2)	(3)	(4)	(5)	(6)
Median Income Percentile	-0.438 (0.007)	-0.696 (0.039)	-0.082 (0.073)	-0.010 (0.068)		-0.068 (0.067)
Adult Education Percentile	-0.501 (0.007)	-0.852 (0.036)	-0.700 (0.075)	-0.787 (0.067)		-0.651 (0.068)
Two Parent Percentile	-0.326 (0.007)	-0.461 (0.045)	-0.099 (0.078)	-0.067 (0.031)		0.081 (0.073)
Native English Percentile	-0.105 (0.007)	-0.668 (0.051)	-0.162 (0.074)	-0.077 (0.067)		-0.107 (0.066)
Home Owner Percentile	-0.307 (0.007)	-0.261 (0.040)	0.148 (0.057)			
ISAT Percentile	-0.334 (0.007)	-0.367 (0.042)	-0.071 (0.062)	-0.025 (0.059)		0.017 (0.058)
Sum of Tract Characteristics					-0.987 (0.054)	
Minority						0.197 (0.022)
Sample	Exam	Apps	Apps	Apps	Apps	Apps
	Apps	≥ 96	≥ 96	≥ 96	≥ 96	≥ 96
Observations	44,962	1,819	1,819	1,819	1,819	1,819
R-Squared			0.238	0.235	0.153	0.268
Misclassification			23.7%	21.9%	25.9%	20.8%

Table A1: Predictions of Free or Reduced Price Lunch Status

Variable	Weight
CPS adult educational attainment	-0.470
Per capita income	-0.339
Median value for owner occupied units	-0.131
Aggregation of adult education	-0.090
Fraction of 18–64 year olds speaking English “not well”	0.076
Fraction of foreign born from Europe	-0.030
Fraction of 25+ with at most 8th grade education	0.010
Fraction of children in households with income 1.5–1.99 poverty	0.009
Median age of males	-0.005

Table A2: Demographic variables selected by LASSO