The Inefficiency of Race-Neutral Affirmative Action: Evidence from Chicago's Exam Schools*

Glenn Ellison and Parag A. Pathak†

September 2015

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), which recently adopted a place-based affirmative action system in place of a race-based system for admission to their highly-competitive exam high schools. We develop a theoretical framework that conceptualizes the efficiency cost of race-neutral policies in terms of the extent to which they distort admissions decisions more than needed to achieve a given level of minority representation. Using data from CPS's exam schools, we find substantial inefficiency in CPS's place-based affirmative action system relative to the previous system based on racial quotas. We also compare the relative efficiency of several alternative race-neutral policies.

*We thank Katie Ellis, Susan Ryan and the staff at Chicago Public Schools for their expertise and help with the data. Vivek Bhattacharya and Alex Olssen provided excellent research assistance. Pathak acknowledges support of National Science Foundation grant SES-1056325.

†Department of Economics, Massachusetts Institute of Technology, Cambridge MA 02142 and NBER, e-mail: gellison@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 related efforts to boost minority participation grew increasingly widespread in public K-12 school systems and universities in the 1970s and 1980s.[1] In June 2007, the US Supreme Court struck down race-based admissions plans for public schools in Seattle and Jefferson County, Kentucky. The court ruled that the US Constitution did not permit those school districts to classify or assign students on the basis of race or national origin. Recent federal guidelines interpret the court as approving generalized race-based approaches such as the overall racial composition of neighborhoods, so long as they do not involve decision-making on the basis of any individual student's race.[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 is in the middle of an experiment with "generalized" race-conscious policies in admissions to one of their selective high schools, including the well-known Northside and Walter Payton high schools. From 1980-2007, a consent decree mandated race-based admissions for magnet and selective enrollment schools, with the goal of creating integrated schools (defined as 15-35% white and 65-85% minority). Admissions decisions are based on a composite score which combines middle school grades, standardized test scores, and an entrance exam. Through 2009 the admissions procedure was a racial-quota system mandated by a decades-old desegregation consent decree: 65% of the seats were reserved for minority students and 35% for nonminority students. In 2010 CPS replaced this system with a socioeconomic-based system. Each census tract in the city was assigned to one of four tiers using an index combining five socioeconomic

[1] California's Proposition 209 banned the use of racial preferences in state's public colleges and universities. Around the same time, the Texas Hopwood decision eliminated the use of racial quotas at public colleges and universities in Louisiana, Mississippi, and Texas. The Supreme Court allowed the use of race once again in 2003, but forbade quotas. Fisher vs. University of Texas provides a broad challenge to using race a factor in undergraduate admissions.

[2] The 2011 DOE 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).

[3] Kahlenberg (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.

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 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 socioeconomic variable (test scores in the local elementary school) was added to the index.

Our focus on Chicago is motivated by several reasons. Chicago's place-based affirmative action plan is held up as a national model for achieving rational and ethnic diversity in selective public schools (e.g., Kahlenberg (2014)). Moreover, Federal Department of Justice and Education Guidelines (OCR 2011) describe the following model alternative for achieving diversity at competitive schools, which mirrors the Chicago 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...*

Finally, Chicago employs a transparent and replicable affirmative action scheme, which can be analyzed quantitatively.

Arguments about the efficacy of race-neutral policies often focus on how minority representation changed when race-based policies were replaced. It is of course possible to eliminate racial gaps in admissions in a race-neutral manner by accepting students at random. However, elite schools do not lottery their admissions because it conflicts with the central reason for their existence: they develop the resource that advanced students represent by providing programs tailored to their knowledge and abilities. Therefore, the "effectiveness" of race-neutral policies inherently requires two components: measurement of the extent to which the policy provides desired diversity (or other) benefits; and measurement of the collateral damage that occurs when race-neutral restrictions make it impossible for schools to admit the students that they would most like to admit.

Our analysis begins in Section 2 with a theoretical model that formalizes the above reasoning and motivates a more specific methodology for assessing policies. In the model students are assumed to benefit both from a curriculum tailored to their knowledge/abilities and from learning within a diverse student body. The social optimal admissions policy will be race-conscious: a racial quota or

a bonus scheme giving some number of points to members of the underrepresented group will both work. If race cannot be considered, policies based on a proxy variable highly correlated with race can be nearly as efficient, but this is not always the case: we present an example in which policies based on a highly correlated proxy are completely ineffective. Accordingly, we argue that assessing the effectiveness of race-neutral alternatives to race-consious policies is inherently an empirical exercise. Our most important theoretical result is a proposition on sufficient statistics. We note that the social welfare produced by any policy, no matter how complex, is a function of two statistics on the allocation it produces: the average preparation level of students selected for the elite school; and minority representation. There is a Pareto frontier of efficient policies in this two-dimensional space. Race-neutral policies will lie inside the frontier. We define a welfare-motivated notion of relative efficiency reflecting policies' distance to the efficient frontier.

We provide more details on the CPS admissions process and our data in section 3. Our empirical analysis begins in section 4 with an examination of the Pareto frontiers which represent the fundamental constraints facing those designing policies for CPS. We focus on the two most selective selective-entry high schools in CPS: Walter Payton College Preparatory High School (Payton) and Northside College Preparatory High School (Northside). The schools are somewhat differently situated: Payton's central location is appealing to students living in a number of predominantly low-income African-American and Hispanic neighborhoods; whereas Northside is far from most predominantly African-American neighborhoods and would naturally draw more of its poor and minority students from surrounding middle-class neighborhoods. This provides us with an opportunity to explore the extent to which race-neutral policies work better in some situations than others, which may be useful for external validity. Our main goal in Section 4 is simply to document the shape of the Pareto frontier. 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. (The composite score combines middle school grades, standardized test scores, and a special entrance exam.) And then we present graphs illustrating the decline in average composite scores that must be incurred if one wants to increase minority populations to various levels. The motivation for the CPS policy is not just racial integration: both the former and current policies are presumably also intended to increase representation of low-income and otherwise disadvantaged students. Accordingly, we also examine 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. Again, policies along the frontier can be implemented in multiple race-conscious ways:

one can use quotas; or one can define an adjusted composite score by adding a particular number of bonus points to the composite index of minority and/or free lunch students and then select the students with the highest adjusted composite scores.

Section 5 examines Chicago's socioeconomic-based affirmative action plan and variants that reserve more or fewer slots for students from low-SES areas. When we examine minority representation our primary focus is on the relative efficiency of the CPS plan as a race-neutral method for increasing minority representation. Here, we note that 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. And there are also limits to what can be accomplished with these policies even if one is not concerned with efficiency: minority representation would decrease even if one were to allocate 100% of seats to the tier-specific quotas. A second finding is that the CPS policy is more effective in some situations than others: it is more efficient at Payton than at Northside. SES-based affirmative action plans may provide integration on other dimensions that could be regarded as offsetting their decreased efficiency as a tool for racial integration: they are designed to increase the number of students from low-income and otherwise disadvantaged backgrounds. Our examination of this effect in Chicago yields a third striking result: the CPS plan is only slightly more efficient than a racial quota as a means for increasing the number of low-income students at Walter Payton; and it may even be decreasing the number of low-income students at Northside.

Section 6 provides some additional detail on how the switch from race-based to race-neutral affirmative action affects the class composition at Payton and Northside. One observation is that the CPS shift to a race-neutral policy increases within-school heterogeneity. The majority of students have 99th percentile scores under either plan, but the lower tail of students who do has become much larger and now extends much lower. A second observation concerns the within-school racial gap in scores. In principal 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 the average scores of whites and Asians. But a second effect working in the opposite direction is that some of the highest scoring African American and Hispanic students (including some low-income students qualifying for free lunch) fail to gain admission which also brings down the minority average. In practice we find that the second effect is stronger than the first and the SES-based plans lead to a larger within-school racial gap in composite scores.

Section 7 asks whether the inefficiency of the socioeconomic policy is due to a suboptimal implementation of the general idea or whether it is inherent to any race-neutral policy. The CPS index of socioeconomic status 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 may not have this property for two obvious reasons: in a multiple regression one will not want to put equal positive weight on all variables that are individually correlated with minority status; and variables that are correlated with minority status in the full sample will not necessarily be correlated with it once one conditions on the (in some cases extremely unusual) event of a student having a 98th percentile score. With this motivation we consider alternative race-neutral policies that differ from the CPS policy in three ways: we weight the CPS tract-level variables and let them enter the index nonlinearly; we allow the policy to consider an student's free-lunch status in addition to tract-level variables; and we implement affirmative action via a continuous bonus scheme instead of using discrete. Our most important finding for Chicago is that socioeconomic affirmative action can be made more efficient than the current CPS policy. A second finding of perhaps broader importance is that the efficiency that can be obtained even from a sophisticated policy depends on the school environment: in the case of Payton our improved race-neutral policy eliminates about half of the inefficiency; but at Northside even complex policies remain highly inefficient as a method for increasing racial diversity.

Section 8 examines an alternate race-neutral approach – policies along the lines of Texas's top 10% rule. 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. We find that policies of this type are much less efficient than the CPS policy. Intuitively, such policies are highly inefficient when used for programs serving extreme high achieving students because we are dealing with outliers and there will be many cases in which unit A has two candidates who are much stronger than the top candidate from unit B, even when units A and B serve students of similar socioeconomic status.

2 A Simple Model of Elite Schools and Affirmative Action

In this section we discuss a simple model of affirmative action at exam schools. Our model assumes that exam schools benefit students by providing a curriculum tailored to their ability and preparation. We then add the assumption that diversity is valuable and establish the optimality of an affirmative

action plan. The majority of the section discusses second-best affirmation plans when the school system cannot explicitly advantage members of some group and must instead target an imperfectly-correlated proxy variable. One observation is that race-neutral alternatives sometimes work very badly. The model also motivates a method for quantifying efficiency losses from race-neutral affirmative action plans.

2.1 Elite schools

Our first building block is the simple model of Ellison and Pathak (2014) in which it is desirable to have elite schools because students benefit when the curriculum is tailored to their ability/preparation.

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 students' ability, preparation, etc.) The school system operates two schools with curricula $c_1, c_2 \in \mathbb{R}$. Suppose that the expected outcome for a student of type θ if assigned to school i depends on the student's type and on the match between the students' type and the curriculum:

$$V_i(\theta) = h(\theta) - k(\theta - c_i)^2.$$

Note that curricula are described by the type of student they serve best and k indexes the importance of student-curriculum matching. Suppose that the school system's goal is to maximize the sum of the outcomes of its students.

This model provides a simple rationale for forming an elite school: 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.[4]

Proposition 1 *In the above model the optimal school system design involves choosing a cutoff $\hat{\theta}$ and sending students with $\theta > \hat{\theta}$ to school 1 and students with $\theta < \hat{\theta}$ to school 2. The curricula are $c_1 = E(\theta|\theta > \hat{\theta})$ and $c_2 = E(\theta|\theta < \hat{\theta})$. Further, the optimal $\hat{\theta}$ is the solution to*

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

The argument is straightforward. Given any assignment scheme the curricula c_1 and c_2 would be

[4]It's worth noting that this rationale for elite schools contrasts with those motivated by an increasing-differences assumption implying assortative allocations are efficient as in, for example, Epple and Romano (1998)

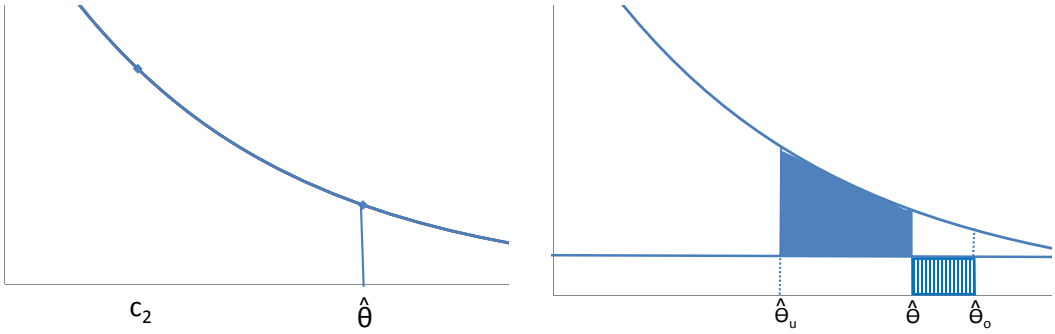


Figure 1: Optimal admissions policies with and without a diversity benefit

the mean θ's of the students assigned to each school. Given any curricula with $c_1 < c_2$ we must assign each student optimally, which implies that each student with $\theta > \frac{c_1+c_2}{2}$ is assigned to school 1, i.e. that the allocation is as described in the proposition.

The left 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 should be thought of as continuing substantially further to the right.[5] 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 within each school. The value of c_2 is shown. (The value of c_2 is not shown because it is to the right of the the area we have graphed.) The cutoff $\hat{\theta}$ is equidistant between c_2 and c_1.

2.2 Optimal affirmative action with diversity benefits

We now add a diversity benefit and show that the optimal assignment involves affirmative action.

Suppose that the school system serves two populations indexed by $j = o, u$. The labels reflect that population o will be overrepresented and population u 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. We suppose that due to some source of disadvantage the type distribution is higher in population o: $1 - F_o(\theta) > 1 - F_u(\theta)$ for all θ. We add a diversity benefit to the model by assuming that outcomes are highest when each school is representative of the population:

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

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

where m_i is the fraction of students in school i from population u. Here d indexes the importance of diversity. We assume here that diversity is valued because it improves all students educational outcomes, e.g. by exposing them to differnet viewpoints. Similar results could be derived in models motivated by other assumptions, e.g. if younger generations benefit from minority students' serving as role models. The model provides a rationale for affirmative action.

Proposition 2 *In the above model the optimal design sends population j students to school 1 if and only if $\theta > \hat{\theta}_j$. The optimal cutoffs $\hat{\theta}_o$, and $\hat{\theta}_u$ satisfy $\hat{\theta}_o > \hat{\theta}_u$ and*

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

Students are perfectly sorted on θ within each population, 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) and smaller when educational outcomes are more sensitive to student-curriculum mistmatch (k is larger) and when the curricula are farther apart ($c_1 - c_2$ larger).

Such affirmative action could be implemented equally well in either of two ways. The school system could use quotas and reserve $m(1-F_u(\hat{\theta}_u))$ seats at school 1 for members of population u. Or it could be implemented under a race-as-a-factor scheme in which students from population 2 are admitted over students of population 1 if their scores are close enough, where "close enough" means within $\hat{\theta}_o - \hat{\theta}_u$.

The argument for this proposition is straightforward. We first argue by contradiction that cutoffs must satisfy $\hat{\theta}_u < \hat{\theta} < \hat{\theta}_o$ and the fractions of underrepresented students must satisfy $m_1 < m < m_2$. The formula for the difference in cutoffs comes directly from the first-order conditions. The details are in the appendix.

Figure 1 provides an illustration. We have divided the student population from the left panel into two subpopulations: the lower part reflects the distribution of types within population o and the upper portion reflects the distribution within population u. If the same cutoff $\hat{\theta}$ is applied to all students then population u will be substantially underrepresented within 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 will always improve

welfare. The right panel of the figure shows the optimal affirmative action plan.[6] Here, the shaded areas show the larger sets of added and displaced students. At the cutoffs $\hat{\theta}_u$ and $\hat{\theta}_o$ it is no longer welfare-improving to add more students from population u (or remove students from population o) from the elite school. There would still be an improvement in diversity. But the mismatch costs would exactly offset the benefit.

2.3 Affirmative action with imperfect proxies

We now discuss second-best affirmative action policies when a school system cannot explicitly consider group membership but can use a correlated proxy.

Consider first a two-population model as above in which a school system cares about the fraction m_i of students from population u at each school. Suppose that each student has an observable characteristic $z \in \{o, u\}$ that is equal to the population j that the stduent belongs to with probability $p > \frac{1}{2}$ (and that this is independent of θ). Suppose that the school system cannot consider j in admissions, but can consider the proxy z.

The left panel of Figure 2 provides a simple illustration of the optimal policy in an example with $p = \frac{3}{4}$. The dashed lines show how the students within each population j are divided on z. In the lower rectangle of students from population o three fourths of the students are in the lower $z = o$ bar. And in the upper set 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 we are adding 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 we are displacing by raising $\hat{\theta}_o$ (shaded by vertical bars) also include students from both populations.

The left panel of Figure 2 illustrates what optimal cutoffs might look like in a model of this type. They are closer together than were the cutoffs with unrestricted affirmative action because only a fraction of the students who are added to school 1 by an incremental reduction in $\hat{\theta}_u$ are from population u. Hence, the per-added-student diversity benefit is smaller. The per-student mismatch costs are the same, so we optimally lessen the degree of affirmative action. Note that this means that minority representation 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 are from the underrepresented population; and (2) because affirmative action is a less effective tool, the school

[6] Again, the figure is constructed by solving a numerical example. We again take $k = 1$, $d = 1$, and θ exponentially distributed and assume that the mass of overrepresented students of type θ is $\min(0.09, e^{-\theta})$.

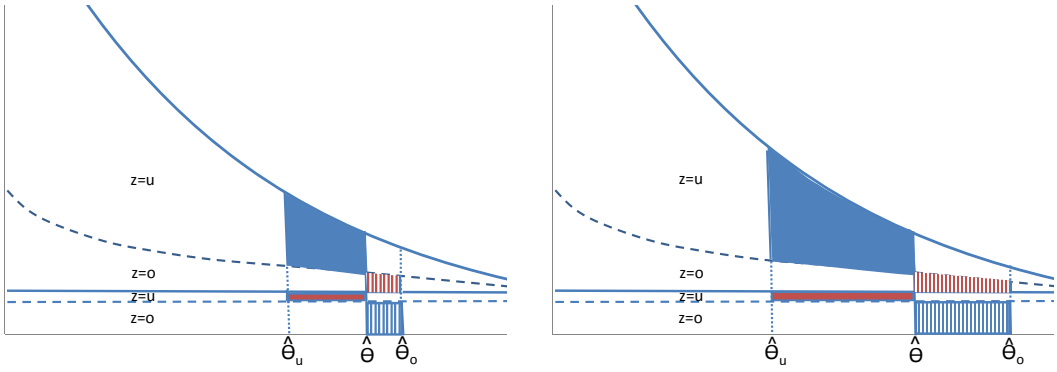


Figure 2: Affirmative action with an imperfect proxy

system uses it less aggressively.

The result that a school system will respond to the limited effectiveness of a proxy-based affirmative action by placing less weight on the proxy depends on the assumed form of diversity benefit. In practice, school systems may feel that the marginal benefits of diversity are greater when they are farther from the ideal point. In this case, the school system may instead institute policies like those illustrated in the right panel of Figure 2. In this example, the cutoffs $\hat{\theta}_o$ and $\hat{\theta}_u$ are farther apart than they were under unrestricted affirmative action. Minority representation does not end up higher than under unrestricted affirmative action – in the graph we have displaced the high-scoring minority students in the region with red bars and also failed to add the slightly below $\hat{\theta}$ students in the white region between the solid red and blue regions. The distortions in match-quality are larger. And there is more heterogeneity in types wtihin the elite school. The red regions in the figure highlight the inefficiency of the policy: some students from population u with very high θ are denied admission, while some students from population o with low θ are admitted.

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

One seemingly innocuous assumption in the above model may be both unrealistic and more important than it appears. We assumed that the distribution of the proxy was independent of θ within each population. This would not be expected to hold in practice if the proxy itself has a causal impact on student preparation. We show now that when this is the case it is possible that proxy-based affirmative acton can be completely ineffective as a means for increasing minority representation.

To make this point we present a simple example in which θ, best thought of here as a measure of

preparation reflected in a test score, is lower in the underrepresented population for two reasons: it is income-related 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 any variable that can be used in admissions. Specifically, assume that student ℓ's type is

$$\theta_\ell = Income_\ell + 0.1 \ Over_\ell,$$

where $Income_\ell$ is student ℓ's income and $Over_\ell$ is a dummy variable for being a member of the overrepresented population. Suppose also that the income distribution differs across populations: income 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 the 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 will go to members of the underrepresented population. The admissions cutoff will be $\hat{\theta} = 0.9$. 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.[7]

Suppose that the school system knows each student's income and can use this variable in its admissions process. Low income is correlated with underrepresentation and has a causal effect on the test score θ so it seems natural to use income as a proxy variable. Suppose that the school system does this by dividing students into four equal-sized tiers on the basis of income and then selecting the top 20% of students from each income tier.

Proposition 3 *In example described above, no students from the underrepresented group gain admission to the elite school under the tier-based affirmative action policy.*

The algebra for the result is straightforward. The bottom tier consists of 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 $\hat{\theta}_4 = 0.3$. All students above

[7] If 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 also gets 37.5% of the seats if the income differences are present but the group-related difference is not.

this cutoff are from the overrepresented group. Calculations for the other tiers are similar. The cutoffs are $\hat{\theta}_3 = 0.55$, $\hat{\theta}_2 = 0.80$, and $\hat{\theta}_1 = 1$. Overrepresented students get all of the seats in the elite school.

While this may seem paradoxical at first it is intuitive. As Chan and Eyster (2003) noted one method for increasing representation of an underrepresented group is to add noise to the decision process. In the completely-noisy limit no group will be underrepresented. In this example, we can think of income differences as a source of noise that lets some students from the underrepresented group overcome their other disadvantage. When we switch to an income-based affirmative action procedure we are removing one source of disadvantage that disproportionately affects the underrepresented group. But we are also removing the source of noise. The example is constructed so that the latter effect is more important than the former.

A number of other income-based affirmative action policies would work equally badly in this enviroment. 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 get any seats under a policy which selected students in order of adjusted scored after defining the adjusted score to be $\theta_k - \alpha\, Income_k$ score for any $\alpha > \frac{2}{3}$. An unrealistic feature of this example is that there is no heterogeneity in θ within each population other than the income-related heterogeneity. Adding some such heterogeneity would make affirmative action work better and eliminate the starkest result. But the noise-reduction effect would still be present and tier-based plans would be counterproductive as a method to address underrepresentation in a range of specifications.

In the model of the previous subsection proxy-based affirmative action can approximate unrestricted affirmative action if the proxy is sufficiently correlated with group membership. But its efficiency depends on the correlation. The example here points out that a correlated proxy can be of no help if it is separately related to admissions score. This range of possibilities motivates proxy-based affirmative action as an empirical topic.

2.5 Welfare, Pareto frontiers, and measuring inefficiency

In this section we discuss the welfare consequences of affirmative action policies in the context of our basic model and present some results on the tradeoff between curriculum-matching and diversity and on welfare losses from proxy-based affirmative action.

Consider again our basic model with two groups and outcomes $V_i(\theta) = h(\theta) - k(\theta - c_i)^2 - d|m_i - m|$. Throughout this section we 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: 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 we choose an optimal plan by maximizing over m_1. Let m_1^s be the population u representation at school 1 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 to this interval for the remainder in the subsequent discussion.

A nice feature of this two-step approach is that the first step is very simple. The optimal plan conditional on m_1 is simply to send the $n_1 m_1$ underrepresented students with the highest θ's to school 1 along with the $n_1 - n_1 m_1$ top students from the overrepresented population. The cutoff scores that do this are $\hat{\theta}_o^{FB}(m_1) = F_o^{-1}\left(1 - \frac{n_1 - n_1 m_1}{1-m}\right)$ and $\hat{\theta}_u^{FB}(m_1) = F_u^{-1}\left(1 - \frac{n_1 m_1}{m}\right)$. The optimal curricula are

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

where we have written E_o an E_u for expectations of θ's drawn from F_o and F_u, respectively.

In the discussion below it is useful to have a notation for the difference between the mean types of the students at the two schools. Using ℓ as an index for the students and writing $s(\ell)$ for any function giving the school to which each student is assigned, this is formally defined by

$$\Delta(s(\ell)) \equiv E(\theta_\ell | s(\ell) = 1) - E(\theta_\ell | s(\ell) = 2)$$

We often omit the argument of Δ when the policy under discussion is clear. We 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 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

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

The proposition below gives simple formulas for these that apply whenever curricula are optimal given the student assignment. Its 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,*

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

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

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

$$W - W^{sb} = 2n_1d(m_1 - m_1^{sb}) - kn_1n_2(\Delta^{sb\,2} - \Delta^2).$$

An important observation about the formula for W_e is that it is a monotone increasing function of Δ. This is intuitive: the benefits from dividing the sytem into two schools comes from differentiating instruction. When this can be done to a greater extent students are better off. The expression for W_d reflects the linear diversity benefit. The formula comparing welfare under an affirmative action plan to welfare under a purely score based plan brings out the basic tradeoff. Affirmative action provides diversity benefits which 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. 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 depend on the degree to which curriculum differentiation Δ is reduced.

In the unrestricted affirmative action model we can think of school system as as choosing among a set of points in (W_d, W_e) space. Each level m_1 for a minority quota gives a point on a downward sloping frontier in this space. The relative importance of curriculum matching k and diversity d will determine which point on the frontier a school system chooses. We saw above 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, $E(\theta_\ell|s(\ell) = 1)$. Accordingly, one can equivalently represent affirmative action plans as points in a two-dimensional set with more concrete 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 $(m_1, E(\theta_\ell|s(\ell)=1))$ 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 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. Although they are simple, we think that these observations are economically important. The zero-slope result brings out 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 brings out 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 play, regardless of how it is constructed, will produce an outcome corresponding to some point in $(m_1, E(\theta_\ell|s(\ell)=1))$ below the Pareto frontier. The distance between the plan and the frontier can be thought of as a quantitative reflection of the efficiency losses incurred by adopting the proxy-based plan istead of a race-based plan. The corollary below provides a formal justification for this thought: it notes that the efficiency loss relative to the 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 assigned to school 1 under $s(\ell)$ are from population u with $m_1 < m$. Let $\Delta = E(\theta_\ell|s(\ell)=1) - E(\theta_\ell|s(\ell)=2)$. Then, the welfare loss from having adopted 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*

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

In our empirical work we will usually present a measure of "relative efficiency" when discussing plans. Consider any plan $s(\ell)$ under which a fraction m_1 of the students at school 1 are from population u and with the difference in the average θ across the two schools is $\Delta(s(\ell))$. Given the way that weighted averages work, 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:

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

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 being from population u. Relative efficiency will be lower than 100% but still positive under any plan other than the optimal race-conscious plan. Note that the corollary above suggests that welfare losses are quadratic in the reduction in Δ which so that the 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 80% efficient plan to a 79% efficient plan is greater than the loss from moving from a 90% efficient plan to an 89% efficient one. But we felt that the scaling of the relative efficiency variable was easier to interpret than a quadratic scale and the corollary does imply that welfare in our model will be monotonically related to efficiency as we have defined it.

3 Affirmative Action at Chicago's Exam Schools

Data

TBD

4 The Pareto Frontier: What is Feasible in Chicago?

In this section we present empirical estimates showing what is feasible at Payton and Northside if there are no restrictions on using race and/or free-lunch status in the admissions process. We present graphs of Pareto frontiers illustrating the extent to which average tests scores must be reduced in order to boost minority/free-lunch representation to various levels and also discuss within-school inequality.

4.1 The tradeoff between composite scores and minority representation

We first examine the Pareto frontier in the dimension of average composite scores vs. minority representation. The left panel of Figure 3 graphs the highest average composite score that can be obtained within each school at each level of minority representation. The minority percentage is along the x-axis and the average composite score is along the y-axis. The red squares are for Payton and the blue triangles are for Northside.

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 98.9 at Payton and 99.1 at Northside.[8] About 20% of the students at each school would be underrepresented minorities. The flatness of the graphs just to the right of these points reflects 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 Peyton decline to 98.6 at 40% minority, 98.2 at 50%, 97.8 at 60%, and 97.1 at 70%. The Northside frontier has a virtually identical shape although average scores will be slightly lower at any given level of minority representation.

4.2 The tradeoff between composite scores and low-income representation

We now turn to tradeoffs between average composite scores and the representation of low-income students, which we examine by looking at the number of students qualifying for free or reduced-price lunch. The right panel of Figure 3 presents a graphs of the maximum average composite score that can be maintained if one wants to increase the number of low-income students to various levels. Again, the score-maximizing outcome can be implemented either by setting a free/reduced lunch quota or by adding an appropriate number of bonus points to the composite scores of students on free/reduced lunch. The shapes of the frontiers are generally similar to those examining the tradeoffs inherent in inreasing minority representation. Low-income representation can be increase to about 10% above its

[8] To provide some scale, a student would receive a composite score of 98.9 if they achieved a score deemed to be at the national 98th percentile on the admissions test, had 97th percentile scores in both English and Math on the ISAT, and a perfect middle school GPA.

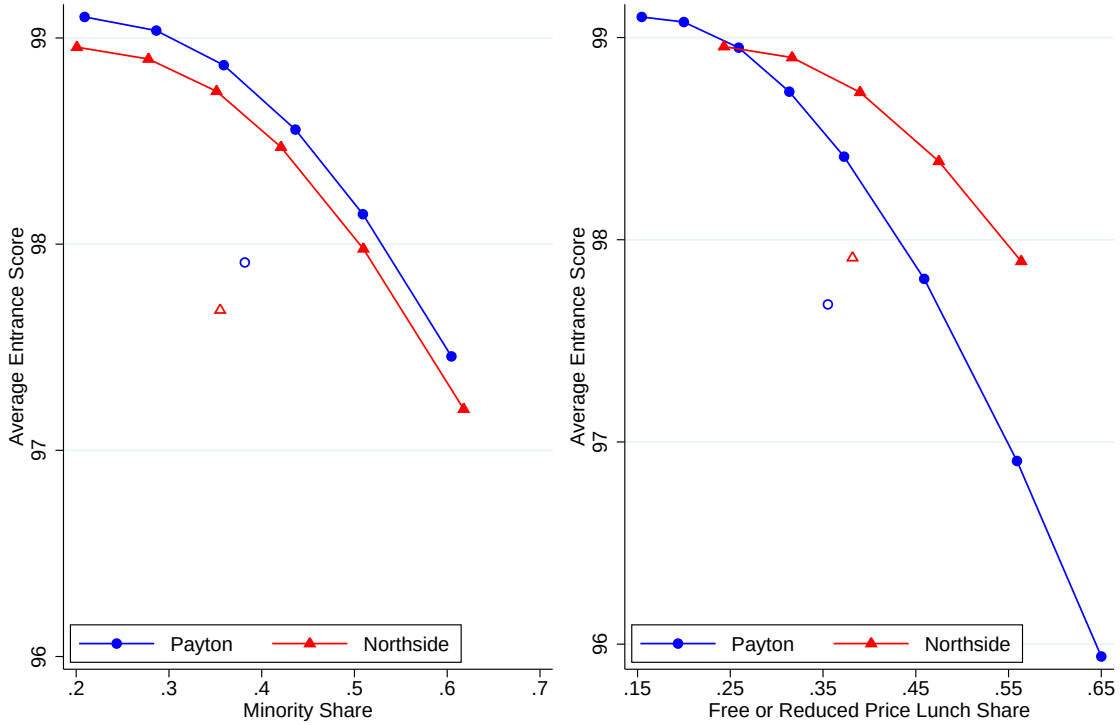


Figure 3: Composite scores vs. minority and free-lunch representation under unrestricted affirmative action

level under a purely score-based policy with only a slight decrease in average composite scores. But the tradeoffs becomes more pronounced at higher levels of low-income representation.

One clear difference between the minority and low-income graphs is the differences across schools. In the minority case the two curves were very similar. Here, they are not. The leftmost points on the low-income graphs record the outcomes under a purely score-based admissions policy. They show that Northside has many more students qualifying for free lunch under such a policy (24% vs. 14%). And 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 in order to increase free-lunch representation by any given amount is smaller.

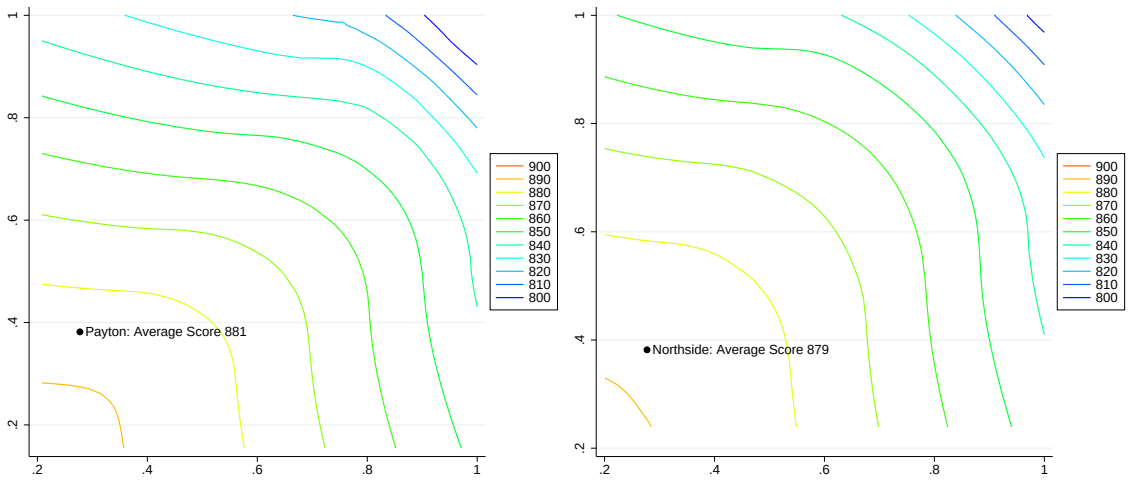


Figure 4: Levels of minority and low-income representation that can be simultaneously achieved at various score levels

4.3 The tradeoff between low-income and minority representation wtih scores held fixed

Schools systems may want to maintain both minority and low-income representation at elite schools. In this section we present some evidence 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 what we do here is to present graphs that have a few slices from this surface. Specifically, we suppose that each school wishes to maintain an average composite score of at least 99 or 98 or 97, 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 4 presents the graph for Payton. The star reflects the class composition under a purely score-based admissions policy: 20.2% of the students are minorities and 13.3% are eligible for free or reduced lunch. The two curves then describe the highest levels of minority and low-income representation that can be simultaneously maintained if one wants to keep the average

composite score at 98 or 97.[9] The primary observation we would like to make about these curves that policies targeting race will increase low-income representation and vice-versa, but that the effects are far less than one-for-one. For example, the rightmost points of these curves we can see that pure racial preferences policies will have the side effect of increasing low-income representation, but that the increase is much less than one-for-one. For example, the rightmost points on these curves correspond to pure racial preference policies. Comparing the rightmost point on the outer curve to the star indicates that a racial quota policy that reduces average composite scores from 98.9 to 97.0 will increase the percentage of minority students increases from 20% all the way up to 72%. The percentage of low-income students be increased as a side effect, but the increase is much more modest and the proportion of low-income students would still just be 32%. The leftmost points on the curves show that a pure low-income preference will similarly have the side-effect of increasing minority enrollment. But again, the effect is less than one-for-one: a low income preference that reduces the average composite score to 97.0 will increase low-income representation to 58%, but the minority representation to just 43%.

A second relevant observation about the curve is that they are quite curved. The fact that the curves are nearly vertical near the right endpoint indicates that a modification of a pure racial-preference that bases portion of the preference on income rather than race can substantially increase low-income representation while only slightly reducing minority representation. (For example, if we wish to maintain average composite scores at 97 we can do this and increase low-income representation from 32% to 40% while only reducing minority representation from 72% to XXX.) And 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 only a small effect on low-income representation (and while keeping average composite scores from dropping at all). Accordingly, policy makers who valued 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 4 presents a comparable graph for Northside. Again the star indicates the class composition under a purely score-based admissions policy: 21.0% if the students are minorities and 25.4% are on free or reduced lunch. There three curves show low income-minority compositions that are feasible if one wants to keep the average composite score at 99, 98, or 97. The distance of the innermost curve from the star is another illustration of how much scope there is for altering class

[9] Recall that the average composite score is 98.9 under a purely score-based policy.

composition while reducing average composite scores just from 99.1 to 99.0. The fact that the entire inner curve is above and to the right of the star reflects that once again racial preferences will have the side effect of raising low-income representation and vice versa. And once again the curves are quite flat near the left endpoints and quite steep near the right endpoints which suggests that policy-makers might prefer an interior point.

Implementing a policy that achieves an interior point requires making explicit use of both a free-lunch eligibility and minority status. Assuming this is possible, policies can be implemented using simple bonus schemes. Students have x of bonus points added to their score if they qualify for free/reduced lunch, y of bonus points added if they are a minority, and assignments are then made simply by accepting students in the order of their bonus-adusted scores.[10] 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 (highest minority representation) point on the average-score-98 curve is obtained by giving XXX points for minority status and no points for being low-income whereas the leftmost (highest low-income representation) point on the same curve is obtained by giving YYY points for low-income status and no points for being a minority.

5 How Efficient Are Chicago's Policies?

As noted earlier CPS policy may be most sophistcated existing race-neutral policy that a school system has adopted in an attempt to achieve goals that were formerly achieved via explicit race-based policies. Accordingly, it is of broader interest to explore whether the CPS policy is an effective substitute for a race-based policy. In this section we provide evidence on one conceptualization of "effective substitute" in examining how the admitted classes at Payton and Northside compare with the classes that could have been admitted if there were no restrictions on the form of affirmative action. Obviously, any restriction on affirmative action will require that CPS admit fewer minority students and/or reduce the average composite scores of the admitted students by some degree. The interesting unknown is what are the magnitudes of these impacts. We also discuss effects on low-income representation and within-school achievement gaps.

[10] Low-income minority students would receive both bonuses.

5.1 The policy as a tool for racial integration

We begin by examining the effectiveness of CPS's policy as a means of increasing minority representation at Payton and Northside. The left panel of Figure 5 provides data on what can be acheived at Payton. The red squares are same Pareto frontier we saw in Figure 3: they illustrate the inherent tradeoff the school faces when trying to increase minority representation and maintain high aveage composite scores. The green triangles graph the constrained Pareto frontier of average composite score-percent minority combinations that are possible when one is restricted to policies that involve reserving some fraction of seats for students in each of the CPS SES-tiers.

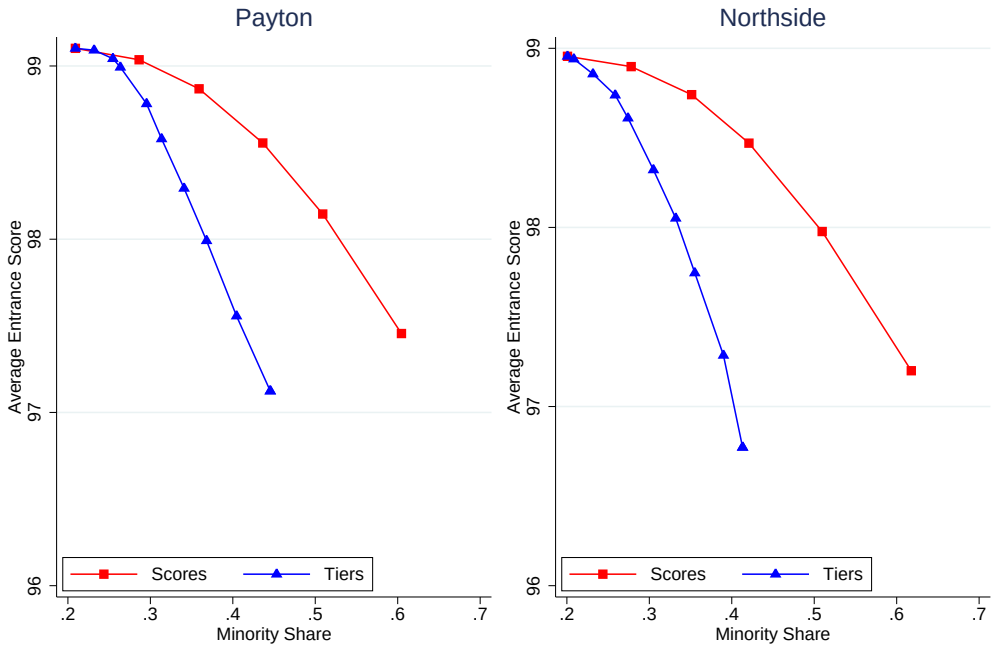


Figure 5: Feasible composite scores at various levels of minority representation: the effect of restricting affirmative action to the CPS race-neutral policy

The upper left points on the two curves are identical: reserving zero seats for each SES tier is equivalent to reserving zero seats for minorities. At first the CPS policy is nearly efficient: if one wants to increase minority representation from 20% to 25% the CPS policy can do this while maintaining just about the same average composite score as would be possible given any race-based policy. But the CPS policy becomes much less efficient if one wants to further increase minority

representation. With unrestricted affirmative action, increasing minority representation from 20% to 40% only requires reducing average composite scores from 98.9 to 98.6. Using the CPS policy it would require reducing average composite scores to 98.0. Another way to think of the difference is in terms of the difference in minority representation that can be achieved if any given reduction in average composite scores is considered acceptable. At the cost of reducing average composite scores to 98.0 an unrestricted policy could increase minority representation to 55% instead of 40%.

In practice a school system must decide which point to choose on its feasible frontier. The larger outlined square shows the level of underrepresented minority representation that CPS chose at Payton in the final year of the race-based policy. The two larger triangles show the two choices that were made under the new policy. Initially CPS reserved 60% of seats to be allocated by SES tier. After noting that minority representation had declined it increased this to 70%. At this point underrepresented minorities are 40.8% of the Payton class.

As we noted in the theory section, one way to quantify the efficiency of an affirmative action policy is in terms of what fraction of the reduction in composite scores under the policy was necessary to achieve a given goal. The remainder of the score decline can be seen as the implicit cost of restricting the set of affirmative action instruments. If one regards the goal of the CPS policy as being simply to increase minority representation, then the efficiency calculation just needs to identify the highest average score at which a minority representation of 40.8% would have been achievable via a race-conscious minority quota policy. Looking straight up from the 40.8% minority point to the efficient frontier we find that a minority quota could have achieved this level of representation with an average composite score of 98.6. Hence, the CPS policy can be regarded as only 26% efficient as a tool for increasing minority representation: almost three-fourths of the reduction in composite scores, $\frac{98.6-97.8}{98.9-97.8}$, could have been avoided by using a race-conscious policy.

The positions of the two large triangles (the 60% reservation vs. the 70% reservation) is also a nice illustration of the observation that the efficiency consequences of restricting affirmative action tends to be more severe when larger changes in school composition are desired. CPS's decision to shift an extra 10% of seats to the tier-based allocation caused a substantial 0.3 point drop in average composite scores, but did not increase minority representation very much (from 39.3% to 40.8%). Under an unrestricted affirmative action policy the decrease in average scores necessary to increasing minority representation from 39.3% to 40.8% is just 0.05 points. Hence, in an incremental sense the expansion of the policy can be thought of as only about 16% efficient.

The right panel of Figure 5 presents a similar graph for Northside. The shape of the Pareto frontier

is virtually identical. The shape of the frontier describing what happens as we increase the number of seats reserved for each CPS tier is also generlly similar, but here there are a couple of visible differences reflecting ways in which the CPS plan is less effective. Near the upper left endpoint of the graph we see that the initial flat segment is almost nonexistent. Increasing minority representation from 25% to 35% requires a drop in the average composite score that is more than twice as large as was needed at Payton (1.3 points vs. 0.5 points). And the Northside graph becomes more convex at high levels of minority representation reflecting that it becomes increasingly difficult to add minority students. No policy based on the CPS tiers, not even increasing the SES Tier set aside to 100%, will get minority representation as high as 43%. CPS has always chosen to use the identical tier-reservation policy (first 60% then 70%) at each school. As a result, shift to the race-neutral policy has resulted in somewhat lower test scores than under the racial quota regime and a substantial reduction in minority representation. Again, the 60% and 70% points are marked by large triangles. The 70% policy results in a class that is 33.6% underrepresented minorities with an average composite score of 97.7.

A minority representation of 33.6% could have been achieved while maintaining an average composite score of 98.9 if CPS had used a minority quota. Hence, the CPS policy is only 13% efficient as a method for increasing minority representation. Intuitively, CPS policy does worse at Northside than at Payton because many of the strong minority applicants to Northside live in the middle class neighborhoods that surround the school. A system which 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.

5.2 The policy as a tool for admitting low income students

The CPS affirmative action policy can be motivated as a tool to achieve socioeconomic as well as racial/ethnic integration. In this section we examine a related question: to what extent is the policy increasing the number of low-income students attending Payton and Northside. A striking result is that the CPS SES-focused plan is not actually an improvement on the old racial quota system as a tool for increasing the number of low-income students.

The left panel of Figure 6 contains three graphs related to potential class compositions at Payton. First, the blue circles represent a Pareto frontier. For each desired level of representation of students eligible for free/reduced lunch they graph the highest average composite score than can be maintained. The red squares reflect the efficacy of minority quotas as a tool to increase the representation of low-

income students: we plot that percent free/reduced lunch and the average composite score that results from using minority quotas of between 20% and 70%. The green triangles show the effect of CPS-like policies that reserve varying fractions of the seats for students in each SES tier.

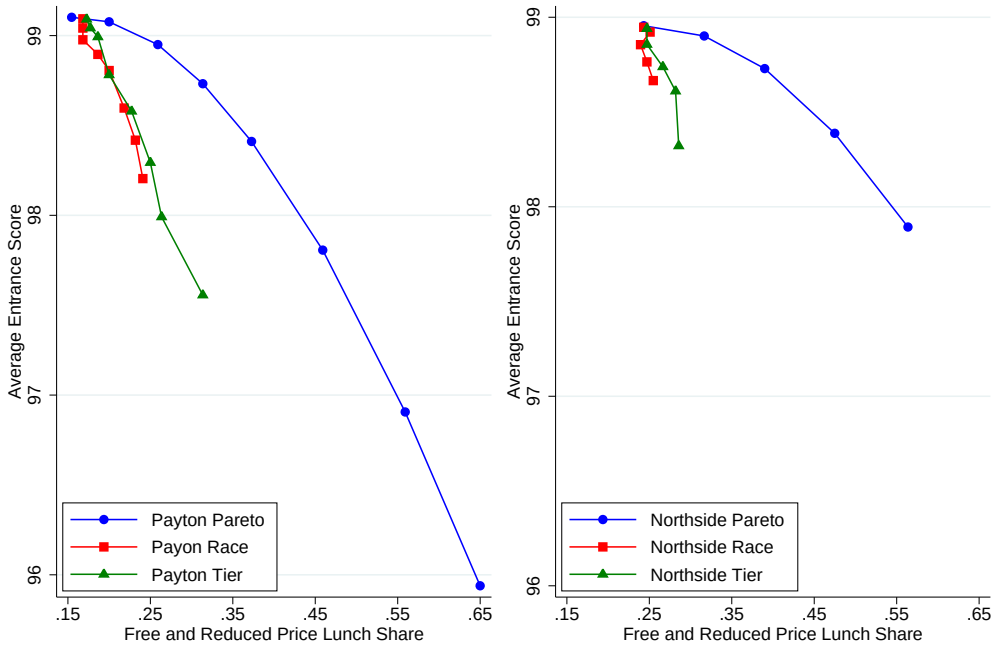


Figure 6: Feasible composite scores at various levels of free/reduced lunch eligibility: CPS and minority quotas vs. the Pareto frontier

Comparing either of the latter two policies to the Pareto frontier indicates that both are fairly inefficient tools for increasing the number of low-income students at Payton. The current CPS policy, for example, can be seen as increasing the proportion of students on free/reduced lunch from 13.8% to 27.6% at the cost of reducing average composite scores from 98.9 to 97.8. However, a simple free/reduced lunch quota could have achieved the same increase in free/reduced lunch representation while maintaining an average composite score of 98.7. Accordingly, our efficiency metric would say that the CPS policy is only 16% efficient as a method of increasing low-income representation. An alternate way to think about the efficiency of the policy is that given the realized decrease in average composite scores, one could have increased the proportion of students on free/reduced lunch all the way to 47% rather than just to 27.6% had this been the only goal.

Increasing low-income representation was clearly not CPS's only goal, so it is not surprising or damning that the CPS policy is not very efficient on this dimension. An observation that is more surprising, however, is that the CPS policy is only a little better than a racial quota policy as a tool for increasing the representation of low-income students (XX% vs. XX% for the racial quota policy. Apparently, even though the CPS policy is explicitly targeting low-income students and the racial quota system was not, the fact that a student is a minority must be roughly as good a predictor for being poor as is living in a low-SES census tract.

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 even less effective here than it was at Northside: our efficiency metric says that it is just XX% efficient. What is most striking in the graph is that the CPS tier-based system now appears to no advantage at all over the minority-quota. This indicates that for any given level of average composite scores, the SES-based policy is not only admitting fewer minorities but has no offsetting advantage in terms of admitting more low-income students.

5.3 Position relative to the three-dimensional Pareto frontier

In this section we use graphs similar those in section 4.3 to describe the efficiency of the CPS system as a method to simultaneously increase both minority and low-income representation.

We noted in section 4.3 that a range of class compositions are possible with the average composite score held fixed at any given level below the highest possible: we can maximize minority representation by using an appropriate minority quota, maximize low-income representation by using an appropriate free/reduced lunch quota, or achieve a range of outcomes in between by using bonus schemes that add appropriate numbers of points to the composite scores of minority and low-income students. The left panel of Figure 7 graphs two Pareto frontiers for Payton. Focus first on the upper frontier, which gives the set of class compositions possible at Payton when one reduces the average composite score to 97.8 (as CPS did under its 70% set aside policy).

We have also highlighed two points in the figure. The star that is embedded in the lower curve indicates the class composition at Payton under the CPS policy: the class is 40.8% minority and about 27.6% free/reduced lunch eligibile. The fact that this point is well below the frontier reflects inefficiency of the policy. We earlier discussed the CPS policy as a method for just increasing minority representation (or just increasing free/reduced lunch eligible students). The CPS policy is clearly

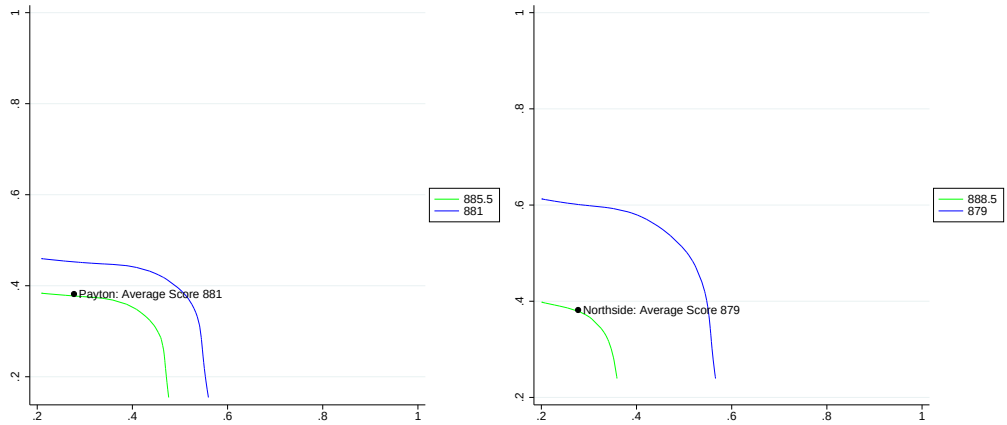


Figure 7: The CPS race-neutral policy relative to the Pareto frontier in minority/low-income space

designed to do more than just increasing minority representation. Accordingly, we feel that a better way to assess its inefficiency is to ask to what degree average composite scores had to be reduced if one wanted both to increase minority representation to 40.8% and to increase the fraction of students on free/reduced price lunch to 27.6% as the CPS policy does. The answer to this question is XXX. The dashed line going through the star is the Pareto frontier for Payton holding the average composite score fixed at XXX. Hence, as a device to serve this dual goal the CPS policy can be regarded as YYY percent efficient: YYY of the reduction in composite scores (from 98.9 to XXX) was necessary with the remainder (from XXX to 97.6) being collateral damage from restricting the form of affirmative action.

Another way to think about the magnitude of the inefficiency is by thinking about what else is possible if one is willing to reduce the average composite scores to 97.6. If one moves out along the line connecting the large circle (which reflects class composition under purely score-based admissions) to the star, the line intersects the Pareto frontier at a point where 56% of the class are underrepresented minorities and 37% are eligible for free lunch.

The right panel contains a similar figure for Northside. The star again marks the class composition under the current CPS policy: Northside is XXX% minority and XXX% free-lunch eligible. Again, this point is well inside the Pareto frontier. In this case we find that the CPS policy is XXX % efficient as a method for reaching the levels of minority and free-lunch representation that it achieves.

6 Within-School Heterogeneity and Achievement Gaps

In this section we discuss some effects of the shift to race neutral affirmative action on the composition of the student bodies at Payton and Northside

6.1 Within-school score heterogeneity

As noted earlier a primary motivation for elite public schools is that they make it feasible for schools to implement a curricula suited to the knowlege and abilities of the city's most advanced students. Within-school heterogeneity in preparation/ability can hamper schools' abilitites to realize this potential benefit. In this section we present some simple histograms illustrating the effect of racial quotas on within-school heterogeneity in composite scores.

The top three histograms in the left panels of Figure 8 illustrate how the distribution of composite scores within Payton changes as admissions shifts from a purely score based model, to a system with 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). The top left panel illustrates that in a purely score-based system there is very little heterogeneity in composite scores: about XXX percent of students have scores of at least 99, XXX percent have scores of 98 to 99, and the remaining XXX percent all scored at least 97. The second histogram corresponds to the 20% set aside. Roughly, a set of students with scores between 98 and 99 have been replaced by a set of students with scores between 96 and 97. Increasing the set-aside to 40% as in the bottom panel causes an even more notable increase in heterogeneity. On the high end most of the students with scores between 98 and 99 have been displaced. And on the low end the bottom tail extends slightly below 93.

The bottom histogram on the left-side of Figure 8 illustrates the distribution of composite scores at Payton under the current CPS policy. The most striking feature is that the left tail of low-scoring students now extends all the way down to 89. Concretely, this means that teachers must teach classes that are substantially more heterogeneous in student preparation. Note also that this occurs despite the fact that the class is only 36.8% minority – a racial composition much closer to that with the 20% set aside that to that with the 40% set aside. The difference also helps one understand why the curve of class compositions feasible under the CPS tier approach is so steeply sloped – replacing an extra unrestricted slot with a tier 4 slot will be replacing a student with 99th percentile scores with a student with 89th percentile scores and there are no guarantees that the replacement will actually improve minority representation.

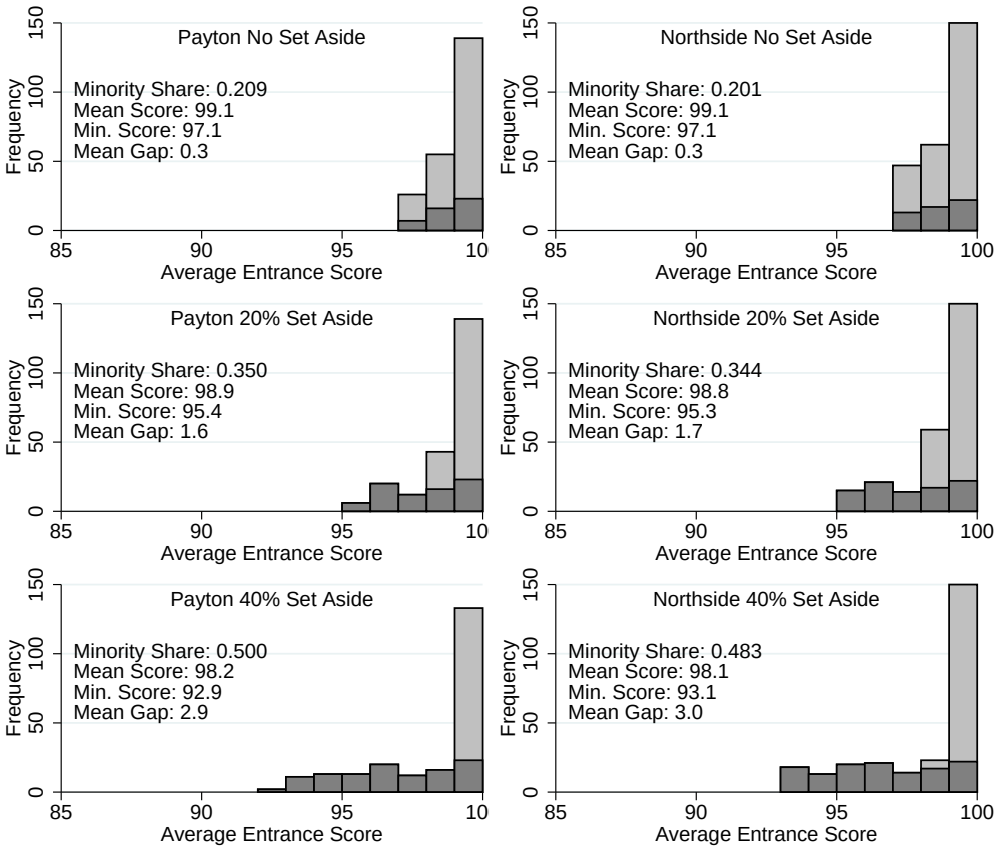


Figure 8: Score distributions with various minority quotas

The right panels present similar graphs for Northside. They show roughly similar patterns. The upper histograms corresponding to the use of minority set asides are very similar – the set of minority students interested in attending Payton and Northside are very similar so in both cases the lower tail extends down to about 93. The bottom histogram on the impact of the CPS policy illustrates our earlier comment that there is more of an efficiency loss from the CPS policy at Northside. The lower tail under the CPS policy extends further down at Northside (to 87) than at Payton.

6.2 Within-school achievement gaps

A concern about affirmative action policies is that they may contribute to within-school achievement gaps. Such gaps can foster the development of racial stereotypes that harm minority students. 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 achievement gaps would be very, very small under a purely score-based admissions policy: all admitted students have scores of at least 97 and there are minority students within all three bands. The second row of histograms, corresponding to a 20% minority set aside, show a developing score gap. Almost all majority students have scores above 98, whereas the minority students scores are more spread out with a median of 98. The gaps widens again when we move to a 40% minority set aside. Almost all of the majority students now have scores of 99. And the minority scores are spread throughout the 90s.

Switching from a racial quota to CPS's tier-based affirmative action will affect 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 whites and Asian American students who live in low SES neighborhoods. Working in the opposite direction, however, is the fact that a number of relatively high-scoring minority students living in medium to high SES neighborhoods are now being denied admission and many of the lowest-scoring admits (all of whom come from tier 4) 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 is that the shift increases the within-school. A comparison fo the second and fourth rows in Figure 8–recall that minority representation is similar under these two

processes–brings this out pretty clearly. At both schools it turns out that essentially all admitted students with scores below 93 under the CPS plan are underrepresented minorities and there is a substantial reduction in the number of admitted minority students with scores between 96 and 98. The CPS plan is adding some majority students with scores between 93 and 97, but not nearly enough to offset the reduction in average minority scores.

Figure 9 provides another easy-to-read illustration. For each level of minority representation the figure graphs the difference between the average composite score of admitted majority and minority students at Payton and Northside. All four curves have a value of around 0.2 at the left endpoint reflecting the very small difference in the average scores of admitted majority and minority students under purely score-based admissions process. The solid blue lines illustrate what happens under racial quota policies with varying numbers of reserved slots. The 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 a little less than 3 points at 50% minority.

The dashed red lines are give corresponding figures for policies based on the CPS tiers: for each desired level of minority representation we find the CPS tier reservation percentage that will produce that level of minority representation; and we graph the difference between average majority and minority composite scores under that tier-based policy. The primary observation we would make about the graphs is that within-school majority-minority achievement gaps are larger under the tier-based policy given any fixed level of minority representation. At Payon the differences are very 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. This can be regarded as an additional disadvantage of switching from a race-based to a race-neutral policy.

7 How Efficient Can "Race Neutral" Policies Be?

Although the CPS socioeconomic affirmative policy is among the most sophisticated plans that have been implemented, there are a number of ways in which it seems crude. The tract-level SES measure is an unweighted sum of the tract's rank on six characteristics; the SES measures are used only to divide the tracts into four tiers which ignores variation within each tier; and the assignment does not include any student-level covariates including an obvious correlate of disadvantage: whether the student is eligible for free or reduced lunch. In this section we explore the CPS design and ask how much more efficient the CPS plan could be made with a slightly more sophisticated use of these variables.

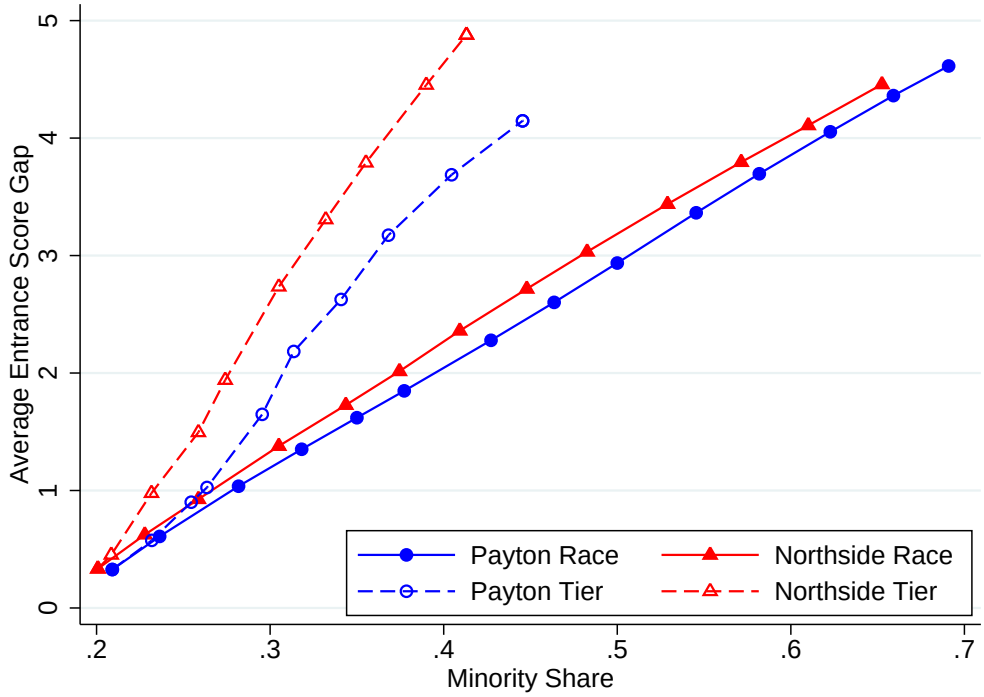


Figure 9: Within school gaps in average composite scores: quotas vs. CPS tier-based selection

7.1 Correlates of underrepresented minority status

An SES-based affirmative action plan will increase minority representation if there are more minorities among the students it advantages than among the students who are displaced. While impacts can be subtle in a complex assignment scheme, a rough intuition is the extent to which this occurs will be greater if the index is more correlated with minority status within the population of students close to the admission margin. There are a number of reasons to think that the CPS plan may not be optimal from this perspective: (1) variables individually correlated with minority status would not receive equal (or even necessarily positive) weight in a multivariate regression; (2) it is not clear if variables correlated with minority status in the full population remain positively correlated when conditioning on students' having achieved very unusual test scores; and (3) linear functional forms may not be the most appropriate. In this section we examine relatioships between minority status and the SES variables used by CPS.

The first column of Table 1 presents estimates from six logit regression run on the full dataset of students applying to CPS selective entry high schools. The dependent variable in each regression an indicator for being an underrepresented minority. Each has a single explanatory variable: one of the SES indicators in the current CPS formula. Each of these variables is entered as a percentile within all census tracts with higher percentiles corresponding to being of higher SES (at least as CPS uses the variable). Each of the estimates is negative verifying that the variables are individually associated with minority status. The second column of Table 1 presents estimates from a logit regression of the minority indicator on all six variables. The results suggest CPS's equal-weights index may be quite different from the index that would be most aligned with minority status: only four of the six variables are significantly negative in this regression and there are substantial differences in the weights they receive. One notable feature is that the coefficient on the "percent owner occupied" variable is significantly *positive* in this regression. This suggests that including this variable in addition to the others may tend to disadvantage minority students. Chicago has some predominantly African American 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.

The third column of Table 1 presents regressions more directly relevant to the admissions policies of Payton and Northside by restricting the sample to students with composite scores or at least 96.

	Dep. Variable: Minority				
	(1)	(2)	(3)	(4)	(5)
MedianIncome	-4.304	-0.023	0.602	0.433	
	(0.097)	(0.189)	(0.629)	(0.626)	
AdultEducation	-3.755	-2.589	-3.492	-3.782	-0.823
	(0.091)	(0.184)	(0.629)	(0.621)	(0.070)
DualParent	-4.250	-1.753	-0.996	-1.011	
	(0.095)	(0.148)	(0.486)	(0.483)	
HomeOwnership	-1.445	0.865	1.266	1.244	
	(0.074)	(0.124)	(0.403)	(0.399)	
NativeEnglish	-0.640	-0.537	-1.248	-1.121	
	(0.071)	(0.182)	(0.683)	(0.671)	
ISATScores	-5.231	-3.233	-2.048	-1.900	
	(0.109)	(0.146)	(0.440)	(0.434)	
StudentFreeLunch			0.719		0.127
			(0.187)		(0.032)
NativeEngMax70					-0.663
					(0.112)
ISATilelt8					0.087
					(0.027)
TwoParentilelt80					0.044
					(0.079)
N	16,381	16,381	1,021	1,021	1,021
R2		0.289	0.204	0.192	0.224

Table 1: Predictors of minority status

Our first observation–visible in the R^2 being lower in the third column than in the second–is that it is harder to use demographics to predict whether high-scoring students are minorities. Intuitively, the students have very unusual test scores for their neighborhood, so it should not be so surprising if they are also unusual demographically. The fourth column adds an indicator variable for whether the student qualifies for free or reduced lunch. Adding this variable seems clearly desirable if one wants to increse low-income representation and turns out that it is also a strong predictor of minority status. Among the demographic variables the Adult Education variable has the largest estimated coefficient, and the variables reflecting ISAT scores at the local elementary school, the fraction of native English speakers, and the fraction of two-parent households also remain signficant.

The previous regression suggests that including a homeownership variable and a high-income variable and treating these as "disadvantages" could increase minority representation, but we assume that this would not stand up to popular scrutiny and might also be legally problematic. To provide more a more detailed view of what the variables are doing and incorporate the constraint that students who are are "disadvantaged" in some way cannot be punished for this, Figure 10 presents graphs of coefficient estimates from a related constrained linear model. In this model we again use a minority indicator as the dependent variable and estimate the model on the set of students with composite scores of at least 96. The explanatory variables are dummy variables for whether each CPS variable belongs to each decile, and we constrain the coefficient estimates to be such that they are monotone decreasing.

Comparing the magnitudes of the effects the adult education variable appears to be the strongest predictor of minority status. The pattern of the coeffiicent estimates suggests that a linear functional form would be reasonable. The native English speakers variable has the second strongest effect. Note though that the model suggests that tracts with very high fractions of native English speakers should not be disadvantaged relative to 75th percentile tracts. This reflects that African American neighborhoods tend to have fewer nonnative English speakers than do neighborhoods that are predominantly white and Asian American. The ISAT and dual-parent variables also appear to be of some value. In each case the estimates suggest that a nonlinear functional form may be preferable. The coefficient estimates on the per-capita and home ownership variables are small. This suggests that including these using these variables in an SES index will not increase minority representation.

Our impresson is that school systems would prefer that affirmative action policies be based on very simple formulas. Accordingly, we estimate a very simple linear model motivated by the above graph. We regress an indicator for whether a student with a score of at least 96 is a minority on a

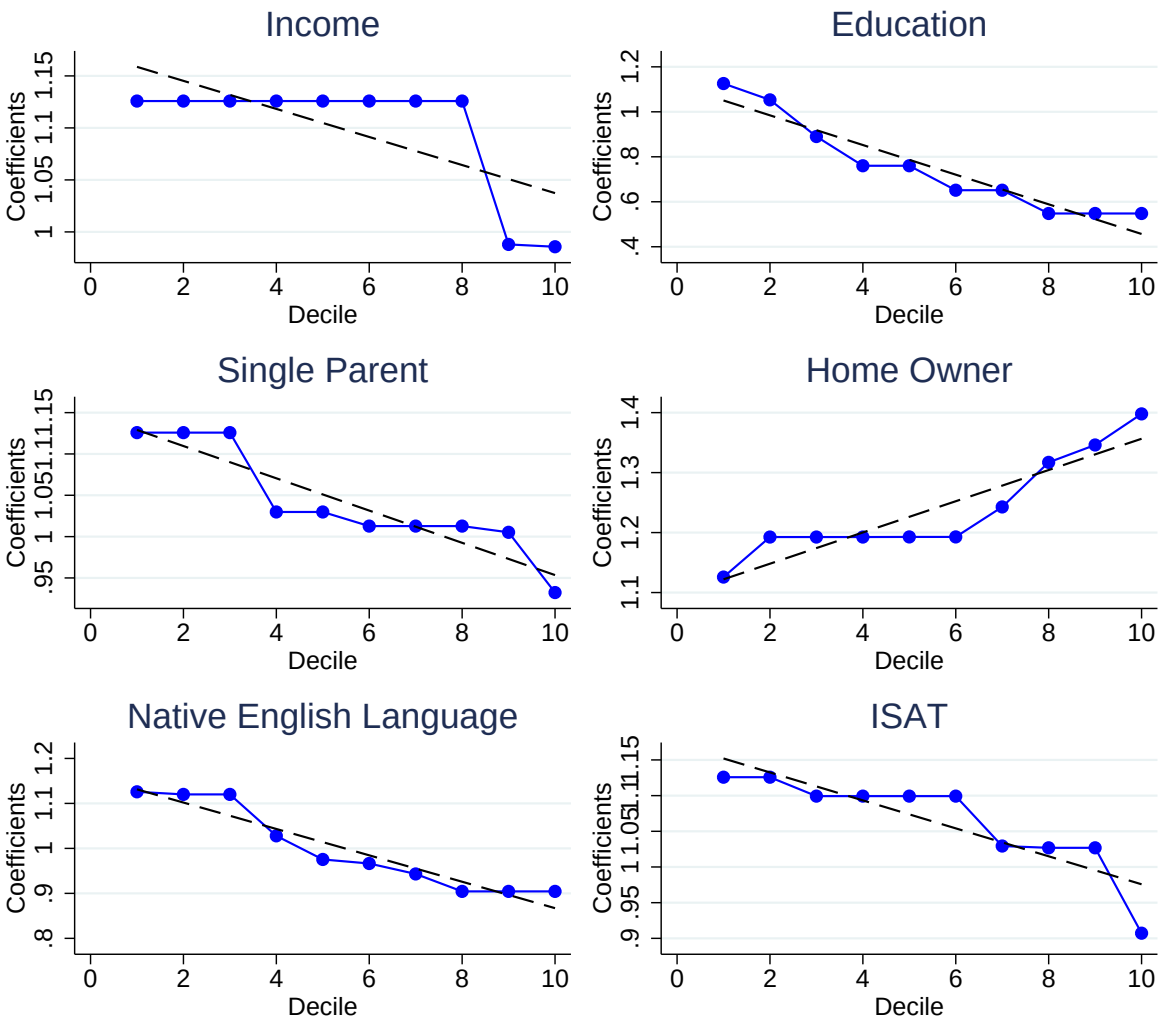


Figure 10: The CPS variables as predictors of minority status in a flexible multivariate regression

student-level free/reduced lunch indicator and four tract-level variables. One of these is the CPS's current adult education variable. Another the native English speaker percentile with its value capped at 70. A third is a dummy variable for a tract's ISAT percentile being in the lowest 80 percent. The fourth is an indicator for the dual-parent family percentile being in the lowest 80 percent. Coefficient estimates from this regression are presented in the final column of Table 1. The R^2 of this regression is XXX indicating that it has more predictive power than the linear regression and retains most of the predictive power of the regression with a large number of dummy variables.

7.2 A more efficient race-neutral policy

In this section we examine the performance of a race-neutral affirmative action policy motivated by the above regressions. First, we define a *Tract SES* variable that reflects disadvantage by taking the negative of the predicted values of the above regression and rescaling them so that they range from 0 in the most advantaged census tract (tract XXX in XXX) to 1 in the most disadvantaged tract (tract XXX in XXX). In terms of the tract characteristics this is:

$$TractSES = XXX - XXX\,AdultEd - XXX\,NatEngMax70 - XXX\,LowISAT - XXX\,LowTwoParent,$$

with the last two variables being indicators for whether the tract is below the 80th percentile. The right panel of Figure 11 presents a histogram of the *TractSES* measure across census tracts. Note that few tracts are at either extreme and that there is a large number of low-SES tracts with indexes around 0.2. The applicant-weighted histogram on the left looks somewhat different reflecting that students from high SES tracts are more likely to apply to selective entry high schools.

(SHOULD WE CUT THE HISTOGRAMS?)

We then assign schools using a simple bonus scheme. We define an adjusted composite score by

$$AdjustedScore = CompositeScore + w_1 Free/ReducedLunch + w_2(1 - TractSES)$$

for some weights w_1 and w_2. And we assign students to schools by going down the list from the highest adjusted score to the lowest and assigning each student to their most preferred school that is not already full.

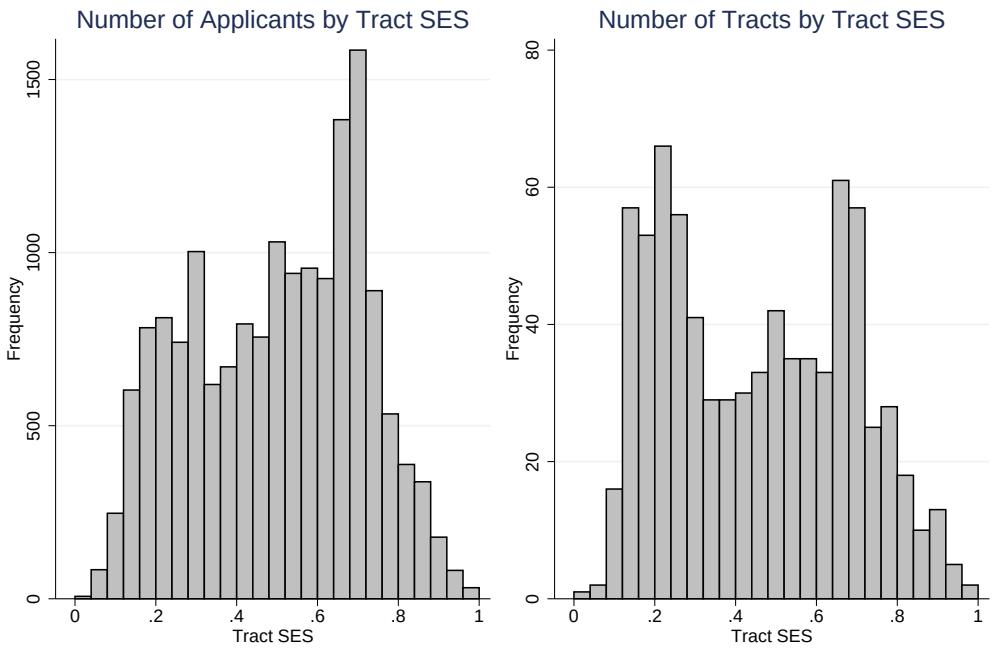


Figure 11: Histograms of Tra

7.2.1 Minority representation

To illustrate the power of such schemes to increase minority representation, Figure 12 reproduces Figure 5 with the addition of thick dashed lines that describe the outcomes when implementing the bonus scheme various weights w_1 and w_2. The left panel graphs average composite scores possible given each level of minority representation at Payton. The Tract SES/Free Lunch bonus schemes are noticeably closer to the efficient frontier than are the CPS Tier-based schemes. The right panel provides a corresponding graph for Northside.

Figure 12: Feasible composite scores at various levels of minority representation for the Tract SES/Free Lunch bonus scheme

7.2.2 Meeting multiple goals: minority and low-income representation

In the previous section we provided graphs of the Pareto frontier of levels of Minority/Free Lunch representation that are feasible while maintaining average composite scores at various levels. Achieving class compositions on those frontiers requires that a school system use race-conscious policies assigning bonus points to minorities and to students qualifying for free lunch. A school system can similarly choose from a range of class compositions with scores held fixed by varying the weights in our bonus scheme. (Roughly think of varying relative weight on the two bonuses while keeping the sum fixed although it is harder than this in practice.) Figure provides graphs of the Pareto frontier and of the feasible-frontier of compositions that can be obtained by varying the weight in our scheme.

The left panel provides graphs for Payton. The solid line is the Pareto frontier with the average composite score held fixed at XXX – its value under the CPS Tier70 plan. The dashed line is the frontier of minority/free lunch combinations that are feasible using bonus schemes with scores held fixed at the same level. The curves coincide at the upper left point – if one only wants to maximize free-lunch representation a bonus scheme putting zero weight on the tract score is efficient. The frontier from our bonus scheme slopes downward much more quickly than the feasible frontier and the most it could achieve is to raise minority representation to XXX, whereas a racial quota could reach XXX. A positive result is that the dashed frontier is well above the point marked with a star, which is the class composition under the current CPS plan. This indicates that substantial increases in minority and/or free lunch representation could be made without lowering average composite scores.

The right panel provides comparable graphs for Northside. In this figure the dashed curve goes almost straight down from its upper-left endpoint. This indicates that the schemes that we have described cannot do much more to increase minority representation beyond what could be achieved simply by giving bonuses to students qualifying for free/reduced lunch. The maximum feasible minority representation of XXX is somewhat better than the XXX that is obtained by the current CPS plan. The fact that both numbers are low is another indication of the difficulty of finding race-neutral alternatives that can effectively raise minority representation in an environment like that of Northside.

7.3 Declined perfect scorers

One potential concern with bonus schemes is that some students from highly "advantaged" tracts may be unable to gain admission regardless of how well they score.This is particularly a worry when admissions criteria topcode scores on admissions tests in a way that prevents exceptionally talented students from gaining much advantage.[11] Chicago's formula, which is based on adding percentiles, is worrying from this perspective. No matter how well a student does on the test it is impossible to gain any advantage over any other 99th percentile student and 99th percentile students take almost all of the unrestricted slots under the current system. Once one starts to add bonus points to some students scores, it is very possible that a student who is not receiving any bonus points will be denied admission even if he or she has a perfect composite score.

In practice whether students who had perfect composite scores would all be admitted to their first choice school under the plan we described depends on the degree of minority/free lunch representation one is trying to achieve. To illustrate this consider two potential implementations we label "Bonus Scheme A" and "Bonus Scheme B". Bonus Scheme A with weights $w_1 = XXX$ and $w_2 = XXX$ uses the greater efficiency of the bonus scheme to increase average composite scores at Payton and Northside while being satisfied with roughly matching the current minority and free lunch percentages. In this plan we are close to, but just short of the point where some perfect scorers would be denied admission: the minimum adjusted scores necessary for admission turn out to be XXX at Payton and XXX at Northside. Bonus Scheme B with with weights $w_1 = XXX$ and $w_2 = XXX$ uses the greater efficiency of the bonus scheme to increase minority and free lunch representation while being satisfied with roughly matching the current average composite scores. Under this plan the cutoff scores

[11] On the SAT, in contrast, the very top of the 99th percentile can score 2400 whereas the bottom is just 2200. This 180 point spread is greater than the difference between the lowest 99th percentile score (2220) and the lowest 95th percentile score (2050).

become XXX and XXX so that some students with perfect average composite scores would be denied admission.

The current Chicago plan does not have this concern because in a school with N slots any student who has one of the top 0.475 N scores will be admitted either via an unreserved slot or via a slot reserved for thier tier. The tract-based bonus scheme we have described could be modified in a similar manner to ensure that this is not a concern here either: a modified plan would first assign some fraction of the slots to the students with the highest composite scores and to then assign the remainder on the basis of adjusted scores. If the fraction reserved for high scoring students was XXX or less, then the outcome under Bonus Scheme A would have been identical to that we have described in the data for the year we studied. Implementing some reservation of this type seems advisable. It would not have any real effects if score distributions and application patterns are similar to what we see in the 2013 data. And it would reassure parents and avoid undesirable outcomes that could arise if distributions turn out to be somewhat different.

8 "Top 10%"-Style Rules

Texas's "Top 10% Rule" has probably received more attention than any other race-neutral 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.[12] Similar plans have been enacted or considered elsewhere In this section we examine how such plans might work in the Chicago context and find that they seem much worse than the current CPS procedure.

We consider here plans that allocate a fraction Y of the seats at Payton and Northside to the students with the highest composite scores reserving the remaining $1 - Y$ fraction of seats to be allocated via 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 200 to 250 students per year. Chicago contains roughly 800 census tracts and CPS operates 478 elementary schools. Accordingly a plan to admit the top student from each elementary school or from each census tract will be infeasible. It would be particularly infeasible to consider a plan which like Texas' focuses purely on grade-point-averages: each school has many, many students with perfect GPAs. Motivated

[12] More 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 which is the subject of recent course cases.

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 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 school 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.[13] Seats are allocated to the students with the lowest composite score rank until the reserved seats are filled. Note that students from relatively small tracts not gain admission even if they have the single highest score in their tract.[14]

2. 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.[15] 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 13 presents our standard illustration of the efficiency of these schemes as a method to increase minority representation. Again, we have put percent minority on the x-axis and average composite score on the y-axis. The solid black line is the efficient frontier which looks a little different here because have changed the scaling of the y-axis. The performance of the two top 10% plans are described by the thick dashed lines. Note that each of these plans fares very poorly. (Indeed, the poor performance that forced us to change the y-axis scale from that which we have used throughout.) Even allocating a very small fraction of the seats leads to substantially lower composite score. At Payton, for example, even increasing minority representation to 30% entails reducing the average composite score to 96.7 – a level at which the CPS tier approach would have 50% minority representation and at which the minority share could be over 75% under a racial quota. The series stop at XXX% minority because the plans cannot increase minority enrollment beyond this amount: this is the result

[13] If two students are tied for the highest composite score each would receive a rank of 1.5.

[14] Given the way in which we are treating ties even students living in large tracts can be denied admission despite a perfect score if there are enough other perfect scorers in their tract.

[15] Census 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.

of allocating 100% of the seats according to the top $X\%rule$. The plan based on neighborhood areas performs somewhat better–it can admit two students from one census tract in a neighborhood and none from the others if this is where the best-students are found–but it is still much less efficient than the current CPS plan.

Figure 13: Feasible composite scores at various levels of minority representation: Top 10%-style rules

9 Recap and Conclusions

We have discussed a number of different affirmative action plans in this paper. Table 2 provides some statistics on each of them. The first row gives statistics on a purely score-based admissions system. As we noted earlier about 20% of each school would be underrepresented minorities and the majority-minority score gap would be very small. The second row gives statistics on two race-conscious plans: the 20% set aside plan produces about the same levels of minority representation as the current CPS plan; and the 40% set aside produces abou the same levels as in the final year of the old quota plan. The next two rows contain statistics on the CPS tier-based plan when reserving 60% and 70% of the seats to be allocated by SES tier. The next two rows present data on the more efficient bonus schemes we described. And the final two row present data on our Top 10%-style plans.

Admissions Criterion	Average Composite Score	Majority-Minority Score Gap	Percent Minority	Percent Free/Red. Lunch	Measure of Efficiency: Minority	Min.+Lunch
Walter Payton College Prep						
Pure score-based	892	2.9	0.21	0.15		
Minority set aside 20%	890	14.6	0.35	0.19		
Minoirty set aside 40%						
Racial quota	890	14.6	0.35	0.19		
CPS Tiers 60%	885	23.6	0.34	0.25		
CPS Tiers 70%	882	28.6	0.37	0.26		
Bonus Scheme A						
Bonus Scheme B						
Top-10 Tracts						
Top-10 Neighborhoods						
Northside College Prep						
Pure score-based	891	3.0	0.20	0.24		
Minority set aside 20%	889	15.5	0.34	0.25		
Minority set aside 20%						
CPS Tiers 60%	882	29.8	0.33	0.31		
CPS Tiers 70%	880	34.1	0.36	0.33		
Bonus Scheme A						
Bonus Scheme B						
Top-10 Tracts						
Top-10 Neighborhoods						

Table 2: Comparison of various admissions criteria

References

Chan, J., and E. Eyster (2003): "Does Banning Affirmative Action Lower College Student Quality?," *American Economic Review*, 93(3), 858–872.

Epple, D., and R. Romano (1998): "Competition Between Private and Public Schools, Vouchers, and Peer Group Effects," *American Economic Review*, 88(1), 33–62.

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," .

——— (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.

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 order effect on mistmatch.[16] 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.[17] The fraction n_1 of students at school 1 have a change of $-d\delta/n_1$, and the fraction n_2 students in school 2 have a change of $-d\delta/n_2$, 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. □

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

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

The sum of the two variances expands as

$$n_1\left(E(\theta_\ell^2|s(\ell) = 1) - E(\theta_\ell|s(\ell) = 1)^2\right) + n_2\left(E(\theta_\ell^2|s(\ell) = 2) - E(\theta_\ell|s(\ell) = 2)^2\right)$$

$$= E(\theta_\ell^2) - n_1(E(\theta_\ell|s(\ell) = 1)^2) - n_2(E(\theta_\ell|s(\ell) = 2)^2).$$

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

[16] This 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 $(m - m_i)$.

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

into the above formula gives

$$\begin{aligned} W_e &= E(h(\theta)) - k\left(E(\theta^2) + n_1 E(\theta)^2 + 2n_1 n_2 E(\theta)\Delta + n_1 n_2^2\Delta^2 + n_2 E(\theta)^2 - 2n_1 n_2 E(\theta)\Delta + n_1^2 n_2 \Delta^2\right) \\ &= E(h(\theta)) - kE(\theta^2) + kE(\theta)^2 + kn_1 n_2(n_1 + n_2)\Delta^2. \\ &= E(h(\theta)) - k\mathrm{Var}(\theta) + kn_1 n_2 \Delta^2. \end{aligned}$$

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. □