

EFFORT, RACE GAPS AND AFFIRMATIVE ACTION: A STRUCTURAL POLICY ANALYSIS OF US COLLEGE ADMISSIONS

BRENT R. HICKMAN

ABSTRACT. Using the college admissions model of Hickman (2010), I study the implications of Affirmative Action (AA) in US college admissions for student academic achievement (prior to college) and college placement outcomes. I argue that the competition among high-school students for college seats is similar to a multi-object all-pay auction. The link to auctions provides access to a set of tools which create tractability and empirical power. This allows me to compare various college admissions policies (*i.e.*, allocation mechanisms) in terms of three criteria: (*i*) the induced level of overall academic achievement, (*ii*) the racial achievement gap, and (*iii*) the college enrollment gap.

To estimate the model I first develop a method for measuring AA practices in the entire college market without the need for student-level application data. Then I recover distributions over student heterogeneity using empirical auctions techniques which avoid imposition of distributional assumptions. These estimates facilitate a set of counterfactual experiments to compare the effects of the estimated US policy with alternatives not observed in the data: color-blind admissions and quotas. AA policies as implemented in the US significantly diminish the enrollment gap, but at the cost of lower academic effort on average, and particularly among talented minorities. A ranking between the color-blind rule and the US policy is ambiguous without knowledge of a social choice function assigning weights to criteria (*i*)-(iii). In contrast, a quota system produces a substantial improvement on all 3 criteria, relative to both alternatives. However, quotas are illegal in the US and cannot be implemented as such. Nevertheless, I propose a variation on the AA policy already in place that is outcome-equivalent to a quota, is simple to implement, and automatically adjusts according to the amount of asymmetry across demographic groups.

Date: **Original version: October 2009.**

Key words and phrases. Affirmative Action; all-pay auctions; admission preferences; quotas; racial achievement gap; approximate equilibrium

JEL subject classification: D44, C72, I20, I28, L53.

I am grateful to Srihari Govindan, Harry J. Paarsch, and Elena Pastorino for continual feedback and support throughout this project. I have also received helpful input from Timothy P. Hubbard, Antonio F. Galvao, Ronald Wolthoff, N. Eugene Savin, Guillame Vandenbroucke, Gustavo Ventura, and Quang Vuong. Any errors are mine alone.

I thank the following people for helping me to acquire the data I used in this project: Robert Ziomek and Justine Radunzel of ACT, Andrew Mary of the Integrated Postsecondary Education Data System of the US Department of Education, and Wayne Camara and Sherby Jean-Leger of The College Board. Brian Christensen provided research assistance in organizing the data. SAT test scores were derived from data provided by the College Board. Copyright ©1996 The College Board. www.collegeboard.com. ACT test scores and ACT-SAT concordances were provided by ACT. Copyright ©1996 ACT. www.act.org. The views expressed in this research are not the views of either The College Board, ACT or the US Department of Education.

DEPARTMENT OF ECONOMICS, UNIVERSITY OF CHICAGO, 1126 E. 59th St., Chicago, IL 60637, USA;

E-mail address: HickmanBR@UChicago.Edu;

1. INTRODUCTION

For several decades, race-conscious admission policies have been used by American colleges and universities with the objective of aiding underrepresented racial minority groups to overcome competitive disadvantages. Two persistent academic disparities among different race groups have been widely studied, and are often cited as a rationale for AA in college admissions. The first, which I shall refer to as the *enrollment gap*, has to do with racial representation in post-secondary education: among students who attend college, minorities are under-represented at selective institutions and over-represented at low-tier schools.¹ Using institutional quality measures for American colleges, I show in Section 3 that although minorities made up 17.7% of all new college freshmen in 1996, they accounted for only 11% of enrollment at schools in the top quality quartile. In the bottom quartile, minorities accounted for 29.7% of enrollment.²

The second academic disparity, known as the *achievement gap*, is typically measured in terms of standardized test scores.³ In 1996, the median SAT score among minority college candidates was at the 22nd percentile among non-minorities.⁴ These circumstances are viewed by many as residual effects of past social ills, and race-conscious college admission policies have been targeted toward addressing the problem.

Despite its intentions, much debate has arisen over the possible effects of AA on the incentives for academic achievement. Supporters claim that it levels the playing field, so to speak. The argument is that AA motivates minority students to achieve at the highest of levels by placing within reach seats at top universities—an outcome previously seen by many as unattainable.⁵ In this way, it makes costly effort investment more worthwhile for the beneficiaries of the policy. Critics of AA argue just the opposite: by lowering the standards for minority college applicants, AA creates adverse incentives for

¹Ultimately, the policy-makers care about AA because of persistent racial wage gaps. These wage gaps are related to the college admissions market in two ways: first, relatively few minorities enroll in college, and second, among minority college matriculants, relatively few end up at elite institutions. Although both are interesting aspects of the college admissions problem, in this paper the enrollment gap on which I focus concerns college placement outcomes *conditional on participation in the college market*. The implications of AA for college enrollment decisions is left for future research.

²Here, the working definition of the term “minority” is the union of the following three race classifications: Black, Hispanic and American-Indian/Alaskan Native. Institutional quality measures are based on data and methodology developed by US News & World Report for its annual *America’s Best Colleges* publication. For a more detailed discussion of the figures cited here, see Section 3. See also Bowen and Bok [1] for a discussion of racial representation among top-tier colleges and Universities.

³Although it may seem at first glance that the enrollment gap and the achievement gap are two sides of the same coin, economic theory has shown that the two are not equivalent: some admission policies can narrow one gap while widening the other. See for example Fu [8] or Hickman [11].

⁴See Section 3 for a more detailed discussion of the figures mentioned here. An extensive study of the black-white test score gap is given in Jencks and Phillips [12].

⁵Fryer and Loury [7] put forth this argument as a possible rebuttal to their “Myth 3: Affirmative Action Undercuts Investment Incentives.”

them to exert less effort in competition for admission to college. By making academic performance less important for one's outcome, they argue, AA creates a tradeoff between promoting equality and achievement. Some critics of AA go even further, bringing into question whether such policies are capable of improving outcomes for disadvantaged market players, or whether the benefits go disproportionately to economically privileged members of the targeted demographic group.⁶

While the arguments on both sides of the debate seem intuitively plausible, satisfying answers to the questions surrounding AA require an economic framework which allows for rigorous quantification of the various social costs and benefits involved. Economic theorists have weighed in on this issue before, but existing models do not facilitate quantitative comparisons among admission policies as presently implemented and competing alternatives.⁷ Moreover, existing work has also primarily focused on contrasting race-conscious admissions and race-neutral admissions, with distinctions among alternative implementations of AA being ignored. Finally, most of the existing theory has favored the viewpoint that no tradeoff exists between equality and academic output, but current models commonly violate the *Wilson doctrine*, requiring that the policy-maker be able to observe individual ability traits in order to determine ideal policy choices.

More recently, Hickman [11] has developed a model of college admissions to study the tradeoffs faced by a policy-maker with only limited information who seeks to address the problem of race gaps, while preserving incentives for academic achievement. The resulting picture is less one-sided and it indicates that the arguments of both supporters and opponents of AA are correct at some level. On the one hand, a tradeoff does exist in the sense that AA always decreases achievement by some segment of the population. Also, certain forms of AA are very ineffective at improving market outcomes for minorities. On the other hand, some varieties of AA can indeed overcome discouragement effects for disadvantaged minorities, potentially producing an increase in average achievement within the minority group. It is even possible to achieve academic performance gains among the population as a whole, while producing a more representative college admissions profile. The model also indicates that different forms of AA vary widely by their induced effort incentives, and by their effect on market outcomes. Thus, the relevant policy question is not merely whether to implement AA, but also *how* best to implement it. As it turns out, a comparison of the social costs and benefits under different admission policies cannot be resolved theoretically, and remains an empirical question.

⁶Sowell [19] expounds this argument in considerable detail; an extensive discussion of the opposite viewpoint is offered by Bowen and Bok [1].

⁷For a discussion of the previous economic theory on AA, see Hickman [11].

In this paper I estimate a structural econometric model derived from the theoretical framework of Hickman [11] in an effort to better inform the policy debate. The ultimate objective of the empirical exercise is a set of counterfactual policy experiments to compare status-quo AA practices with alternative policies not observed in the data. Using data on colleges and college entrance test scores, I empirically measure AA practices in the US college market. I then estimate the distributions of students' academic ability with tools from the structural auction econometrics literature developed by Guerre, Perrigne and Vuong [9, henceforth, GPV]. These structural estimates enable the counterfactual experiments.

The policies I analyze are "admission preferences," as implemented in the American higher education system; quotas, as implemented in India; and race-neutral admissions. However, a problem arises because it is difficult for a researcher to compare alternative policies in this context without knowledge of the social choice function. As this information is obviously out of reach, I proceed carefully by evaluating the alternatives in terms of 3 criteria: (i) academic performance, as measured by equilibrium grade distributions; (ii) the racial achievement gap, as measured by cross-group differences in grade distributions; and (iii) the enrollment gap, as measured by differences in the distributions of college seats awarded by the admissions mechanism to each demographic group in equilibrium. The only assumptions I make concerning the policy-maker's preferences are that I) she values academic achievement, II) she wishes to minimize the racial achievement gap, and III) she wishes to close the college enrollment gap. However, I make no assumptions about how much weight the policy-maker places on each objective. Therefore, the primary research objective is to characterize the costs and benefits associated with each policy, with establishing rankings between policies as a secondary objective, as it is not clear *ex ante* whether this will be possible.

The results of the empirical analysis indicate that actual AA practices in the United States significantly improve market outcomes for minority students. If AA were eliminated from college admissions decisions in the US, minority enrollment in the top quartile of colleges and universities would decrease by 33.3%. AA does decrease the quality of schools attended by non-minorities, but the change to the group as a whole is relatively smaller, amounting to a 4.2% reduction in non-minority enrollment in the top quartile. AA practices in US college admissions narrow the gap between median SAT scores among minorities and non-minorities by 14%. They discourage achievement among minority students at the upper and lower extremes of the score distribution, while encouraging students in the middle to score higher. The two effects balance each other out, so that virtually no change occurs for average minority SAT scores.

As for policy comparisons, it turns out that no clear ranking can be established between a color-blind admission scheme and the status-quo admission preference system without further information on the policy-maker’s preferences. The latter narrows the achievement gap and the enrollment gap, but the former results in higher academic achievement in the overall population of students. On the other hand, it can be reasonably argued that a quota system is superior to *both* of the other two policies on all three objectives: it produces the highest academic performance, a substantial narrowing of the achievement gap, and, by design, it closes the enrollment gap completely. Explicit quotas are illegal in the United States and cannot be implemented as such.⁸ Nevertheless, using insights from the workings of a quota mechanism, I propose a simple variation on the AA scheme currently in place, which delivers the same performance along the three policy objectives, and can be implemented using only information on race and grades. Another interesting property of this alternative policy is that it is a self-adjusting AA rule that naturally phases itself out as the racial asymmetry diminishes.

The rest of this paper has the following structure: in Section 2, I briefly outline the theoretical model on which the econometric exercise is based. In Section 3, I describe the data that will be mapped into the model. In Section 4, I outline a two-stage estimator for the structural model, similar to that of GPV. I also incorporate techniques developed by Karunamuni and Zhang [14] on boundary-corrected kernel density estimation, to overcome certain technical problems in the estimation. In Section 5, I discuss the results of estimation and the counterfactual exercise. In Section 6, I propose the alternative admission policy and I conclude. An Appendix contains technical details on certain data issues, as well as results from a robustness check.

2. THE THEORETICAL MODEL

In this section I outline the theoretical model of college admissions, and the equilibrium equations that characterize academic achievement under a given admission policy. The discussion here will be brief, but a full detailed analysis is provided in Hickman [11].

2.1. Costs of Achievement. Decision makers in the model are a set $\mathcal{K} = \{1, \dots, K\}$ of students competing for admission to college, each being characterized by a privately-known study cost type $\theta \in [\underline{\theta}, \bar{\theta}]$. The choices available to each student are grades, denoted $s \in \mathbb{R}_+$, but in order to achieve grade level s , they must incur a utility cost

⁸The US Supreme Court Ruling in *Regents of the University of California v. Bakke*, 438 U.S. 265 (1978) established the unconstitutionality of explicit quotas in the US.

$\mathcal{C}(s; \theta)$, which depends on their type. The cost function is assumed to satisfy the following regularity conditions for each $s \in \mathbb{R}_+$, $\theta \in [\underline{\theta}, \bar{\theta}]$:

$$\frac{\partial \mathcal{C}}{\partial s} > 0; \quad \frac{\partial \mathcal{C}}{\partial \theta} > 0; \quad \frac{\partial^2 \mathcal{C}}{\partial s^2} \geq 0; \quad \text{and} \quad \frac{\partial^2 \mathcal{C}}{\partial s \partial \theta} \geq 0.$$

In one interpretation, this cost structure can be thought of as resulting from an underlying labor-leisure tradeoff, and private cost types can be thought of as subsuming various external factors affecting students' academic performance such as home conditions, affluence, school quality, or access to things like health-care and tutors. However, they are not a proxy for the "effort" a student chooses to put forth, or the cost he bears in order to achieve a grade.

2.2. Rewards of Achievement. There is a set of prizes

$$\mathbf{P}_{\mathcal{K}, \mathcal{K}} = \{p_k\}_{k=1}^{\mathcal{K}},$$

where p_k denotes the utility of consuming the k^{th} prize. Students have single-unit demands, and the prizes represent college seats for which they compete. There is a seat open for every student who wishes to go to college, but not all seats are considered equally desirable; *i.e.*, $p_k \neq p_j$, $k \neq j$. Although prize values are *ex-ante* observable before effort decisions are made, for convenience they are modeled as being independently generated as random draws from an interval $\mathcal{P} = [\underline{p}, \bar{p}]$ according to a commonly-known prize distribution, $F_P(p)$. Moreover, I assume that the prize distribution has a density $f_P(p)$ which is strictly positive on \mathcal{P} . The value in framing prizes this way will become clear later on.

As a side note, it is not essential to the model for all students to place the same value on a seat at a given college.⁹ The important assumption here is that students rank prize values the same. Without this assumption, a policy discussion concerning admission outcomes is either impossible or trivial: either it will be the case that students' preferences cannot be empirically disentangled from their private costs; or the researcher is left with the unsatisfying conclusion that fewer minorities attend elite institutions simply because they prefer it that way. An alternative view of the uniform ranking assumption is that students have similar preferences over school attributes such as per-pupil spending, graduation rates, student-faculty ratios, *etc.*

⁹This model is equivalent to one in which achievement is uniformly costly for all competitors, but individuals value prizes differently and have different marginal utilities of upgrading to slightly higher ranked schools.

2.3. Demographics. Each student observably belongs to one of two groups: $\mathcal{M} = \{1, 2, \dots, M\}$ (minorities), and $\mathcal{N} = \{1, 2, \dots, N\}$ (non-minorities), where $M + N = K$. Each competitor views his opponents' private costs as independent random variables following commonly-known, group-specific cost distributions $F_{\mathcal{M}}(\theta)$ and $F_{\mathcal{N}}(\theta)$ with strictly positive densities $f_{\mathcal{M}}(\theta)$ and $f_{\mathcal{N}}(\theta)$. Although the number of competitors from each group is *ex-ante* observable, the "asymptotic mass" of minority competitors is modeled by a number $\mu \in (0, 1)$. In other words, nature assigns each student to group \mathcal{M} with probability μ , after which a private cost is drawn from the appropriate distribution. As the number of competitors becomes large, the mass of minorities approaches μ with certainty. It will become clear shortly why this assumption is useful.

Before moving on, I should mention that by allowing for asymmetry I do not intend to suggest that there are fundamental differences in inherent ability across races: private costs may reflect a myriad of environmental factors as well. Rather, this feature of the model is in keeping with arguments made by proponents of AA regarding systemic competitive disadvantages for minorities. The idea is that, on average, minority students must expend more personal effort to overcome environmental barriers (*e.g.*, poverty, poor health-care, lower quality K-12 education, *etc.*) to the learning process. There is some empirical evidence consistent with this view. Neal and Johnson [16] find that for the Armed Forces Qualification Test, "family background variables that affect the cost or difficulty parents face in investing in their children's skill explain roughly one third of the racial test score differential" (pg. 871). Fryer and Levitt [6] analyze data on racial test-score gaps among elementary school children in an attempt to uncover the causes. They find that by controlling for socioeconomic status and other environmental factors which vary substantially by race, test-score gaps significantly decrease, but not entirely. They test various hypotheses to explain the remainder of the gap, and find that disparities in school quality is the only one not rejected by the data.

In theory, cross-group academic achievement differentials can arise independently from either asymmetry or effort disincentives created by certain AA policies. Therefore, at the end of the day the data can be relied upon to indicate how much (if any) cross-group heterogeneity exists in private cost types.

2.4. College Admissions. Access to the set of prizes $\mathbf{P}_{\mathcal{K},K}$ is controlled by a College Admissions Board (henceforth, "the Board"), which receives a grade from each student and allocates prizes according to some rule.¹⁰ A simple *color-blind* rule involves assortative

¹⁰The allocation determined by The Board can be thought of as a centralized implementation of the outcome achieved by a matching market in which college candidates submit applications and individual schools reply with acceptance/rejection letters. Later on, the estimator of the college admission policy will allow for an evaluation of the assumption that all admissions boards behave similarly enough to consider them as a single entity. As it turns out, this assumption is empirically innocuous. See Section 4 for details

matching of prizes with grades: the student submitting the highest grade gets the most valuable prize, and so on. As for race-conscious admissions, a *quota* rule, similar to the one in place in India, is a mandate that a representative sample of M prizes be reserved for allocation only to minorities. Thus, under a quota the competition is split into two separate games where students compete only with members of their own group. The idea that the reserved set of prizes is “representative” can be accomplished by either randomly selecting M prizes from the set $\mathbf{P}_{\mathcal{K},\mathcal{K}}$, or it can be by ordering the elements of $\mathbf{P}_{\mathcal{K},\mathcal{K}}$ and selecting out every m^{th} prize, where $m = \frac{M+N}{M}$. Either way, when the set of prizes is large, the overall effect is the same.

Finally, American-style AA takes the form of what is referred to as an *admission preference* rule. This rule is modeled as a grade transformation function $\tilde{S} : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ that the Board uses to match prizes assortatively with non-minority grades and transformed minority grades

$$\{s_{\mathcal{N},1}, \dots, s_{\mathcal{N},N}, \tilde{S}(s_{\mathcal{M},1}), \dots, \tilde{S}(s_{\mathcal{M},M})\}.$$

Assumption 2.1. $\tilde{S}(s)$ is a strictly increasing function lying above the 45°-degree line.

Assumption 2.2. $\tilde{S}(s)$ is continuous.

Assumption 2.3. $\tilde{S}(s)$ is differentiable.

Assumption 2.1 corresponds to the notion that the policy is geared toward assisting minorities, effectively moving each minority student with a grade of s ahead of each non-minority student with a grade of $\tilde{S}(s) \geq s$. Monotonicity means that a policy-maker will not choose to reverse the ordering of any segment of the minority population, so that some students are awarded prizes of lesser value than other students within their own group whose grades were lower. Assumptions 2.2 and 2.3 imply that the policy-maker does not make sudden jumps in either the grade boost or the marginal boost. These assumptions are regularity conditions which facilitate derivation of the equilibrium.

Once an admissions rule, $\mathcal{R} \in \{cb \text{ (color-blind)}, q \text{ (quota)}, ap \text{ (admission preference)}\}$, is specified, an agent’s decision problem defines a strategic Bayesian game. Under the payoffs induced by a particular admission rule, students optimally choose grades based on their own private costs and their opponents’ optimal behavior. A (*group-wise*) *symmetric equilibrium* of the Bayesian game $\Gamma(M, N, \mathbf{P}_{\mathcal{K},\mathcal{K}}, \mathcal{R})$ is a set of group-specific achievement functions $\gamma_i : [\underline{\theta}, \bar{\theta}] \rightarrow \mathbb{R}_+$, $i = \mathcal{M}, \mathcal{N}$ which generate optimal grades, given that one’s opponents behave similarly.

2.5. Equilibrium. As it turns out, the college admissions model defined above is strategically equivalent to a special type of game known in the contests literature as an *all-pay auction*. Using analytic tools borrowed from auction theory, Hickman demonstrates

existence, monotonicity, and uniqueness of the symmetric equilibrium. The fact that equilibrium grades are decreasing in private costs implies that grade distributions, denoted by G_i , $i = \mathcal{M}, \mathcal{N}$, are generated from the equilibrium achievement functions in the following way:

$$G_i(s) = 1 - F_i(\psi_i(s)), \quad i = \mathcal{M}, \mathcal{N},$$

where ψ_i denotes the inverse achievement function for group i . At the end of the day, the grade distributions are the objects of interest to the policy-maker, as these are sufficient to evaluate policy performance along all three criteria outlined in Section 1.

One drawback to the current model is that for even moderately large $K = M + N$, the equilibrium is analytically and computationally intractable, because a decision-maker's objective function is a complicated weighted average of all the prizes, where the weight on the k^{th} best prize is one's probability of being the k^{th} lowest order statistic among K competing private costs. However, Hickman shows that if agents and prizes are generated according to the natural processes outlined above, then for large K the equilibrium of the game can be accurately approximated by considering the limiting decision problem as $K \rightarrow \infty$, effectively treating prizes and competitors as being represented by a continuum, rather than a finite set.

This simplification is useful since the number of new freshmen enrolling at American colleges and universities every year is well over a million. Under the above assumptions on model primitives, Hickman shows that the maximizers of the limiting objective functions constitute what is referred to as an *approximate equilibrium*, or a set of functions that approximate equilibrium strategies and payoffs to arbitrary precision for large enough K . Hickman shows that as the number of players grows, the sequence of finite objective functions under color-blind admissions, quotas, and admission preferences converge uniformly to

$$(1) \quad \Pi^{cb}(s; \theta) = F_p^{-1}[G(s)] - \mathcal{C}(s; \theta),$$

$$(2) \quad \Pi_i^q(s; \theta) = F_p^{-1}[G_i(s)] - \mathcal{C}(s; \theta), \quad i = \mathcal{M}, \mathcal{N}, \text{ and}$$

$$(3) \quad \begin{aligned} \Pi_{\mathcal{M}}^{ap}(s; \theta) &= F_p^{-1} \left[\mu G_{\mathcal{M}}(s) + (1 - \mu) G_{\mathcal{N}}(\tilde{S}(s)) \right] - \mathcal{C}(s; \theta) \\ \Pi_{\mathcal{N}}^{ap}(s; \theta) &= F_p^{-1} \left[\mu G_{\mathcal{M}}(\tilde{S}^{-1}(s)) + (1 - \mu) G_{\mathcal{N}}(s) \right] - \mathcal{C}(s; \theta), \end{aligned}$$

respectively

The intuition is simple. Under a color-blind rule (see equation (1)), the Board rewards achievement by mapping the quantiles of the observed population grade distribution into the corresponding quantiles of the prize distribution. Under a quota (see equation (2)), the Board maps the quantiles of the observed group grade distributions

into the corresponding prize quantiles.¹¹ To understand the workings of an admission preference, recall that for a random variable S distributed according to $G(s)$, the distribution of $T = \tilde{S}(S)$ is given by $G(\tilde{S}^{-1}(T))$. Under an admission preference, the Board rewards minorities by mapping the quantiles of the distribution over minority grades and *de-subsidized* non-minority grades into the corresponding prize quantiles. For non-minorities, the Board maps the quantiles of the distribution over non-minority grades and *subsidized* minority grades into the corresponding prize quantiles. Thus, a minority's standing relative to the opposite group changes in a positive direction and a non-minority's standing relative to the opposite group changes in a negative direction. For both, standing relative to their own group remains the same. Finally, one's payoff is the value of the prize received, minus the cost of achievement.

I shall conclude by outlining the approximate-equilibrium equations under an admission preference. Rewriting equations (3) in terms of equilibrium strategies, I get the following objective functions:

$$\begin{aligned}\Pi_{\mathcal{M}}(s; \theta) &= F_P^{-1} \left[1 - \left(\mu F_{\mathcal{M}} [\psi_{\mathcal{M}}(s)] + (1 - \mu) F_{\mathcal{N}} [\psi_{\mathcal{N}}^{ap}(\tilde{S}(s))] \right) \right] - \mathcal{C}(s; \theta) \\ \Pi_{\mathcal{N}}^{ap}(s; \theta) &= F_P^{-1} \left[1 - \left(\mu F_{\mathcal{M}} [\psi_{\mathcal{M}}^{ap}(\tilde{S}^{-1}(s))] + (1 - \mu) F_{\mathcal{N}} [\psi_{\mathcal{N}}^{ap}(s)] \right) \right] - \mathcal{C}(s; \theta).\end{aligned}$$

The approximate equilibrium is partially characterized by the first-order conditions, being

$$(4) \quad \frac{(1 - \mu) f_{\mathcal{N}} [\psi_{\mathcal{N}}^{ap}(\tilde{S}(s))] (\psi_{\mathcal{N}}^{ap})'(\tilde{S}(s)) \tilde{S}'(s) + \mu f_{\mathcal{M}} [\psi_{\mathcal{M}}^{ap}(s)] (\psi_{\mathcal{M}}^{ap})'(s)}{f_P \left(F_P^{-1} \left[1 - \left((1 - \mu) F_{\mathcal{N}} [\psi_{\mathcal{N}}^{ap}(\tilde{S}(s))] + \mu F_{\mathcal{M}} [\psi_{\mathcal{M}}^{ap}(s)] \right) \right] \right)} = \mathcal{C}'(s; \psi_{\mathcal{M}}^{ap}(s))$$

for minorities and

$$(5) \quad \frac{(1 - \mu) f_{\mathcal{N}} [\psi_{\mathcal{N}}^{ap}(s)] (\psi_{\mathcal{N}}^{ap})'(s) + \mu f_{\mathcal{M}} [\psi_{\mathcal{M}}^{ap}(\tilde{S}^{-1}(s))] (\psi_{\mathcal{M}}^{ap})'(\tilde{S}^{-1}(s)) \frac{d\tilde{S}^{-1}(s)}{ds}}{f_P \left(F_P^{-1} \left[1 - \left((1 - \mu) F_{\mathcal{N}} [\psi_{\mathcal{N}}^{ap}(s)] + \mu F_{\mathcal{M}} [\psi_{\mathcal{M}}^{ap}(\tilde{S}^{-1}(s))] \right) \right] \right)} = \mathcal{C}'(s; \psi_{\mathcal{N}}^{ap}(s))$$

for non-minorities. However, there are some caveats involved.

For example, suppose the admission preference function is such that $\tilde{S}(0) = \Delta > 0$. Then the non-minority objective function as stated above is only valid for students achieving grades exceeding Δ (for grades below this point, \tilde{S}^{-1} is negative). Letting θ_{Δ} denote the non-minority type achieving an equilibrium grade of Δ , for non-minorities

¹¹recall that, in reserving a representative prizes for minorities, the Board randomly samples M prizes from the set of K prizes generated according to F_P . By the law of large numbers, the resulting distributions of prizes allocated to each group converge in probability to F_P .

in the interval $[\theta_\Delta, \bar{\theta}]$ the policy effectively places them behind *every* minority student. Hickman shows that their limiting objective function is the same as under a quota rule, resulting in the following first-order condition:

$$(6) \quad (\gamma_{\mathcal{N}})'(\theta) = -\frac{f_{\mathcal{N}}(\theta)}{f_P \left(F_P^{-1} (1 - F_{\mathcal{N}}(\theta)) \right) \mathcal{C}'(\gamma_{\mathcal{N}}(\theta); \theta)}$$

Also, a boundary condition is needed in order to complete the solution for non-minority achievement. It is pinned down by the following assumption on the relationship between $\bar{\theta}$ and \underline{p} :

Assumption 2.4 (*Zero Surplus Condition*). $\underline{p} = \mathcal{C}(0; \bar{\theta})$

As explained in Hickman [11], the zero surplus condition can be thought of as resulting from free-entry in the market which supplies post-secondary education services and unskilled jobs to new high school graduates. Prize values are the additional utility one gains from going to college versus opting out, and $[\underline{\theta}, \bar{\theta}]$ is the set of individuals who choose to participate in the college market. If colleges and firms can freely enter the market and supply either college seats or unskilled jobs, agent type $\bar{\theta}$ —the highest cost type who decides to become educated—will be just indifferent between attending college and entering the work force as an unskilled laborer. This point highlights a limitation of the current model: it attempts to characterize student behavior *conditional on participation in the post-secondary education market*, and it is not intended to provide insights into the decision of whether to acquire additional education. Although interesting, this aspect of the college admissions problem is left for future research.

By monotonicity, a student with cost type $\bar{\theta}$ is sure to be awarded the lowest quality prize, so the assumption implies the following boundary condition:

$$(7) \quad \gamma(\bar{\theta}) = \mathcal{C}^{-1}(\underline{p}; \bar{\theta}).$$

Finally, Hickman also shows that if the admission preference policy is such that $\tilde{S}(0) = \Delta > 0$, then there may be a mass point of minorities achieving grades of zero, depending on the derivative of the transformation function at zero. If this is true, then the minority achievement function (conditional on positive effort) will have a different boundary condition than the non-minority one. This initial condition can be characterized using the following equation—obtained by substituting equation (5) into equation (4)—which relates behavior across race groups:

$$(8) \quad \mathcal{C}'(s; \psi_{\mathcal{M}}(s)) = \mathcal{C}' \left(\tilde{S}(s); \psi_{\mathcal{N}} \left[\tilde{S}(s) \right] \right) \tilde{S}'(s).$$

Thus, non-minority achievement is characterized by a piece-wise differential equation defined by (5) and (6), with boundary condition (7), and minority achievement is given by differential equation (4) with a boundary condition given by evaluating equation (8) at $s = 0$. To be complete, the functions $G_{\mathcal{M}}$ and $g_{\mathcal{M}}$, which show up in equations (4) and (5) are actually the distribution and density of minority grades *conditional on positive effort*, rather than the overall distribution (which may have a mass point) and density (which may not exist). With that understood, it will simplify the discussion to abuse terminology and simply refer to them as the distribution and density of minority grades. The distinction shall be explicitly made later on when necessary.

2.6. Policy Objectives. Equilibrium achievement functions and private cost distributions induce a set of group-specific grade distributions, $G_{\mathcal{M}}$ and $G_{\mathcal{N}}$ and a population grade distribution G . These are ultimately the objects of interest from a policy standpoint, as they fully characterize achievement, achievement gaps and enrollment gaps in equilibrium. Henceforth, the achievement gap shall be formally represented by a function $\mathcal{A} : [0, 1] \rightarrow \mathbb{R}$ defined by

$$\mathcal{A}(q) \equiv G_{\mathcal{N}}^{-1}(q) - G_{\mathcal{M}}^{-1}(q).$$

In words, \mathcal{A} characterizes the difference between minority and non-minority achievement at each quantile of the grade distributions. Thus, to eliminate the achievement gap is to accomplish an outcome where $\mathcal{A}(q) = 0, \forall q \in [0, 1]$.

As for the enrollment gap, let $F_{P_i}(p), i = \mathcal{M}, \mathcal{N}$ denote the distribution of prizes awarded to either group in equilibrium. Then the enrollment gap is a function $\mathcal{E} : [0, 1] \rightarrow \mathbb{R}$ defined by

$$\mathcal{E}(q) \equiv F_{P_{\mathcal{N}}}^{-1}(q) - F_{P_{\mathcal{M}}}^{-1}(q).$$

Once again, to eliminate the gap is to accomplish an outcome where $\mathcal{E}(q) = 0, \forall q \in [0, 1]$. Finally, the overall profile of academic achievement is represented by the population grade distribution,

$$G(s) = \mu G_{\mathcal{M}}(s) + (1 - \mu) G_{\mathcal{N}}(s).$$

Measures of cost and benefit cited in the policy debate over AA are often related to, or derived from \mathcal{A} , \mathcal{E} , or G . For example, a statement about the test score gap that “the median minority SAT score lags behind the non-minority median by 150 points,” is equivalent to the statement $\mathcal{A}(.5) = 150$. The reason for defining race gaps and achievement in such general terms is that it avoids imposing strong assumptions on what policy-makers care about. To wit, if preferences place the same weight on the enrollment gap at every point of the college quality spectrum, then \mathcal{E} could be reduced

to $\mathcal{E} = \int_0^1 \left(F_{P_N}^{-1}(q) - F_{P_M}^{-1}(q) \right) dq$; but if the policy-maker cares more about the enrollment gap at elite schools, then this would be inappropriate.

Having formalized my notion of race gaps and achievement, I shall proceed under the light assumptions below regarding the policy-maker's preferences. These in turn establish a partial ordering on the space of policy functions.

Assumption 2.5. For two achievement gap functions, \mathcal{A}^* and \mathcal{A} ,

$$\mathcal{A}^*(q) \leq \mathcal{A}(q) \quad \forall q \in [0, 1] \quad \Rightarrow \quad \mathcal{A}^* \succcurlyeq \mathcal{A},$$

and $\mathcal{A}^* \succ \mathcal{A}$ if in addition

$$\exists q^* \in [0, 1] \quad \text{s.t.} \quad \mathcal{A}^*(q^*) < \mathcal{A}(q^*).$$

Assumption 2.6. For two enrollment gap functions, \mathcal{E}^* and \mathcal{E} ,

$$\mathcal{E}^*(q) \leq \mathcal{E}(q) \quad \forall q \in [0, 1] \quad \Rightarrow \quad \mathcal{E}^* \succcurlyeq \mathcal{E},$$

and $\mathcal{E}^* \succ \mathcal{E}$ if in addition

$$\exists q^* \in [0, 1] \quad \text{s.t.} \quad \mathcal{E}^*(q^*) < \mathcal{E}(q^*).$$

Assumption 2.7. For two population grade distributions, G^* and G ,

$$G^*(s) \leq G(s) \quad \forall s \in \mathbb{R}_+ \quad \Rightarrow \quad G^* \succcurlyeq G,$$

and $G^* \succ G$ if in addition

$$\exists s^* \in \mathbb{R}_+ \quad \text{s.t.} \quad G^*(s^*) < G(s^*).$$

3. DATA

I now proceed to the empirical exercise by describing the data that will be used to recover each component of the model. Ultimately, the objects of principal empirical interest are the group-specific private cost distributions, $F_{\mathcal{M}}(\theta)$ and $F_{\mathcal{N}}(\theta)$; the demographic parameter μ ; the prize distribution $F_P(p)$; and the cost function $\mathcal{C}(s; \theta)$. These objects will enable the counterfactual experiments, which are the ultimate goal of the policy analysis. However, it will first be necessary to obtain estimates of some intermediate objects: the group-specific grade distributions, $G_{\mathcal{M}}(\theta)$ and $G_{\mathcal{N}}(\theta)$; the distributions of prizes allocated to each group under the actual AA policy, $F_{P_{\mathcal{M}}}(p)$ and $F_{P_{\mathcal{N}}}(p)$; and the actual AA policy $\tilde{S}(s)$, corresponding to the data-generating process. To identify the various model components, I use data on quality measures for colleges and universities in the US, freshman enrollment, and student-level college entrance test scores.

I use data for the academic year 1995-1996 primarily because one can reasonably assume that, prior to that year, AA policies determining payoffs were stable and known to decision-makers. In the summer of 1996 the outcome of a federal lawsuit *Hopwood v.*

Texas (78 F.3d 932, 5th Cir. 1996) was finalized, marking the first successful legal challenge to AA in US college admissions since 1978, nearly two decades before.¹² Subsequently, other potentially important changes occurred, including state laws banning AA being passed in Texas, California, and Michigan.

3.1. Prize Data. Institutional quality measures are derived from data and methodology developed by US News & World Report (henceforth, USNWR) for the purpose of computing their annual *America's Best Colleges* rankings (see Morse [15]). USNWR collects data on fourteen quality indicators for American colleges and universities each year; the sample size for 1996 was 1,314 schools. They classify the 14 indicators into 6 categories: *selectivity*, comprised of application acceptance rate, yield (% of accepted students who choose to enroll), average entrance test scores, and % of first-time freshmen in the top quartile of their high school class; *faculty resources*, comprised of % of full-time instructional faculty with a PhD or terminal degree, % of instructional faculty who are full-time, average faculty compensation, and student/faculty ratio; *financial resources*, comprised of education spending per student and non-education spending per student; *retention*, comprised of graduation rate and freshman retention rate; *alumni satisfaction*, comprised of % of living alumni contributing to annual fund drives; and *academic reputation*, comprised of a ranking measure taken from a survey of college administrators. A single index of quality is determined by computing empirical distributions for each indicator, and taking a weighted average of the 14 empirical CDF values for a given school. The Data Appendix summarizes weights and descriptive statistics for each the 14 quality indicators.

One drawback of using the USNWR method for my purpose is that it separates schools by Carnegie classification (*i.e.*, national/regional universities and national/regional liberal arts colleges) and geographic region (*i.e.*, northern, southern, midwestern and western; see Morse [15] for more details). Therefore, I alter the method slightly by combining all schools into the same set. This does not pose a problem for most of the quality indicators, except one: the academic reputation score. This score is determined by asking college administrators to rank the schools in their Carnegie class and region. Since the reputation score loses its meaning when taken outside of these smaller subsets of schools, I drop it from the list and generate the quality index with the remaining 13 indicators, uniformly spreading the reputation weight among the remaining 5 categories.

¹²On March 18, 1996 the US Fifth Circuit Court disallowed race-conscious admissions decisions at the University of Texas law school, but appeals continued for several months afterward. The outcome of the case was finalized in July when the Supreme Court declined to review the Fifth Circuit's ruling. The last successful legal challenge before *Hopwood* was in 1978, when the Supreme Court declared quotas unconstitutional in *University of California v. Bakke* (438 U.S. 265 1978).

This is of little consequence for the overall rankings, due to the high degree of correlation among the quality indicators.

With the modified USNWR quality measure in hand, I establish the uniform prize ranking by interpreting a school's quality index as a measure of prize value. More precisely, I assume that there is a linear relationship between the USNWR quality index for each school and the utility derived from occupying a seat there. I argue that interpreting the quality index as a meaningful measure of value is sensible for two reasons. First, acquiring information to rank schools and judge one's chances for admission is a costly exercise for an inexperienced high school student, but USNWR solves this problem by providing large quantities of data on many schools, along with advice on how to interpret the data. Second, the validity of the USNWR rankings is presumably reinforced in the student's mind by the enthusiasm with which so many schools advertise their status in the *America's Best Colleges* rankings. One need not search long through undergraduate admissions web pages to find multiple references to USNWR.

The other relevant data on school characteristics is freshman enrollment, provided by the US Department of Education through the National Center for Education Statistics' Integrated Postsecondary Education Data System tool. For each school in the sample, I obtained a tally of all first-time freshman enrollment (including full-time and part-time), for the following 7 racial classifications: White, Black, Hispanic, Asian or Pacific Islander, American Indian or Alaskan Native, non-resident alien, and race unknown. The data representing schools are $\{Q_u, M_u, N_u\}_{u=1}^U$, where for the u^{th} school Q_u is the modified USNWR quality index, M_u is the number of seats awarded to minorities, and N_u is the number awarded to non-minorities. There are 1,056,580 total seats open at all schools; for individual schools the median is 451 seats, the mean is 804.09, and the standard deviation is 934.78.

The above data characterize the sample of prizes and the samples of prizes allocated to each group, given by

$$\begin{aligned} \mathbf{P}_{\mathcal{K},\mathcal{K}} &= \{p_k\}_{k=1}^K = \left\{ \{p_{ui}\}_{i=1}^{M_u+N_u} \right\}_{u=1}^U, \quad p_{ui} = Q_u, \\ \mathbf{P}_{\mathcal{M},\mathcal{M}} &= \{p_m\}_{m=1}^M = \left\{ \{p_{uj}\}_{j=1}^{M_u} \right\}_{u=1}^U, \quad p_{uj} = Q_u, \\ \mathbf{P}_{\mathcal{N},\mathcal{N}} &= \{p_n\}_{n=1}^N = \left\{ \{p_{ul}\}_{l=1}^{N_u} \right\}_{u=1}^U, \quad p_{ul} = Q_u. \end{aligned}$$

The fact that there are multiple prizes in the data with the same value represents a departure from the theory, but it is a small one given that the largest school in the sample (in terms of enrollment) has a mass of only 6.6×10^{-3} , while the mean and median schools have masses of 7.61×10^{-4} and 4.268×10^{-4} , respectively. Another possible

TABLE 1. Racial Representation Within Different Academic Tiers

Tier	% of Total Enrollment in Tier	Black 11.2%	Hispanic 5.7%	American Indian/ Alaskan Native 0.8%	White 72%	Asian/ Pacific Islander 5.7%	\mathcal{M} 17.7%	\mathcal{N} 82.4%
I	35.9%	5.6%	5%	0.5%	74.8%	9.3%	11.1%	88.9%
II	26.8%	10.1%	5.3%	0.8%	75.1%	4.2%	16.1%	83.9%
III	20.4%	14.2%	6.2%	0.9%	70%	4.3%	21.3%	78.7%
IV	16.8%	21.3%	7.3%	1.2%	63.4%	2.2%	29.7%	70.3%


criticism of this approach is that the rankings are dependent upon an arbitrary weighting scheme. Critics sometimes accuse USNWR of manipulating the weights assigned to the different quality indicators, in order to alter the relative standings of elite schools. However, this objection is inconsequential if one takes the larger picture into account. Because of the high degree of correlation among the 13 quality indicators, the overall prize distribution is remarkably robust to substantial changes in the weighting scheme. While it is possible that the relative rankings of the top 10 schools are affected somewhat by such changes, the bigger picture is very stable.

Finally, I have yet to specify the distinction between groups \mathcal{M} and \mathcal{N} . I shall define the minority group as the union of the race classes Black, Hispanic, and American Indian or Alaskan Native; non-minorities are all others. This corresponds to the notion that AA policies are targeted toward groups that are under-represented at elite universities and over-represented at lower-quality schools. Table 1 provides a clearer picture of this criterion. Similarly to what is done in *America's Best Colleges*, I have sorted the schools in descending order of quality index and separated them into four tiers, each containing one quarter of the schools in the sample. Tier I comprises the schools with the highest quality indices, and so on. I compute the mass of each race group within each tier to show representation; I also list the population mass of each race group under its name. The final two columns contain figures for the aggregated race groups.¹³ I also list the percentage of students in each tier, as quality quartiles are different from quartiles in terms of enrollment.

Each of the minority race classes is under-represented in the top two tiers and over-represented in the bottom two. For whites it is the opposite, and for Asians/Pacific Islanders the difference is even more pronounced: they are heavily under-represented in every tier except the top. Similar observations hold when the groups are aggregated.

¹³The table does not list the race unknown and resident alien groups, which is why the first five population masses do not quite sum to one. However, these groups are included in the calculations for group \mathcal{N} , so the final two masses do sum to one.

FIGURE 1. The Empirical Enrollment Gap
Prize Allocation Distributions: Data




The difference in allocations is captured graphically in Figure 1. The dotted line is the empirical distribution of $\mathbf{P}_{\mathcal{K},\mathcal{K}}$, the lower dashed line is the empirical distribution of $\mathbf{P}_{\mathcal{N},\mathcal{N}}$, and the upper solid line is the empirical distribution of $\mathbf{P}_{\mathcal{M},\mathcal{M}}$. As Figure 1 shows, the non-minority allocation dominates the minority allocation in the first-order sense. For example, roughly one half of minority students attend schools with a quality index of .5 or less, while the fraction of non-minorities in that same lowest segment is only one third.

3.2. Academic Achievement Data. The remaining data used to estimate the model are college entrance test scores. For the 1996 graduating seniors cohort, I have individual-level data on composite SAT scores, race, and other characteristics for a random sample of 92,514 students, with 73,361 non-minority observations and 19,153 minority observations. SAT scores range between 0 and 1,600 in increments of 10, but I drop the final digit and treat them as ranging from 0 to 160 in increments of 1. Before moving on, it will be necessary to address a preliminary technical concern: the interpretation of an achievement level of “zero.”

Theoretically, it is possible for a student to score zero by answering all test questions incorrectly. However, such a feat is extremely difficult unless one knows enough to achieve a nearly perfect score: with a probability of virtually one, an uninformed student will get a positive score, due to the multiple-choice format of the test. A reasonable interpretation of a student with an academic achievement of zero is one who engages in random responding to all test questions. As it turns out, a randomized test-score simulation exercise indicates that the SAT score one can expect from such an uninformed student is 58 (See appendix for details).

FIGURE 2. Test Score Distributions



For the remainder of the paper, SAT scores will be normalized by subtracting 58 (observed scores below 58 are normalized to zero), and the samples of normalized test scores are denoted by

$$\mathbf{S}_{\mathcal{M}, T_{\mathcal{M}}} = \{s_{\mathcal{M}, t}\}_{t=1}^{T_{\mathcal{M}}}, \text{ and } \mathbf{S}_{\mathcal{N}, T_{\mathcal{N}}} = \{s_{\mathcal{N}, t}\}_{t=1}^{T_{\mathcal{N}}},$$

where T_i is the number of grade observations on group $i = \mathcal{M}, \mathcal{N}$. The academic achievement gap is illustrated in Figure 2 where the empirical distributions of normalized SAT scores are displayed. The median for non-minorities is 44, and the median for minorities is 29, which corresponds to the 22nd percentile for non-minorities. Figure 2 also suggests a small mass-point of minorities with scores of zero. This will serve as a partial specification test later on. The theory indicates that a necessary condition for mass-points in minority achievement is $\tilde{S}(0) > 0$.

4. THE EMPIRICAL MODEL

As Hickman [11] points out, the theoretical model outlined in Section 2 is strategically equivalent to an all-pay auction. An all-pay auction is a strategic interaction in which agents compete for a limited resource by incurring some non-recoverable cost before learning the outcome of the game. In the college admissions model, the Board is analogous to an auctioneer, who auctions off a set of heterogeneous prizes according to a pre-determined mechanism. Students are similar to bidders, and the grades they achieve are analogous to sunk payments tendered to the auctioneer, since they cannot recover lost leisure time or disutility incurred by study effort.

Empirically, this is an attractive framework since the auction econometrics literature has emerged as one of the foremost successes in empirical industrial organization over

the past two decades. Since the founding work of Paarsch [17], auction econometricians have exploited the parsimonious, one-to-one link between observable behavior and private information to recover empirically the distributions over bidder heterogeneity. The key assumption underlying the *structural approach* to estimating these models is that the theoretical equilibrium is consistent with the data-generating process. Said differently, the assumption is that observed behavior was purposefully generated by rational decision makers. This assumption shall form the basis of my estimation strategy as well.

Another landmark paper in empirical auctions is Guerre, Perrigne and Vuong [9, henceforth, GPV], who devised an estimation strategy for auctions which is computationally inexpensive and does not rely on distributional assumptions. The main idea of the paper comes from an observation about equilibrium equations in auction models which express bids as functions of private information and the (unobserved) distribution of private information. GPV recognized that these equations could be rearranged so as to express a bidder's private information as a function of his observable behavior and the (observable) distribution over all bidders' behavior.

As I will shortly demonstrate, equations (4), (5), and (6) from the college admissions model can be similarly manipulated so as to allow the econometrician to recover a sample of private costs implied by observed test scores and the distributions over test scores. However, the form of the policy function \tilde{S} plays a crucial role in defining those equations and determining how estimation should proceed (recall that equation (6) applies only if there is a positive grade boost for a minority score of zero). Therefore, I shall begin by proposing an estimator for \tilde{S} .

4.1. Estimating the Grade Transformation Function. The rules of college admissions as set forth by the Board are exogenous to the model which I have defined. I shall assume that some function \tilde{S} , as described in Section 2, is consistent with the data-generating process. This is an empirically attractive construct, because it nests a broad range of policies as special cases, including a quota and a color-blind rule. From the policy-maker's perspective, grades and race are mapped into outcomes via the following reward functions for each group:

$$(9) \quad \begin{aligned} \pi_{\mathcal{M}}(s) &= F_p^{-1} \left[(1 - \mu)G_{\mathcal{N}}(\tilde{S}(s)) + \mu G_{\mathcal{M}}(s) \right], \quad s \geq 0, \text{ and} \\ \pi_{\mathcal{N}}(s) &= F_p^{-1} \left[(1 - \mu)G_{\mathcal{N}}(s) + \mu G_{\mathcal{M}}(\tilde{S}^{-1}(s)) \right], \quad s \geq \tilde{S}(0). \end{aligned}$$

One key observation here allows for identification of the policy function:

$$(10) \quad \pi_{\mathcal{M}}(s) = \pi_{\mathcal{N}}(\tilde{S}(s)).$$

Using this fact, one can recover \tilde{S} by determining what rule could have produced allocations $\mathbf{P}_{\mathcal{M},\mathcal{M}}$ and $\mathbf{P}_{\mathcal{N},\mathcal{N}}$ from the observed grade distributions. More specifically, for $r \in (0,1)$ let $s_{\mathcal{N}}(r) \equiv G_{\mathcal{N}}^{-1}(r)$ denote the r^{th} quantile in the non-minority grade distribution. For minorities, let $r_{\mathcal{M}}(r) \equiv G_{\mathcal{M}}\left(\tilde{S}^{-1}(s_{\mathcal{N}}(r))\right)$ denote the quantile rank of the *de-subsidized* version of $s_{\mathcal{N}}(r)$ within the minority grade distribution. By Assumption 2.1 and by observation (10), it immediately follows that

$$\begin{aligned}
 (11) \quad & F_{P_{\mathcal{M}}}^{-1}(r_{\mathcal{M}}[r]) = F_{P_{\mathcal{N}}}^{-1}(r) \\
 & \Rightarrow G_{\mathcal{M}}\left(\tilde{S}^{-1}[s_{\mathcal{N}}(r)]\right) = F_{P_{\mathcal{M}}}\left(F_{P_{\mathcal{N}}}^{-1}[r]\right) \\
 & \Rightarrow G_{\mathcal{N}}^{-1}(r) = \tilde{S}\left(G_{\mathcal{M}}^{-1}\left[F_{P_{\mathcal{M}}}\left(F_{P_{\mathcal{N}}}^{-1}[r]\right)\right]\right),
 \end{aligned}$$

where the second and third lines follow from substituting and from monotonicity. Equation (11) above provides a moment condition that forms the basis of a simple policy function estimator. For a given specification of \tilde{S} one can choose a set of quantile ranks and pick the parameters of the policy function so as to minimize the distance between the left-hand side and right-hand side.

More formally, assume that

$$\tilde{S}(s) = \sum_{i=0}^I \Delta_i s^i,$$

(i.e., the true policy function is linear in parameters) and estimate the parameter vector $\Delta = (\Delta_0, \Delta_1, \dots, \Delta_I)$ semiparametrically by the generalized method of moments as follows.¹⁴

Step 1: Choose the largest set of quantile ranks that can be gleaned from the data, or

$\mathbf{r} = \{r_u\}_{u=1}^U$, where $r_u = \hat{F}_P^{-1}(p_u)$, $u = 1, \dots, U$, and \hat{F}_P is the Kaplan-Meier empirical distribution function for the set of prizes $P_{\mathcal{K},\mathcal{K}}$.

Step 2: For $I \leq U$, let $\tilde{S}(s) = \sum_{i=0}^{\infty} \Delta_i s^i$, and define

$$\hat{\Delta} = \arg \min \left\{ \sum_{u=1}^U \left[\hat{G}_{\mathcal{N}}^{-1}(r_u) - \tilde{S}\left(\hat{G}_{\mathcal{M}}^{-1}\left[\hat{F}_{P_{\mathcal{M}}}\left(\hat{F}_{P_{\mathcal{N}}}^{-1}[r_u]\right)\right]\right) \right]^2 \right\},$$

where $\hat{G}_{\mathcal{M}}$, $\hat{G}_{\mathcal{N}}$, $\hat{F}_{P_{\mathcal{M}}}$, and $\hat{F}_{P_{\mathcal{N}}}$ are the Kaplan-Meier empirical distributions of $\mathbf{S}_{\mathcal{M},T_{\mathcal{M}}}$, $\mathbf{S}_{\mathcal{N},T_{\mathcal{N}}}$, $\mathbf{P}_{\mathcal{M},\mathcal{M}}$, and $\mathbf{P}_{\mathcal{N},\mathcal{N}}$, respectively.

¹⁴The assumption that \tilde{S} is a polynomial of order I need not be a strong restriction on the empirical model. By Assumption 2.2, the Weierstrauss Approximation Theorem implies that the true policy function can be expressed as an infinite polynomial series $\tilde{S}(s) = \sum_{i=0}^{\infty} \Delta_i s^i$. Alternatively, one could choose a truncation point I to grow at a rate no faster than the data; this would eventually allow for recovery of the true, unrestricted \tilde{S} as the number of available sample moments grows.

Step 3: Using the standard errors from Step 2, test $H_0 : \Delta_0 = 0$. If H_0 is rejected, remove from \mathbf{r} any r_u such that $H_0^* : \Delta_0 \geq \widehat{G}_{\mathcal{N}}^{-1}(r_u)$ is rejected and repeat Step 2. ■

Step 3 in the above process comes from the fact that Step 2 is defined by equation (10), which is only valid for $s \geq \widetilde{S}(0)$.

There are two senses in which this estimator is semiparametric. First, there are no assumptions imposed on the form of the distributions of grades and prizes. Empirical CDF inverses can be recovered via “nearest neighbor” interpolation of the Kaplan-Meier empirical distributions. Second, the polynomial specification allows for the order of the policy function to be chosen as high as desired, given enough data. In that sense, the above proposal could be classified under the broad umbrella of estimation by the *method of sieves*.¹⁵ Said differently, this flexible form for \widetilde{S} allows for a simple estimation procedure within a finite parameter space, while virtually avoiding imposition of *a priori* restrictions on the behavior of the grade markup rule. If I is chosen to be large enough so that numerical stability is an issue, \widetilde{S} could alternatively be specified as a weighted sum of orthogonal basis polynomials, rather than the standard polynomial basis.

Another advantage is that minimization in Step 2 is greatly simplified by the polynomial specification of \widetilde{S} , since $\widehat{\Delta}$ can be found by simply regressing $\mathbf{Y} = \left(\widehat{G}_{\mathcal{N}}^{-1}(r_1), \dots, \widehat{G}_{\mathcal{N}}^{-1}(r_U) \right)^\top$ on the matrix of explanatory variables,

$$\mathbf{X} = \begin{pmatrix} 1 & \widehat{G}_{\mathcal{M}}^{-1} \left[\widehat{F}_{P_{\mathcal{M}}} \left(\widehat{F}_{P_{\mathcal{N}}}^{-1}(r_1) \right) \right] & \widehat{G}_{\mathcal{M}}^{-1} \left[\widehat{F}_{P_{\mathcal{M}}} \left(\widehat{F}_{P_{\mathcal{N}}}^{-1}(r_1) \right) \right]^2 & \dots & \widehat{G}_{\mathcal{M}}^{-1} \left[\widehat{F}_{P_{\mathcal{M}}} \left(\widehat{F}_{P_{\mathcal{N}}}^{-1}(r_1) \right) \right]^I \\ 1 & \widehat{G}_{\mathcal{M}}^{-1} \left[\widehat{F}_{P_{\mathcal{M}}} \left(\widehat{F}_{P_{\mathcal{N}}}^{-1}(r_2) \right) \right] & \widehat{G}_{\mathcal{M}}^{-1} \left[\widehat{F}_{P_{\mathcal{M}}} \left(\widehat{F}_{P_{\mathcal{N}}}^{-1}(r_2) \right) \right]^2 & \dots & \widehat{G}_{\mathcal{M}}^{-1} \left[\widehat{F}_{P_{\mathcal{M}}} \left(\widehat{F}_{P_{\mathcal{N}}}^{-1}(r_2) \right) \right]^I \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 1 & \widehat{G}_{\mathcal{M}}^{-1} \left[\widehat{F}_{P_{\mathcal{M}}} \left(\widehat{F}_{P_{\mathcal{N}}}^{-1}(r_U) \right) \right] & \widehat{G}_{\mathcal{M}}^{-1} \left[\widehat{F}_{P_{\mathcal{M}}} \left(\widehat{F}_{P_{\mathcal{N}}}^{-1}(r_U) \right) \right]^2 & \dots & \widehat{G}_{\mathcal{M}}^{-1} \left[\widehat{F}_{P_{\mathcal{M}}} \left(\widehat{F}_{P_{\mathcal{N}}}^{-1}(r_U) \right) \right]^I \end{pmatrix}.$$

This implies the familiar estimator $\widehat{\Delta} = (\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X}^\top \mathbf{Y}$, along with the familiar variance-covariance matrix for linear regression models.¹⁶ Using well-known results, it follows that the above GMM estimator is consistent, asymptotically normal, and converges at rate \sqrt{U} .

4.2. Estimating Pseudo-Private Costs. I now turn to the primary task of estimating the distributions over heterogeneity among competing students. Throughout this section I

¹⁵A sieve is a sequence of nested, finite-dimensional parameter spaces whose limit contains the true parameter space. For an in-depth discussion on estimation by the method of sieves, see Chen [3].

¹⁶For improved efficiency, one could incorporate an optimal weighting matrix \mathbf{W} into Step 2 and minimize

$$(\mathbf{Y} - \mathbf{X}\Delta)\mathbf{W}(\mathbf{Y} - \mathbf{X}\Delta)^\top$$

instead. Using the current data set it will become clear later that there is little to be gained in this case.

shall consider the case where $\tilde{S}(0) = \Delta_0 > 0$ since estimation in the opposite (simpler) case is similar, but with fewer caveats. Recall from Section 2 that in this case non-minority achievement is given by a piecewise differential equation. For minorities with equilibrium grades $s \in [0, \Delta_0]$, equilibrium achievement is characterized by differential equation (6). By monotonicity of the equilibrium, I have the following two identities,

$$G_{\mathcal{N}}(s) = 1 - F_{\mathcal{N}}(\psi_{\mathcal{N}}(s)), \text{ and}$$

$$g_{\mathcal{N}}(s) = -f_{\mathcal{N}}(\psi_{\mathcal{N}}(s)) \psi'_{\mathcal{N}}(s) = -f_{\mathcal{N}}(\theta) / \gamma'_{\mathcal{N}}(\theta).$$

Using this, I can re-write equation (6) to get the following

$$(12) \quad C'(s; \theta) = \frac{g_{\mathcal{N}}(s)}{f_P(F_P^{-1}[G_{\mathcal{N}}(s)])} = \bar{\xi}_{\mathcal{N}}(s).$$

For non-minorities submitting grades above Δ_0 , something similar can be done using differential equation (5). Recall that for a random variable S distributed according to $F(s)$, the distribution of $Z = \zeta(S)$ is simply $F(\zeta^{-1}(Z))$. Minority grades are distributed

$$S_{\mathcal{M}} \sim G_{\mathcal{M}}(s) = 1 - F_{\mathcal{M}}[\psi_{\mathcal{M}}(s)],$$

from which it follows that subsidized minority grades are distributed according to

$$\tilde{S}(S_{\mathcal{M}}) \sim \tilde{G}_{\mathcal{M}}(s) = G_{\mathcal{M}}[\tilde{S}^{-1}(s)] = 1 - F_{\mathcal{M}}(\psi_{\mathcal{M}}[\tilde{S}(s)]).$$

Note that $\tilde{G}_{\mathcal{M}}$ and its derivative show up in equation (5), along with $G_{\mathcal{N}}$ and its derivative. Therefore, the differential equation for non-minority achievement above grade level Δ_0 can be re-written as

$$(13) \quad C'(s; \theta) = \frac{(1 - \mu)g_{\mathcal{N}}(s) + \mu\tilde{g}_{\mathcal{M}}(s)}{f_P[F_P^{-1}((1 - \mu)G_{\mathcal{N}}(s) + \mu\tilde{G}_{\mathcal{M}}(s))]} = \tilde{\xi}_{\mathcal{N}}(s), \quad s \geq \Delta_1.$$

Similarly for minority achievement (conditional on positive output), (4) can be re-written as

$$(14) \quad C'(s; \theta) = \frac{(1 - \mu)\tilde{g}_{\mathcal{N}}(s) + \mu g_{\mathcal{M}}(s)}{f_P[F_P^{-1}((1 - \mu)\tilde{G}_{\mathcal{N}}(s) + \mu G_{\mathcal{M}}(s))]} = \tilde{\xi}_{\mathcal{M}}(s), \quad s \geq 0,$$

where $\tilde{G}_{\mathcal{N}}(s) = G_{\mathcal{N}}(\tilde{S}(s))$ is the distribution of *de-subsidized* non-minority test scores and $\tilde{g}_{\mathcal{N}}$ is its derivative. Equations (12), (13), and (14) provide a simple basis for an estimator of the private cost distributions, as they express a student's unobservable private cost type in terms of objects which are all observable to the econometrician. This will allow for recovery of sample of *pseudo-private costs* for each group, which in turn facilitate estimation of the underlying distributions.

The advantages of this method are two-fold. First, the resulting estimation procedure is computationally inexpensive, since equilibrium equations need not be repeatedly solved as in, say a maximum likelihood routine. Second, estimation requires no *a priori* assumptions on the form of the distributions F_M and F_N . However, there is one drawback: without parametric assumptions, it is impossible to identify private cost types for the potential mass point of minorities whose equilibrium achievement is zero. Under circumstances one might consider to be reasonable, this concern will only apply to a small portion of the sample, but it must be addressed. The policy function estimate and equations (8) and (12) can be used to recover the minority boundary condition

$$\theta^* = \inf \{ \theta : \gamma_M(\theta) = 0 \}$$

by computing the solution to

$$(15) \quad \mathcal{C}'(0; \theta^*) = \mathcal{C}'(\Delta_0; \theta_{\Delta_0}) \tilde{S}'(0),$$

where θ_{Δ_0} solves

$$\mathcal{C}'(\Delta_0; \theta_{\Delta_0}) = \bar{\zeta}_N(\Delta_0).$$

By comparing the resulting estimate of θ^* with the estimate of $\bar{\theta}$ recovered from equation (12) (where $s = 0$), if the interval $[\theta^*, \bar{\theta}]$ has a non-empty interior, then the empirical model implies a mass point, and minority private costs corresponding to a grade of zero are non-identified.

One way of dealing with the non-identification problem is to parameterize the upper tail of the distribution. If the upper tail is sparsely populated, a reasonable option would simply be to spread the mass of minorities uniformly over $[\theta^*, \bar{\theta}]$.¹⁷ With this modification, the equations above allow for recovery of a sample of *pseudo-private costs*

$$\hat{\Theta}_{N, T_N} = \{ \hat{\theta}_{N, t} \}_{t=1}^{T_N} \quad \text{and} \quad \hat{\Theta}_{M, T_M} = \{ \hat{\theta}_{M, t} \}_{t=1}^{T_M}$$

corresponding to each SAT score observation for minorities and non-minorities, respectively. From these, the underlying private cost distributions can be recovered, given some specification of the cost function \mathcal{C} . This leads to the next section.

4.3. Cost Function Estimation. Another advantage to the GPV method is that it provides for a partial specification test of the theoretical model. In any pure-strategy equilibrium, the theory requires that mappings (12), (13), and (14) must reflect a monotonic decreasing relation between private costs and academic achievement in order for the FOC to constitute an equilibrium. Given some specification of costs \mathcal{C} , if the data do not produce monotone decreasing mappings, the model is rejected on the grounds that the

¹⁷The specification error introduced by this parameterization can be assessed by comparing the results with alternative estimates obtained by mapping all zero-score observations for minorities onto either θ^* or $\bar{\theta}$.

data are not consistent with a monotone equilibrium in the specified model. To begin with, one might be inclined to consider a simple linear specification, say $\mathcal{C}(s; \theta) = \theta s$, as this would avoid introducing additional parameters into the model. However, this specification of costs leads to a non-monotone empirical mapping being recovered from (12), (13), and (14). As it turns out, there must be curvature in students' utility in order for the model to be consistent with the data.

I assume that achievement costs take the form

$$\mathcal{C}(s; \theta) = \theta e^{\alpha s}, \alpha > 0.$$

This choice is motivated by several factors, the most important being that it satisfies the regularity conditions required for existence of a monotonic, pure-strategy equilibrium (see Section 2.1). Aside from that, it has other attractive properties as well. Note that the cost of submitting a grade of zero is strictly positive. This corresponds to the notion that students must forego some minimum cost to graduate high school as a prerequisite for participation in the college admissions market. As it turns out, the above cost function allows for a tight fit between the empirical model and the data (at the optimal value of α).

With this specification of private costs, equations (12), (13), and (14) become

$$(16) \quad \theta = \frac{g_{\mathcal{N}}(s)}{f_P \left(F_P^{-1} [G_{\mathcal{N}}(s)] \right) \alpha e^{\alpha s}} = \frac{\bar{\xi}_{\mathcal{N}}(s)}{\alpha e^{\alpha s}}, \quad s \leq \Delta_1,$$

$$(17) \quad \theta = \frac{(1 - \mu)g_{\mathcal{N}}(s) + \mu \tilde{g}_{\mathcal{M}}(s)}{f_P \left[F_P^{-1} \left((1 - \mu)G_{\mathcal{N}}(s) + \mu \tilde{G}_{\mathcal{M}}(s) \right) \right] \alpha e^{\alpha s}} = \frac{\tilde{\xi}_{\mathcal{N}}(s)}{\alpha e^{\alpha s}}, \quad s > \Delta_1, \text{ and}$$

$$(18) \quad \theta = \frac{(1 - \mu)\tilde{g}_{\mathcal{N}}(s) + \mu g_{\mathcal{M}}(s)}{f_P \left[F_P^{-1} \left((1 - \mu)\tilde{G}_{\mathcal{N}}(s) + \mu G_{\mathcal{M}}(s) \right) \right] \alpha e^{\alpha s}} = \frac{\tilde{\xi}_{\mathcal{M}}(s)}{\alpha e^{\alpha s}}, \quad s \geq 0,$$

respectively. The zero surplus condition and equation (16) imply a relation between the curvature parameter α , and the value of the lowest prize, \underline{p} :

$$(19) \quad \begin{aligned} \mathcal{C}(0; \bar{\theta}) = \bar{\theta} &= \frac{\bar{\xi}_{\mathcal{N}}(0)}{\alpha} = \underline{p} \\ \Rightarrow \alpha &= \frac{\bar{\xi}_{\mathcal{N}}(0)}{\underline{p}}. \end{aligned}$$

As discussed in Section 2.5, the zero surplus condition is analogous to broader market forces (not included in the model) that determine participation in the higher-education market. If students have a choice between going to college and some outside option, then the marginal college candidate will be indifferent between going to college and

opting out. If prize values represent the additional utility from going to college over the outside option, then the result is (19). This condition places structure on the relative link between the utility of consumption and the disutility of work.

In work related to this, Guerre, Perrigne and Vuong [10] and Campo, Guerre, Perrigne and Vuong [2] extend the GPV method to first-price auctions where agents' utility functions display some form of curvature. They show that such models are not identified without imposing additional structure, due to the weak restrictions that the game-theoretic model places on observed bids. In fact, simply parameterizing either utility or the distribution of private information alone does not necessarily lead to model identification. Fortunately, the prize distribution in the college admissions model provides some additional structure, so here it is sufficient to parameterize just the cost function. Campo, *et al.* [2] use information on heterogeneity across auctioned objects to identify the utility function. This is conceptually similar to the role that the sample of prizes $\mathbf{P}_{\mathcal{K},K}$ plays, only instead of dealing with many (single-unit) auctions for heterogeneous items, I have a single "auction" with many heterogeneous objects.

Proposition 4.1. *If the cost function is restricted to the parametric class $\mathcal{C}(s; \theta) = \theta e^{\alpha s}$, $\alpha > 0$, then there exists a unique curvature parameter α and a unique set of cost distributions $F_{\mathcal{M}}$ and $F_{\mathcal{N}}$ which rationalize a given set of grade distributions, $G_{\mathcal{M}}$ and $G_{\mathcal{N}}$, a policy function \tilde{S} , and a prize distribution $F_{\mathcal{P}}$.*


Heuristic Proof: A formal proof that the model with exponential costs is identified from a sample of grades, prizes and a known policy function is under construction. However, some intuition can be gleaned by looking at what happens to the model for different values of α . Intuitively, the role of α and θ is to reconcile the prize utilities with the levels of observed achievement. At each grade quantile s , α must be such that the prize value allocated to a student with a score of s justifies the resulting cost $\mathcal{C}(s; \theta(s; \alpha))$, where $\theta(s; \alpha)$ is defined by the inverse equilibrium equations (16), (17), and (18).

The model's ability to reconcile the prize utilities with observed behavior hinges on the cost curvature parameter through the term

$$(20) \quad \alpha e^{\alpha s}$$

in the denominators of the three GPV equations. Specifically, suppose one were to fix a value of α , recover the associated GPV estimates of $F_{\mathcal{M}}$ and $F_{\mathcal{N}}$, and then compute the implied model-generated bid distributions. If α is too high, then the exponent of (20) becomes very important and the model cannot produce the high grades observed in the data because prize values are not enough to compensate for the cost of achievement. In other words, the marginal rate of substitution of prize value for work is too low to rationalize a fixed grade distribution from a fixed prize distribution. If α is very small,

FIGURE 3.



then the exponent is unimportant for low grades (because $e^{\alpha s}$ is close to 1), and the effect of the coefficient in (20) dominates. Costs become nearly linear for low α , and when this happens the behavioral separation in the model diminishes among low-performing students—in fact, mappings (16), (17), and (18) eventually become non-monotonic—and the observed low-score frequencies cannot be rationalized. However, in the middle there is a balance between the two extremes and the whole empirical grade distribution can be rationalized. Figure 3 provides an illustration. This concept motivates the proposed estimator below. ■

The estimator I propose for the utility function parameter is motivated by the fact that the model's ability to rationalize the empirical grade distributions \hat{G}_i , $i = \mathcal{M}, \mathcal{N}$ vanishes as α approaches the two limiting extremes of 0 and ∞ . For fixed α , the restricted GPV estimates of the cost distributions can be recovered from equations (16), (17), and (18). These and the equilibrium equations from Section 2 imply a set of model-generated grade distributions, \check{G}_i , $i = \mathcal{M}, \mathcal{N}$. The goal in choosing α , as with any parametric estimation routine, is to minimize the distance between the data and the model. This leads to the following nonlinear least squares (NLLS) estimator for the utility parameter:

$$(21) \quad \hat{\alpha} = \arg \min \left\{ \sum_{j=1}^J \left(\check{G}_{\mathcal{M}}(s_j; \alpha) - \hat{G}_{\mathcal{M}}(s_j) \right)^2 + \left(\check{G}_{\mathcal{N}}(s_j; \alpha) - \hat{G}_{\mathcal{N}}(s_j) \right)^2 \right\},$$

where $\mathbf{S} = \{s_1, s_2, \dots, s_J\}$ is the set of all grades observed in the data, $\ddot{G}_i(\cdot; \alpha)$ is the model-generated grade distribution for group i given α , and \hat{G}_i is the Kaplan-Meier empirical CDF.¹⁸

While this is an intuitive criterion function, optimization is complicated by the fact that the derivatives $d^n \ddot{G}_i / d\alpha^n$, $i = \mathcal{M}, \mathcal{N}$, $n = 1, 2, \dots$ of the model-implied grade distributions are not readily available due to a lack of closed-form solutions for the equilibrium equations in Section 2. The lack of closed-form solutions also necessitates repeated solution of the model equations during optimization for each guess of the cost curvature parameter. To address these problems, I use the *golden search method*, a derivative-free optimization algorithm.

Golden search begins with an initial guess on the search region, $[\underline{\alpha}, \bar{\alpha}]$ and evaluation of the objective function at two interior points $\alpha < \alpha'$. After comparing the functional values, the sub-optimal interior point is used to replace the nearest endpoint of the search region, and the process is repeated until the length of the search region collapses to a pre-specified tolerance, τ . The algorithm has some unique and attractive characteristics because the interior points are chosen as

$$\alpha = \varphi \underline{\alpha} + (1 - \varphi) \bar{\alpha}, \quad \text{and} \quad \alpha' = (1 - \varphi) \underline{\alpha} + \varphi \bar{\alpha},$$

where $\varphi = (\sqrt{5} - 1)/2$ is the inverse of the golden ratio, a number famously venerated by ancient Greek philosophers (hence, the name “golden search”). By choosing the interior points in this way, with each successive iteration one of the interior points is carried over from the previous iteration, necessitating only one new objective function evaluation. Moreover, at each step the length of the search region contracts by a factor of exactly φ ($\approx .62$), meaning that convergence obtains in a known number of steps equal to $(\log(\tau) - \log(\bar{\alpha} - \underline{\alpha})) / \log(\varphi)$.

Although the proposed semiparametric utility function estimator requires repeated computation of the model equilibrium, this last property of golden search gives the researcher an *a priori* idea of the magnitude of the problem. As for estimation of α with the current data set, it is easy to identify an appropriate search region of length less than 10, so convergence obtains in roughly 33-43 iterations for $\tau \in [10^{-8}, 10^{-6}]$. Of course, there are the usual problems of locating the global minimum as opposed to local minima, but this is not unique to derivative-free optimization methods. The final point left to discuss is the nonparametric density estimates that will be used in equations (16), (17), and (18). This is covered in the next section.

¹⁸An alternative criterion one could adopt is to choose $\hat{\alpha}$ so as to minimize the sup-norm distance between \ddot{G}_i and \hat{G}_i , for $i = \mathcal{M}, \mathcal{N}$.

4.4. Boundary-Corrected Kernel Density Estimation. Established asymptotic theory on GPV-type estimators is based on obtaining kernel-smoothed density estimates, which are known to exhibit excessive variance and bias near the extremes of the sample. GPV-type econometric routines typically address this issue by trimming elements from the sample of pseudo-private information based on kernel density estimates close to the extremes of the sample. However, addressing the problem in this way would cause problems here for several reasons. First, boundary conditions are needed for computation of the model equilibrium; second, the relation between α and \underline{p} is pinned down precisely at the boundary (see equation (19)); and third, the boundary of the minority grade distribution plays a role in estimating interior values of non-minority private costs (see equation (17)). Fortunately, there is a well-established set of tools from the statistics literature for improving the performance of kernel density estimators when the underlying random variables live on a bounded support.

Let f denote a density function with support $[a, b]$ and consider nonparametric estimation based on a random sample $\{Z_1, Z_2, \dots, Z_T\}$ using the standard kernel density estimator $\hat{f}(x) = \frac{1}{Th} \sum_{t=1}^T \kappa\left(\frac{x-Z_t}{h}\right)$, where κ is a unimodal density function and h is a bandwidth parameter strategically chosen to approach zero at a rate no faster than $\frac{1}{T}$. It is well-known that on the set $[a+h, b-h]$ this estimator has bias of order $O(h^2)$, but on the complement of this set, the bias is $O(h)$. In particular, the standard method tends to underestimate density values on the set $[a, b] \setminus [a+h, b-h]$ for an intuitive reason: since it cannot detect data outside the boundaries of the support, it penalizes the density estimate near those boundaries. This is commonly referred to as the *boundary effect*.

Various methods have been developed to address the problem.¹⁹ Two common coping techniques are known as the *reflection method* and the *transformation method*. The former is a simple technique in which the data are “reflected” outside the support near the boundaries, resulting in the following estimator: $\hat{f}(x) = \frac{1}{Th} \sum_{t=1}^T \left\{ \kappa\left(\frac{x-Z_t}{h}\right) + \kappa\left(\frac{x+Z_t}{h}\right) \right\}$. Transformation methods map the data onto an unbounded support via $\lambda : [a, b] \rightarrow \mathbb{R}$, resulting in $\hat{f}(x) = \frac{1}{Th} \sum_{t=1}^T \kappa\left(\frac{x-\lambda(Z_t)}{h}\right)$.

While these methods reduce the bias due to boundary effects, they come at a cost of increased variance in the density estimate. However, Karunamuni and Zhang [14, henceforth, KZ], overcome this problem by constructing a kernel estimator that is a hybrid of the reflection and transformation techniques. Formally, the boundary-corrected

¹⁹See Karunamuni and Alberts [13] for a more in-depth discussion of the various correction methods, as well as for a comparison of their performance.

KZ density estimator is given by

$$\begin{aligned}
\widehat{f}_T(x) &= \frac{1}{Th} \sum_{t=1}^T \left\{ \kappa \left(\frac{x - Z_t}{h} \right) + \kappa \left(\frac{x + \widehat{\lambda}(Z_t)}{h} \right) \right\}, \\
\widehat{\lambda}(y) &= y + \widehat{d}y^2 + A\widehat{d}^2y^3, \\
\widehat{d} &= \log(f_T(h_1)) - \log(f_T(0)), \\
f_T(h_1) &= f_T^*(h_1) + \frac{1}{T^2}, \\
f_T(0) &= \max \left\{ f_T^*(0), \frac{1}{T^2} \right\}, \\
f_T^*(h_1) &= \frac{1}{Th_1} \sum_{t=1}^T \kappa \left(\frac{h_1 - Z_t}{h_1} \right), \\
f_T^*(0) &= \frac{1}{Th_0} \sum_{t=1}^T \kappa_0 \left(\frac{-Z_t}{h_0} \right),
\end{aligned}
\tag{22}$$

where κ is a symmetric kernel with support $[-1, 1]$; $A > \frac{1}{3}$; $h_1 = o(h)$; $\kappa_0 : [-1, 0] \rightarrow \mathbb{R}$ is an optimal boundary kernel, given by $\kappa_0(y) = 6 + 18y + 12y^2$; and $h_0 = \beta h_1$, with

$$\beta = \left\{ \frac{\left(\int_{-1}^1 x^2 \kappa(x) dx \right)^2 \left(\int_{-1}^0 \kappa_0^2(x) dx \right)}{\left(\int_{-1}^0 x^2 \kappa_0(x) dx \right)^2 \left(\int_{-1}^1 \kappa^2(x) dx \right)} \right\}^{1/5}.$$

Interestingly, this estimator reduces to the standard kernel density estimator on the interior of the set $[a, b] \setminus [a + h, b - h]$. Most importantly, KZ show that if f is strictly positive and has a continuous second derivative within a neighborhood of the boundary, then \widehat{f}_T as defined above has $O(h^2)$ bias and $O(\frac{1}{Th})$ variance everywhere on the support.²⁰

The above boundary correction technique applies to the current empirical model of college admissions. A key assumption of the theory is that prizes and private cost types live on compact intervals, which in turn leads to bounded achievement. However, one can reasonably argue that these assumptions correspond to natural characteristics of the data. In the case of achievement, a student cannot put forth negative effort, so a grade of zero naturally forms a lower bound on the support of grades. By design, there is also

²⁰The assumption that the true density is strictly positive at the boundary is not necessary for boundary correction in general, just for the hybrid KZ estimator. If it is known *a priori* that the density attains a value of zero at the boundary, a suitable replacement with similar performance is the locally adaptive—meaning that the bandwidth is adjusted as domain points get closer to the boundary—kernel density estimator of Karunamuni and Alberts [13]. The cost associated with this alternative is that it is more difficult to implement.

a maximum attainable SAT score.²¹ As for the prize distribution, an argument similar to the logic behind Assumption 2.4 establishes bounds on the support. Once again, the set of realized prizes is assumed to be the result of a broader market equilibrium including entry and exit of firms supplying post-secondary education and unskilled jobs to high school graduates. Therefore, $[\underline{p}, \bar{p}] = [\min\{\mathbf{P}_{\mathcal{K},K}\}, \max\{\mathbf{P}_{\mathcal{K},K}\}]$, and the upper and lower bounds arise naturally from the exogenous private cost distributions (including high-cost types who opt out of higher education) and the interaction between employment suppliers and universities.

I can now summarize the structural estimator of the college admissions model in the following three-step process:

Step 1: Obtain the following preliminary estimates

- (i) the population demographic parameter $\hat{\mu} = \frac{\sum_{u=1}^U M_u}{\sum_{u=1}^U (M_u + N_u)}$;
- (ii) \tilde{S} as outlined in Section 4.1;
- (iii) the boundary-corrected KZ prize density $\hat{f}_P(p)$, and its integral, $\hat{F}_P(p)$ from the sample of prizes;
- (iv) the boundary-corrected KZ grade densities $\hat{g}_M(s|S > 0)$ and $\hat{g}_N(s)$, and the corresponding integrals, \hat{G}_M and \hat{G}_N from the samples of SAT scores.

- Step 2:** (i) For a given guess of α , estimate samples of pseudo-private costs $\hat{\Theta}_{N,T_N}$ and $\hat{\Theta}_{M,T_M}$ from equations (16), (17), (18), and (15), where the grade and prize distributions are substituted for the estimates from Step 1. In the event of a mass point at a score of zero for minorities, map minority scores of zero uniformly onto an evenly-spaced grid on $[\theta^*, \bar{\theta}]$, where the spacing between grid points is smaller than h , the bandwidth parameter for minority private costs, conditional on positive achievement.
- (ii) Given part (i) of Step 2, estimate the study-cost parameter $\hat{\alpha}$ via NLLS as outlined in Section 4.3. ■

Step 3: Obtain boundary-corrected KZ density and distribution estimates for private costs, \hat{f}_M and \hat{f}_N , using the samples of pseudo-private costs from Step 2.

4.5. Asymptotic Properties. In a related setting, Campo, Guerre, Perrigne, and Vuong [2, henceforth, CGPV] develop a similar semiparametric estimator of a first-price auction model where competitors' utility exhibits curvature. They parameterize bidder utility

²¹SAT scores are actually a proxy for overall academic achievement, so assuming that the maximal score forms a natural upper bound is an approximation to the truth. However, the data suggest it is a reasonable approximation: the number of students who manage a perfect SAT score make up less than three thousandths of a percent of the overall population.

and use variation in observable auction characteristics to estimate it via a NLLS routine. After that, they recover type distribution estimates similarly as in GPV. CGPV prove asymptotic normality and show that the utility curvature estimator converges at rate $K^{(R+1)/(2R+3)}$, where R is the number of continuous derivatives of the (true) type distributions. Type distribution estimates are shown to converge at the optimal rate for kernel-based estimators.

The estimators I have proposed for α , F_M and F_N are conceptually very similar to CGPV. Like them, I exploit variation in objects being auctioned to identify utility curvature, which I estimate via NLLS. Moreover, my type distribution estimates are conditioned on the utility curvature parameter, just as in CGPV. \hat{F}_M and \hat{F}_N are otherwise nonparametric and estimated via a two-step kernel smoothing procedure which involves analytically inverting the equilibrium equations from the theoretical model.

Henceforth, discussion pertaining to estimates and inference shall assume the asymptotic theory proven by CGPV. The standard error I report for $\hat{\alpha}$ shall reflect the conservative assumption that type distributions have a single continuous derivative, or $R = 1$. This implies that the rate of convergence is $K^{2/5}$.

As a precaution, I also perform a bootstrap exercise to evaluate the role of sampling variability for the estimates (see the appendix for details and diagrams). The histogram of bootstrapped estimates for α appears fairly normal, with variance slightly smaller than the estimate I get by assuming $R = 1$. Moreover, 95% confidence bands for the type distributions are fairly tight (see appendix), suggesting that the large sample size eliminates concerns about sampling variability. Effectively, estimation amounts to an exercise in curve fitting. This will simplify the discussion on the counterfactuals considerably, as one can reasonably focus on policy changes under the estimated distributions while ignoring inferential concerns.

4.6. Practical Issues. Choice of A is generally inconsequential, as long as $A > \frac{1}{3}$, so I have selected $A = .55$ as suggested by KZ. By definition of the boundary-corrected estimator, the Gaussian kernel is not an option, so I have chosen the biweight kernel (also known as the quartic kernel) $\kappa(x) = \frac{15}{16} (1 - x^2)^2 \mathbb{I}[-1 \leq x \leq 1]$, where \mathbb{I} is an indicator function. As proposed by KZ, I have selected bandwidth h via Silverman's [18] optimal global bandwidth rule

$$h = T^{-\frac{1}{5}} \times \left\{ \frac{\int_{-1}^1 \kappa^2(x) dx}{\left(\int_{-1}^1 x^2 \kappa(x) dx \right)^2 \int_{[0,a]} \left(\frac{d^2 f}{dx^2} \right)^2 dx} \right\}^{1/5},$$

where the second term in the denominator is substituted by

$$\int_{[0,a]} \left(\frac{d^2 f}{dx^2} \right)^2 dx \approx \frac{3}{8} \pi^{-\frac{1}{2}} \sigma^{-5},$$

and where σ is the sample standard deviation (see Silverman [18, equation 3.27]). Finally, there are many ways to choose $h_1 = o(h)$, but I use $h_1 = hT^{-\frac{1}{20}}$ as proposed by KZ.

In order to obtain estimates of the distributions, I numerically integrate the boundary-corrected, KZ densities via Simpson's rule. This method has the advantages of being both accurate and easy to implement. Moreover, I strategically choose the grid of points on which the densities are estimated so that the spacing is $\delta = \min\{h, .01\}$; this ensures that the resulting numerical error is of higher order than the statistical bias. The approximation error of Simpson's rule depends on the product of δ^5 and the fourth derivative of the actual integrand. Since the biweight kernel has a constant fourth derivative, the numerical error is actually $c\delta^5$, where c is fixed across domain points.

There are two final practical issues concerning numerical performance during estimation of $\hat{\alpha}$. I reconcile \underline{p} and α via an additive shift using equation (19), but before doing so, I treat $\hat{p} = \min_t \{Q_t\}$ as the numeraire good and divide all prize values by it. This has the effect of scaling up the length of the interval on which the optimal $\hat{\alpha}$ lives (roughly by a factor of 10), to allow for finer adjustments. In order to compute \tilde{G} at each golden search iteration, I solve the model equilibrium equations using a 4th-order Runge-Kutta algorithm. I also take measures to ensure a finer grid of domain points in regions of the function marked by a high degree of curvature. The maximal grid-point spacing for the Runge-Kutta integration is approximately .019, resulting in a numerical error on the order of 10^{-6} . Each iteration required 51 seconds on average, and convergence with a tolerance of 10^{-6} obtained in 28 iterations.

5. RESULTS AND COUNTERFACTUALS

5.1. Estimation Results. For the 1996 freshmen enrollment data, there were a total of 1,056,580 seats, with 186,507 going to minority students. This results in a demographic parameter estimate of $\hat{\mu} = .17652$, with a standard error of .000141. Table 2 displays summary statistics on normalized grades for each group. It also displays summary statistics for prizes awarded to each group. These figures are for USNWR quality indices prior to performing the affine transformations discussed in the previous section. See Figures 2 and 1 for graphical representations of the grade and prize distributions.

I selected an affine specification of the grade transformation function, or $\tilde{S}(s) = \Delta_0 + \Delta_1 s$. As it turns out, higher-order terms are unimportant, and this simple specification

TABLE 2. Summary Statistics for Normalized SAT Scores and Prizes

Sample	# of Obs	Median	Mean	StDev	Min	Max
Minority Grades	18,407	29	29.86	17.76	0	101
Non-Minority Grades	73,361	44	44.3	19.1	0	102
Minority Prizes (Raw USNWR Quality Index)	186,507	0.4875	0.4958	0.189	0.087	0.973
Non-Minority Prizes (Raw USNWR Quality Index)	870,073	0.5877	0.5792	0.1826	0.087	0.973

TABLE 3. AA Policy Function Regression Results

	Δ_0	Δ_1	R^2	Implied Avg. Grade Boost
	3.4218	1.0917	0.99789	6.1611
	(0.00277)	(0.00000199)		
95% Conf. Interval:	[3.3187, 3.5251]	[1.089, 1.0945]		

is enough to achieve a remarkably tight fit for the data.²² Table 3 summarizes the results of the policy function estimation. The regression R^2 value is 0.99789, with both the slope and the intercept being statistically significant at the 1% level. The high R^2 value is not necessarily surprising, given the nature of the sample: I have observations on virtually the entire universe of colleges, and the sample size for SAT scores constitutes a non-trivial fraction of the actual freshman population. More remarkable is the fact that such a tight fit is achieved with only an affine specification of the markup function.

The estimated policy function assesses a grade inflation factor of 9.17%, along with an additive boost of 34 points (in the original SAT score units). For example, a minority student with an SAT score of 1000 would see his grade increased to


$$\widehat{S}(1000) = 1.0917(1000) + 34 \approx 1126.$$

Combining these figures with the sample of normalized minority scores results in an estimated average grade boost of about 62 points in the original SAT score units.

Figure 4 graphically compares the estimated policy function with the data. The solid line is the regression line, and the dots are a scatter plot of the X_u s versus the Y_u s from Section 4.1. The dashed line is the 45°-line, where the policy would lie under color-blind admissions. The dispersion of the data-points around the regression line represents the

²²Higher order polynomial specifications produce coefficients that are statistically significant, but they do not improve the fit of the model in any practical sense. Moreover, the affine estimate and the polynomial estimates differ the most toward the upper extreme of the sample where the data are sparse.

FIGURE 4.
Grade Transformation Function:
Data versus Estimation



mis-specification error introduced by the assumption that individual college admissions boards can be treated as a single entity employing an affine admission preference. The data suggest that there was a remarkable degree of coordination on AA practices among different colleges and universities in 1996.

This view of admission preferences is consistent with previous empirical work on AA. Chung, Espenshade and Walling [5], estimate the average SAT-equivalent grade boost received by minority students at elite universities. They use individual-level data on applications and acceptance decisions at 3 undisclosed institutions from “the top tier of American higher education” to estimate the admission preference assessed to minority students. Chung, *et al.* fit a logistic regression model to the data in order to determine how different factors affect a student’s probability of being accepted. They find that minority students receive a substantial SAT-equivalent boost in admission decisions—230 points for African Americans and 185 points for Hispanics. While these figures are not directly comparable to my measure of the admission preference—Chung, *et al.* measure the added probability of being accepted at a particular college, whereas \tilde{S} measures the increase in school rank for the final placement outcome associated with race—I also find that race plays a significant role in how college seats are allocated. Moreover, Chung, *et al.* find that the admission preference is highest for minority applicants with high scores.

TABLE 4. Summary Statistics for Pseudo-Private Costs

Sample	Min	Max	Median	Mean	StDev
Minorities	.0005255	4.8754	0.5149	0.7452	0.6586
Non-Minorities	.0005255	4.8754	0.2241	0.3698	0.6251496

This is also consistent with my positive grade inflation estimate $\hat{\Delta}_1 = 1.0917$ which implies a larger bonus for higher scores. Among minority SAT scores in the top 5% (*i.e.*, a score ≥ 1200), the average grade boost is 98 points.

Finally, in related work Chung and Espenshade [4] find that the opportunity cost of admission preferences at selective institutions tends to be borne primarily by Asian students, who receive a significant SAT-equivalent penalty, whereas whites do not. The current study offers some explanation as to why. First, recall that Asian students are under-represented in every tier except the top. Moreover, the SAT data suggest that the distribution of Asian SAT scores has a higher mean and a fatter upper tail than that for Whites. Both score distributions are roughly normal, with the former being $N(1039, 213)$ and the latter being $N(1030, 183)$. Since the estimated policy function rewards high minority scores more, by extension it also penalizes high non-minority scores more (see equation 3). This is the reason why Asian applicants are negatively impacted the most: their score distribution has the fattest upper tail.

On the other hand, if one measures the opportunity cost of AA in terms of allocations of college seats, then it may actually *not* be the case that Asian students are most adversely affected. Inasmuch as the college admissions market is consistent with two key assumptions—namely, that (i) the market is effective at matching higher-performing students (of the same demographic class) with higher-quality schools, and (ii) the policy-maker does not attempt to rearrange the relative orderings of students within the same demographic group when devising an AA policy—then it will be the marginal non-minority students that are eliminated from elite institutions due to AA. For example, if a given admission preference produces a 10% reduction in non-minority enrollment within the top quartile of colleges, then only the lowest-scoring 10% within the top segment will be bumped down to schools in the next quartile. Since the best of the best non-minority students are disproportionately Asian—in fact, conditional on scores above the median, Asian SAT results stochastically dominate all other groups—the negative allocational effect of AA would tend to be born predominantly by other non-minorities. This is true despite the fact that Asians are typically assessed the highest effective penalty by admission preferences.

FIGURE 5. Empirical Equilibrium Mapping




FIGURE 6. Type Distributions and Densities




FIGURE 7. Goodness of Fit: SAT Score Distributions





FIGURE 8. Goodness of Fit: Prize Distributions



The cost curvature estimate is $\hat{\alpha} = 0.05758$, with a standard error of 0.001373, making $\hat{\alpha}$ significant at the 1% level. The empirical mappings under $\hat{\alpha}$ between types and achievement are displayed in Figure 5, where log types are on the abscissa and SAT scores on the ordinate. Monotonicity indicates that the empirical model is consistent with the data. Summary statistics for pseudo-types are displayed in Table 4, and the distributions and densities are displayed in Figure 6. The model suggests that types for minorities stochastically dominate non-minorities in the first-order sense.

Figure 7 illustrates the fit between the model and the data. The dashed step functions represent empirical grade distributions, and the solid lines represent model-generated analogs under the above parameter estimates. The right two panes display the percent error of the model from the data, as well as the absolute error. As another way of gauging goodness of fit, Figure 8 displays a similar plot for prize distributions, which were not directly targeted in the NLLS criterion function. These graphs provide suggestive evidence that the parameterizations introduced into the estimation scheme did not impose substantial mis-specification errors.

5.2. Counterfactual Policy Experiments. It is worth emphasizing that the standing assumption in the model is that the cost distributions F_M and F_N are invariant to policy changes. In that sense, the appropriate interpretation of Hickman [11] is a short-run model of policy implications. It is reasonable to assume that individual characteristics which determine academic competitiveness are fixed for children born prior to a policy change. One could certainly conceive of a broader model in which the Board designs a policy today so as to affect the private costs of future generations (*i.e.*, the children of today's college freshmen), but such an undertaking is beyond the scope of the current exercise, and is left for future research. Instead, I shall concentrate on the effects of the policy-maker's choices on actions and outcomes for individuals such as today's college candidates, whose private costs are reasonably viewed as fixed.

With the structural estimates in hand, I am now ready to address the main objective of assessing policy implications. In particular, I wish to compare the effects of three separate policies: the status-quo admission preference, a quota rule, and a color-blind admission scheme. The maintained assumption on the policy-maker is that he cares primarily about three objectives: (i) facilitating academic achievement, (ii) narrowing the racial achievement gap, and (iii) narrowing the enrollment gap. In terms of the objects associated with model equilibria, this means that he prefers (I) a population grade distribution over another if it first-order dominates; (II) a situation in which the separation between group-specific grade distributions is minimized; and similarly, (III) a minimal separation between distributions of prizes allocated to each group in

TABLE 5. %-Changes in Group Scores Relative to Status-Quo Policy

Quantile:		10 th	25 th	Median	75 th	90 th
\mathcal{M} Grades	Color-Blind:	+31.5	-1.32	-5.42	-1.83	+3.74
	Quota:	+64.45	+25.06	+11.67	+2.73	-2.29
\mathcal{N} Grades	Color-Blind:	+8.7	+3.09	+1.49	+0.59	+0.8
	Quota:	3.16	+2.48	+2.3	+2.58	+2.35

equilibrium. With these objectives in mind, I compute model equilibria and allocations under the three distinct policies, and I display the results below.

Figures 9 and 10 graphically represent the results of the counterfactual experiments. Dashed lines denote distributions associated with the status-quo policy, solid lines denote distributions arising from a quota, and dash-dot lines denote a color-blind outcome. Between two lines with the same style, the one lying to the left pertains to minorities. When comparing two grade distributions, keep in mind that if distribution i lies to the right of distribution j in some region, that indicates an interval of students who achieve higher SAT scores under policy i . Table 5 displays percentage-changes in achievement under the two unobserved policies, relative to the status quo. The changes are measured at various quantiles, including the upper and lower deciles, the intermediate quartiles, and the median.

This information produces some intriguing insights into AA. Figure 9 and the first line of Table 5 characterize the effect of an admission preference on academic output. Relative to color-blind admissions, both the highest- and lowest-performing minority students decrease their effort, whereas students in the middle increase it. Although the policy-maker might hope that students will use a grade bonus solely to bolster their competitive edge, in some situations a rational student will react by treating the bonus as a direct utility transfer. In the case of high-performing students, the bonus is not needed and they achieve lower grades.

For low performers, the fixed grade boost Δ_0 improves their standing (the inflation factor is insignificant for scores close to zero), but in so doing, it adversely alters the marginal costs and benefits of achievement. In order to improve their payoff beyond what the grade subsidy achieves, they would have to compete with students whose costs are significantly lower than theirs. Thus, the marginal cost of competing is too high relative to the potential benefits. It is only for students whose costs are low enough—but not too low—that the admission preference entices additional investment in effort.

TABLE 6. %-Changes in Enrollment, Relative to Status-Quo Policy

Tier:		I	II	III	IV
Minorities	Color-Blind:	-35.06	-21.25	+3.71	+41.73
	Quota:	+52.69	+16.17	-14.71	-42.88
Non-Minorities	Color-Blind:	+4.59	+3.75	-1.02	-18.53
	Quota:	-6.88	+2.95	+3.79	+19.37

As for non-minorities, an admission preference creates discouragement effects which discourage achievement among all types of students, relative to the color-blind case.

Figure 9 also illustrates some interesting insights on the effects of a quota. It increases output among low-performing minorities, relative to a color-blind rule, and decreases effort among high-performers. The intuition is simple. A high-cost minority competing against the population at large is subject to a substantial discouragement effect since there is a large amount of competitors with lower costs. On the other hand, if he competes only against his own group (as with a quota) where costs are on average higher, then it is more worthwhile to invest in costly effort, since his relative standing with regards to the competition is improved. For low-cost minorities, the opposite effect occurs: when they only compete against other minorities, there is less need to outperform the competition as aggressively as before. For non-minorities with high- and low-cost types, the reverse effect applies (low-performers back off effort, high-performers increase it) by similar reasoning.

In January of 2008, presidential candidate Barack Obama famously stated in a television interview that his daughters should not be treated as disadvantaged in college admissions decisions, and that poor white children should be given extra consideration. The current empirical model seems to support the intuition behind Mr. Obama's assertion. It is interesting to note that both types of Affirmative Action are detrimental to effort incentives for low-cost minorities *and* high-cost non-minorities.

With conflicting changes in academic output for different segments of the population, one might ask how the overall population grade distribution is effected. Table 5 answers this question: population grade distributions under each policy can be ordered by stochastic dominance. Color-blind admissions dominate the status quo, and a quota dominates the color-blind policy.

The model also shows that race-conscious admissions have a significant impact on college placement outcomes for minority students. Figure 10 displays the distributions of

TABLE 7. %-Changes Relative to Status-Quo Policy


	Quantile:	10 th	25 th	Median	75 th	90 th
Population Grades (Objective I)	Color-Blind:	+6.01*	+1.99*	+1.17*	+0.92*	+0.92*
	Quota:	+10.63**	+4.96**	+2.51**	+2.17**	+2.11**
Achievement Gaps (Objective II)	Color-Blind:	-2.95*	+7.82	+13.53	+6.91	-9.74**
	Quota:	-28.11**	-21.81**	-14.02**	+2.2*	+18.98
Enrollment Gaps (Objective III)	Color-Blind:	+52.19	+61.9	+76.62	+100.13	+117.76
	Quota:	-100**	-100**	-100**	-100**	-100**

prizes allocated to each group under each policy. Note the substantial first-order dominance shift that occurs under either AA policy, relative to color-blind admissions. Table 6 numerically displays the percentage changes in enrollment for each college tier. By shutting down American AA (as in color-blind admissions) minority enrollment within the top quartile would decrease by roughly a third, and within the upper middle quartile it would decrease by just over a fifth. Another striking feature of the table is that the majority of the displaced minority enrollment resulting from elimination of AA would end up in the lowest tier. The cost imposed on non-minorities amounts to roughly 5% and 4% of enrollment in each of the top two quartiles, respectively. Of course, on an individual level the cost and benefit to each group exactly balance out: any quality units reallocated to one student are necessarily transferred from another. Whether such transfers are justified is beyond the scope of economic reasoning.

However, it does appear that the AA policies implemented in real-world settings are effective at improving market outcomes for minorities, as intended by policy-makers. On the other hand, they do not eliminate the enrollment gap completely. For example, under a quota minority enrollment in the top tier would increase by an additional 52%. Loosely speaking, the US admission preference eliminates roughly $\frac{2}{3}$ of the enrollment gap in the top tier.

For a comparison of admissions rules along each of the policy objectives, I turn to Table 6, which tracks changes along objectives I – III. Once again, all figures are stated in terms of percentage changes, relative to the status-quo policy. For example, switching to a color-blind policy would increase the gap between prizes awarded to the median student from each group by 80%. In the table, asterisks are used to denote the preference ranking among the three policies. Two asterisks denote the most preferred outcome, one asterisk denotes the second-most preferred, and no asterisks indicates the least preferred. Interestingly, the status-quo admission preference never achieves the best outcome in any

FIGURE 9. Counterfactual: Group SAT Score Distributions




category. These figures also demonstrate that no ranking between a color-blind rule and the status quo can be established without knowing how the policy-maker's preferences weight objectives *I – III*. The former does strictly better in terms of academic performance, as mentioned above. The latter does strictly better in terms of enrollment gaps (see also Figure 10), and in terms of the achievement gap, the result is a toss-up.

On the other hand, a striking feature of the table is that a quota rule does strictly better than both other policies in *nearly every* category. Not only does it induce the highest academic output from the overall population of competitors, but it also shuts down the enrollment gap completely, by design. A quota also produces a substantial increase in minority achievement, as well as a narrowing of the achievement gap among the majority of the population. The lone drawback to a quota rule is that it produces the widest achievement gap in the upper tail of the grade distribution. However, one can argue that a quota rule appears to be the clear winner among college admission policies for reasonable social choice functions that do not place an extreme amount of weight on minimizing the achievement gap specifically in the upper tail of the grade distribution.

5.3. Alternative Policy Proposal. The counterfactual exercise produced some valuable insights into the costs and benefits of AA. However, the value in knowing that a quota is a substantially superior policy choice would seem to be diminished by the fact that quotas

FIGURE 10. Counterfactual: Prize Allocation Distributions



are illegal in the US, due to a 1978 Supreme Court ruling.²³ One might then ask whether an admission preference system can be modified so as to improve its performance, but without requiring an unreasonable level of information on the part of the policy-maker. As it turns out, the insights gained from the properties of a quota mechanism can be used to design a simple admission preference that performs similarly along the three policy objectives.

Formally, this alternative policy is defined by

$$\tilde{S}^*(s) \equiv G_{\mathcal{N}}^{-1} [G_{\mathcal{M}}(s)].$$

In words, the college admissions board simply announces beforehand that it will map quantiles of the minority grade distribution into the corresponding quantiles of the non-minority grade distribution. For example, the median minority score is reassigned a value equal to the median non-minority score, and so on. The fact that this mechanism is outcome-equivalent to a quota immediately follows from plugging \tilde{S}^* or $(\tilde{S}^*)^{-1}$ into equations (3), which then become equations (2).

Aside from its superior performance, this alternative admission preference has two other advantages worth mentioning. The first is its relative simplicity. In keeping with the *Wilson doctrine*, it does not require the policy-maker to have knowledge of students'

²³See *University of California v. Bakke* (438 U.S. 265 1978).

individual abilities, or their beliefs about the competition they face. Rather, \tilde{S}^* allows the policy-maker to implement an improved outcome using only information on grades and race. The second advantage is that this mechanism is a self-adjusting grade markup rule: the bonus it assesses to minority test scores is proportional to the amount of asymmetry between demographic groups. In fact, if the competition is symmetric, \tilde{S}^* is also equivalent to a color-blind mechanism. This concept is formalized in the following Theorem.

Theorem 5.1. *For a sequence of cost distributions $\{F_{\mathcal{M},k}, F_{\mathcal{N},k}\}_{k=1}^{\infty} \rightarrow (F_{\Theta}, F_{\Theta})$, let \tilde{S}_k^* be defined by*

$$\tilde{S}_k^* \equiv G_{\mathcal{N},k}^{-1}[G_{\mathcal{M},k}(s)],$$

where $G_{i,k}$, $i = \mathcal{M}, \mathcal{N}$ are the equilibrium grade distributions. Then it follows that the induced sequence $\{\tilde{S}_k^*\}$ converges to a color-blind rule, or $\tilde{S}^*(s) = s$.

Proof: As shown in Hickman [11], for each k , achievement under the mechanism defined by \tilde{S}^* is given by the following differential equation:

$$(23) \quad (\gamma_{i,k}^*)'(\theta) = -\frac{f_{i,k}(\theta)}{f_P \left(F_P^{-1} (1 - F_{i,k}(\theta)) \right) \mathcal{C}'(\gamma_{i,k}^*(\theta); \theta)}, \quad i = \mathcal{M}, \mathcal{N},$$

with a boundary condition given by the zero surplus condition. Moreover, achievement under a color-blind mechanism is characterized by


$$(24) \quad (\gamma_{\mathcal{M},k}^{cb})'(\theta) = (\gamma_{\mathcal{N},k}^{cb})'(\theta) = (\gamma_k^{cb})'(\theta) = -\frac{\mu f_{\mathcal{M},k}(\theta) + (1 - \mu) f_{\mathcal{N},k}(\theta)}{f_P \left(F_P^{-1} [1 - \mu F_{\mathcal{M},k}(\theta) - (1 - \mu) F_{\mathcal{N},k}(\theta)] \right) \mathcal{C}'[\gamma_k^{cb}(\theta); \theta]},$$

with the same boundary condition. Note that as $\{F_{\mathcal{M},k}, F_{\mathcal{N},k}\}_{k=1}^{\infty} \rightarrow (F_{\Theta}, F_{\Theta})$, the right-hand sides of equations (23) and (24) above both converge to

$$(25) \quad (\gamma_{\mathcal{M}}^*)'(\theta) = (\gamma_{\mathcal{N}}^*)'(\theta) = (\gamma^*)'(\theta) = -\frac{f_{\Theta}(\theta)}{f_P \left(F_P^{-1} (1 - F_{\Theta}(\theta)) \right) \mathcal{C}'(\gamma^*(\theta); \theta)}, \quad i = \mathcal{M}, \mathcal{N}.$$

Given this fact, we have $\{G_{\mathcal{M},k}, G_{\mathcal{N},k}\}_{k=1}^{\infty} \rightarrow (G, G)$, from which the result follows. ■

Figure 11 plots a comparison of the status quo AA policy function with \tilde{S}^* (as generated by equilibrium grade distributions) and a color-blind policy under the 1996 cost distribution estimates. Several interesting observations arise from the plot. First, \tilde{S}^* overcomes the incentive problem at the lower end of the achievement distribution by closely resembling a color-blind rule for students whose academic output is low. Second, \tilde{S}^* incentivizes higher test scores for low and mid-range students (recall Figure 9) by awarding them an increasing marginal grade markup for low and mid-range scores. Third, the



marginal grade markup eventually decreases as achievement increases (roughly around a grade of 60), corresponding to the notion that lower cost types need less help. Finally, the lone drawback of a quota rule—recall that it results in the lowest minority achievement and the widest achievement gap above the 90th percentile (see Table 6 and Figure 9)—arises from the fact that it gives too much assistance to high-performing students; in fact, it awards a larger grade boost than the status quo. This comes as a result of the information constraints that the policy-maker faces. Once he announces the policy \tilde{S}^* , agents' behavior partially determines the shape of the grade transformation, making it impossible to improve incentives for all students, without observing their private information.

6. CONCLUSION

This work has provided some useful empirical insights into the costs and benefits of Affirmative Action in college admissions. I have documented that significant gaps exist among different races in terms of academic performance and academic outcomes. I have also demonstrated that a policy-maker's choice of what admission rule to implement can have a large impact on both performance and outcomes. Some policies are difficult to compare, while others emerge as being superior in terms of a set of general policy objectives. In particular, a quota rule promotes higher academic performance, and gives rise to a narrower achievement gap than an admission preference or a color-blind policy. By construction, it also shuts down the enrollment gap completely.

Future progress along this line of research can be achieved by studying a dynamic version of the model to explore the implications of college admissions policies in a setting where the policy-maker attempts to affect the long-run evolution of private-cost

distributions. This will help to uncover how/whether AA helps or hinders the ultimate objective of erasing the residual effects of past institutionalized racism. The insights developed here will hopefully serve as a basis for answering these important questions in the future.

APPENDIX

6.1. USNWR Data and Methodology. Table 8 contains descriptions and descriptive statistics of the quality measures used to compute the USNWR quality index. Column 1 contains variable descriptions and column 2 displays the weights placed on each category (within-category weights are uniform). Columns 3-5 display descriptive statistics. Column 5 displays total sample size for each variable. In cases where USNWR lacks a certain datum for some school, it replaces the datum with the lowest value observed for schools within the same region and Carnegie classification. Columns 3 and 4 display means and sample standard deviations for the schools where the variable value is observed. One final note is also worth mentioning: in computing the quality index, USNWR maps average SAT and ACT scores into the corresponding cumulative distribution values within the SAT and ACT score distributions. This allows for comparisons of scores on different tests. The mean and standard deviation for average test scores in the table reflect this transformation.

6.2. Zero Achievement Cutoff. Recall that the working interpretation of a student with zero academic achievement is one who simply engages in random responding to test questions. In order to uncover the distributions over random outcomes, I simulated 100,000 random responses to a published practice test for the ACT—another standardized test widely used in US college admissions. The results are plotted in Figure 12. The upper pane is the unconditional distribution of simulated random responses. The mean of the distribution is 12.1224, with a standard deviation of .9224.

The question of whether 12 or 13 is the appropriate zero-achievement cutoff is addressed in the lower two panes. On the left is a comparison of the simulated distribution and the distribution from the data, conditional on a score of 12 or less; the right pane is the same for a cutoff of 13. On the right side the two distributions are close, with a single crossing at a score of 11; the data distribution stochastically dominates on the left, and the distributions are not as close. Using these insights, I interpret an ACT score of 12 as corresponding to zero academic achievement, or in other words, $S = 0 \Leftrightarrow \text{ACT score} = 12$.


I use score concordance tables to determine the equivalent zero achievement cutoff on the SAT test. Score concordances are jointly computed by the designers of the ACT and SAT using data on students who took both tests. The result is an interval of SAT scores

TABLE 8. USNWR Quality Indicators (Total of 1,314 Schools):

Category/Variable Description	Weight	Mean	StDev.	Observations
SELECTIVITY	15%			
Acceptance Rate		.7597	.1553	1,226
Yield (% accepted students who enroll)		.4428	.1518	1,226
Avg. SAT/ACT Scores of Enrolled Students		.5515	.2101	1,152
% First-Time Freshmen in Top HS Quartile		.5227	.2038	1,008
FACULTY RESOURCES	20%			
% Full-Time Instructional Faculty w/Terminal Degree		.7622	.1665	1,221
% Full-Time Instructional Faculty		.6505	.1891	1,231
Avg. Faculty Compensation		\$52,409.23	\$12,982.11	1,291
Student/Faculty Ratio		14.99	4.2	1,245
FINANCIAL RESOURCES	10%			
Education Spending/Student		\$9,494.56	\$5,283.01	1,193
Non-Education Spending/Student		\$5,951.12	\$8,321	1,292
RETENTION	25%			
Avg. Graduation Rate		.5353	.6581	1,154
Freshman Retention Rate		.7396	.1146	1,224
ALUMNI SATISFACTION	5%			
Alumni Giving Rate		.2105	.1237	1,165
ACADEMIC REPUTATION	25%			
College Administrator Ranking Poll		N/A	N/A	N/A

being mapped into each outcome-comparable ACT score (since SAT scores occur on a finer grid). These indicate typical outcomes one can expect on the SAT for a student with a given score on the ACT, and vice versa. The SAT-equivalent range for an ACT score of 12 is 520-580

The alert reader may wonder why the random responding exercise was not performed using an SAT practice test instead. As it turns out, the mean score from random responding on the SAT is 450, significantly lower than the 520-580 range predicted by the concordance study. Moreover, conditional score distributions for the SAT do not render a similar fit as in the lower left pane of Figure 12: random responding and actual data distributions conditional on low scores differ significantly in shape. However, this is not surprising considering that the SAT is designed to test one's academic aptitude (*i.e.*, ability for abstract reasoning), whereas the ACT is designed to test one's





achievement (*i.e.*, acquisition of knowledge). Although study effort undoubtedly plays a major role in determining scores on both tests, the distinction between achievement versus aptitude becomes more pronounced near the lower extreme. As the concordance study suggests, individuals who choose to acquire low levels of knowledge—*i.e.*, individuals with scores statistically indistinguishable from random responding on the ACT—typically have aptitudes that allow them to beat random responding on the SAT. For more information on the distinction between the ACT and SAT tests, see <http://www.act.org/aap/concordance/understand.html>.

6.3. Bootstrapped Standard Errors. For this exercise, I resampled the data one thousand times and computed $\hat{\alpha}$, \hat{F}_M , and \hat{F}_N each time. Figure 13 displays 95% confidence bands for the distribution and density estimates. Figure 14 displays a histogram of cost curvature estimates. The bootstrapped mean and standard deviation of $\hat{\alpha}$ are $\mu_{\hat{\alpha}} = 0.05622$ and $\sigma_{\hat{\alpha}} = 0.00104$, respectively. A kernel-smoothed density as well as a $N(\mu_{\hat{\alpha}}, \sigma_{\hat{\alpha}})$ density (scaled by histogram bin width) has been superimposed on the histogram for comparison.

It should be mentioned here that the current estimator does not strictly satisfy all of the regularity conditions for the nonparametric bootstrap, since the support of the type distributions is *ex-ante* unknown. However, this is only intended to be a precautionary measure to compliment the asymptotic theory proved in CGPV [2]. From the results displayed in the two figures, two things appear evident: first, it seems plausible that $\hat{\alpha}$ is asymptotically normal, and second, sampling variability seems to play no significant role.

FIGURE 13. Bootstrapped Confidence Bands

FIGURE 14. Histogram of Bootstrapped $\hat{\alpha}$ Estimates

REFERENCES

- [1] William G. Bowen and Derek Bok. *The Shape of the River: Long-Term Consequences of Considering Race in College and University Admissions*. Princeton, NJ: Princeton University Press, 1998.
- [2] Sandra Campo, Immanuel Guerre, Isabelle Perrigne, and Quang Vuong. Semiparametric estimation of first-price auctions with risk averse bidders. *Typescript, The Pennsylvania State University Department of Economics*, 2009.

- [3] Xiaohong Chen. Large sample sieve estimation of semi-nonparametric models. *Typescript, New York University Department of Economics*, 2005.
- [4] Chang Y. Chung and Thomas J. Espenshade. The opportunity cost of admission preferences at elite universities. *Social Science Quarterly*, 86(2):293–305, 2005.
- [5] Chang Y. Chung, Thomas J. Espenshade, and Joan L. Walling. Admission preferences for minority students, athletes, and legacies at elite universities. *Social Science Quarterly*, 85(5):1422–1446, 2004.
- [6] Roland G. Fryer, Jr. and Steven D. Levitt. Understanding the black-white test score gap in the first two years of school. *The Review of Economics and Statistics*, 86(2):447–464, 2004.
- [7] Roland G. Fryer, Jr. and Glenn C. Loury. Affirmative action and its mythology. *Journal of Economic Perspectives*, 19(3):147–162, 2005.
- [8] Quiang Fu. A theory of affirmative action in college admissions. *Economic Inquiry*, 44(3):420–428, 2006.
- [9] Immanuel Guerre, Isabelle Perrigne, and Quang Vuong. Optimal nonparametric estimation of first-price auctions. *Econometrica*, 68(3):525–574, 2000.
- [10] Immanuel Guerre, Isabelle Perrigne, and Quang Vuong. Nonparametric identification of risk aversion in first-price auctions under exclusion restrictions. *Econometrica*, 77(4):1193–1227, 2009.
- [11] Brent R. Hickman. Effort, race gaps, and affirmative action: A game-theoretic analysis of college admissions. *Typescript, University of Iowa Department of Economics*, 2010.
- [12] Christopher Jencks and Meredith Phillips, editors. *The Black-White Test Score Gap*. Washington D.C.: The Brookings Institution, 1998.
- [13] Rhoana J. Karunamuni and Tom Alberts. On boundary correction in kernel density estimation. *Statistical Methodology*, 2:191–212, 2005.
- [14] Rhoana J. Karunamuni and Shunpu Zhang. Some improvements on a boundary corrected kernel density estimator. *Statistics & Probability Letters*, 78:499–507, 2008.
- [15] Robert J. Morse, editor. *America’s Best Colleges*. Washington D.C.: US News & World Report, 1996.
- [16] Derek A. Neal and William R. Johnson. The role of premarket factors in black-white wage differences. *Journal of Political Economy*, 104(5):869–895, 1996.
- [17] Harry J. Paarsch. Deciding between the common and private value paradigms in empirical models of auctions. *Journal of Econometrics*, 51:191–215, 1992.
- [18] Bernard W. Silverman. *Density Estimation for Statistics and Data Analysis*. London: Chapman and Hall, 1996.
- [19] Thomas Sowell. Affirmative action around the world. 2004.

DEPARTMENT OF ECONOMICS, UNIVERSITY OF CHICAGO, 1126 E. 59th ST., CHICAGO, IL, 60637, USA
E-mail address: HickmanBR@UChicago.Edu