

IDENTIFYING THE SOURCES OF RETURNS TO COLLEGE EDUCATION USING AFFIRMATIVE ACTION POLICY

AARON L. BODOH-CREED[‡] AND BRENT R. HICKMAN[†]

ABSTRACT. We estimate a model of a college admissions contest with affirmative action (AA) where students compete for seats at better schools by choosing pre-college human capital (HC) investments. We are able to identify and flexibly estimate the contest structure in the college admissions market, including the tendency for AA to affect admissions profiles on different segments of the US college quality spectrum. We identify the effects of college quality, pre-college HC, and unobserved student characteristics on post-college income using a control function derived using methods from the auctions literature. College quality is the primary income determinant, while unobserved student characteristics play a secondary role. Pre-college HC affects income indirectly through its influence on graduation probability and enrollment. We also estimate the fraction of pre-college HC “wasted” by the rat-race nature of the admissions contest, and run counterfactuals of admissions, graduation rates, and post-college income race gaps under alternate admissions schemes not observed in the data.

Date: Original version: July 2017; Current version: September 2019.

Key words and phrases. Returns to College Quality; Affirmative Action; Racial Achievement Gap; Inequality

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

Special thanks are due to Srihari Govindan, Harry J. Paarsch, Elena Pastorino, Derek Neal, Limor Golan, Ariel Pakes, Isabelle Perrigne, Quang Vuong, Timothy P. Hubbard, and John Eric Humphreys. We are also grateful for helpful inputs from (in no particular order) Philippe Fevrier, Xavier D’Haultfoeuille, Parag Pathak, Steven Durlauf, Jimmy Roberts, Glen Weyl, Alessandra Voena, B. Ravikumar, Guillame Vandenbroucke, and Gustavo Ventura, as well as conference/seminar participants at U of Iowa, Johns Hopkins, U of Chicago, U of Georgia, the Australian National U, U of Wisconsin–Madison, the Milton Friedman Institute, the Stanford Institute for Theoretical Economics, Yale, Brigham Young U, U of Illinois at Urbana-Champaign, Penn State, Northwestern, Berkeley, the Society for Economic Dynamics, Copenhagen Business School, Washington U in St. Louis, Ohio State, Cornell, Harvard, Rice and Sciences Po. Any errors are our own. We thank the following individuals for helping us acquire the data for this project: Andrew Mary of the Integrated Postsecondary Education Data System of the US DOE and Kevin Brown of NORC. Brian Christensen and Joanna Harris provided crucial research assistance in organizing the college-level data.

[‡]HAAS SCHOOL OF BUSINESS, UNIVERSITY OF CALIFORNIA, BERKELEY, acreed@berkeley.edu

[†] OLIN BUSINESS SCHOOL, WASHINGTON UNIVERSITY in ST. LOUIS, hickmanbr@gmail.com;

1. INTRODUCTION

For many Americans, the competition to be admitted to a high quality college is one of the highest stakes contests of their lives. The reason is simple: there is a high degree of college heterogeneity in terms of educational inputs and the outcomes realized by students. For example, in 1988 the interquartile range of spending per student is \$5,931 to \$9,551, while the interquartile ranges of the graduation rates and household income (10 years after graduation) are 38% to 63% and \$54,000 to \$113,000.¹ Separating the causes of these disparate outcomes requires disentangling the influence of college quality, pre-college human capital (HC) investment, and underlying student characteristics on the returns to attending college.

Previous empirical work has attempted to address the issue by instrumenting for the influence of college quality while subsuming unobserved student characteristics into the unexplained error term in the model (see Brewer, Eide, and Ehrenberg [17]; Dale and Krueger [27]; Black and Smith [12, 13]; Arcidiacono [6]; and Long [46]). Moreover, this literature does not account for how affirmative action (AA) shapes HC accumulation by inducing incentive variation into the competition for college seats.² Our goals are to (1) separately measure the roles of college quality, pre-college invested HC, and permanent student heterogeneity in driving post-college income; (2) explore the welfare implications of using high-school achievement as a rank-order index for allocating heterogeneous college seats, thereby introducing a “rat-race” element into HC decisions; and (3) explore the effects of changes to the AA system on student outcomes (e.g., endogenous student HC investment). We estimate a model wherein college applicants with differing unobserved learning costs compete for admission to college through the accrual of HC spanning years of the student’s life prior to the college application process. Each college occupies a different point on the quality spectrum and uses a distinct, endogenous admissions cutoff. AA programs cause these cutoffs to be different for different demographic groups. A structural approach is necessary to identify the students’ private information about characteristics unobserved to the econometrician and (potentially) to colleges.

Within our model, there are two incentives to invest in HC. First, HC is a productive asset, which creates a *productive channel* of incentives. The strength of this channel (i.e., the marginal return to HC) may be a function of the quality of the student’s college or the student’s unobserved type.³ This channel is present in the complete-information assortative matching model of Becker [10]. Second, since colleges wish to admit the best students possible, they rank students using the students’ academic achievement (i.e., HC). Therefore, HC investment plays a second role as

¹The graduation rate data comes from Integrated Postsecondary Education Data System, while the income data comes from the Baccalaureate and Beyond (B&B) survey. Income quartiles have been rounded to the nearest thousand dollars to protect B&B respondents’ anonymity.

²Throughout this paper, when we use the term AA we refer to the practice of granting preferential treatment in college admissions on the basis of race, holding academic merit (grades and exam scores) fixed.

³Throughout the paper, we frequently refer to a student’s “unobserved type” as an essential exogenous characteristic that drives the student’s behavior. Intuitively, this can be thought of as a student-specific fixed effect which subsumes a host of life-shaping factors—e.g., parents’ income and education, home environment, etc.—that influence the ease with which the student acquires new HC throughout childhood and early adulthood.

a rank-order index that allows better students to out-compete their peers and enroll into better schools, and these indirect returns create a *competitive channel* of incentives. The competitive channel creates a strategic interaction between one’s own actions and the actions of others: if rivals never study mathematics and science, one might study less, consume more leisure time, and still place into a top college. The strategic interaction results in over-production of HC in response to competitive pressure. By separating the two channels of incentives, we can also empirically measure the degree to which pre-college HC is overproduced due to the competitive channel of incentives.⁴

The first goal of our analysis is to empirically tease apart the effects of college quality, HC, and unobserved type on post-college household income. We use individual-level data from the Baccalaureate and Beyond (B&B) survey conducted by the U.S. Department of Education on student demography, academic achievement, and post-college household income.⁵ For our purposes there are two key sources of variation in the B&B data. First, although the market is highly assortative in that better colleges tend to attract more accomplished students, it is not perfectly so, and each university’s student body exhibits a distribution of HC within each demographic group. Therefore, college quality and HC achievement are imperfectly correlated, which allows us to separate the influence of these two factors on post-college earnings. The second source of variation is AA practices, which create different investment incentives for students in different demographic groups. Using techniques from the empirical auctions literature (e.g., Athey and Haile [8] and Guerre, Perrigne, and Vuong [36] (GPV)), we derive control functions from the model’s equilibrium conditions that map a student’s choice of pre-college HC into the unobserved learning-cost type that rationalizes this HC choice. Because affirmative action policies cause minority and non-minority students to face different HC accumulation incentives, there is a distinct control function for each demographic group. The differences in these control functions is what allows us to separate the effects of HC and learning cost on household income.

We find that college quality and pre-college HC both have a significant effect on the probability that a student graduates. On the other hand, pre-college HC has very little direct influence on household income 10 years following graduation, whereas college quality and unobserved student characteristics play dominant roles. However, pre-college HC has important indirect effects on household income by increasing the student’s graduation probability and allowing the student to enroll in a higher quality college. We do not take a stand on why college quality influences income. It is likely a combination of higher quality colleges allowing students to more easily accrue college-level HC, the superior alumni networks of high college colleges, peer effects, or the signal that attending a high quality college sends. For example, our results are consistent with a view of the world in which a student’s high school grades and test scores help her to

⁴Since all pre-college HC is productive, the overproduction is not purely waste. We simply mean that the overproduction would not have occurred but for the need to compete with other students for seats at high quality colleges.

⁵These data allow us to provide an analysis that accounts for the strategic interactions at work in the nation-wide admissions contest, which would be impossible with data from a single school or even an entire state. There are other nation-wide surveys we might use, but all other datasets that we are aware of have important shortcomings for our purposes. See Section 2.2.

get into and pass the curriculum of a good college, but where post-college household income is determined by the amount of new HC accrued during college, which depends predominantly on college quality and exogenous learning costs. It is important to stress that our results *do not* imply that HC is unrelated to a household's income generating capability. Rather, the appropriate interpretation of our results is that while *pre-college* HC is useful for helping one to enroll in and graduate from a high quality college, only HC acquired during college appears to drive household income for graduates.

We then analyze the relative magnitudes of the productive and competitive channels of HC accumulation incentives. To the best of our knowledge, ours is the first paper to do so. We find that the competitive channel is stronger than the productive channel for all but the highest-achieving 15% of students overall. The relative strength of the competitive channel increases for students with higher learning costs, being roughly three times as powerful as the productive channel for the middle 50% of the learning cost distribution. This suggests that a great deal of pre-college HC accumulation is "wasted" in the sense that it would not be accumulated in a first-best world where the students' learning costs are publicly known and the resulting student-college match is (mostly) assortative without the need for a contest.

We use our estimates of the drivers of post-college outcomes to conduct counterfactual analyses of the effects of AA schemes on enrollment, graduation rates, and household income. AA has a long history in the United States that stretches back to the Kennedy Administration. It is motivated by its proponents on grounds of racial disparities in college placement and a desire to achieve greater diversity within student bodies at top universities.⁶ For example, in 1988 16% of all new college freshmen were under-represented minorities (Black, Hispanic, or Native American) who accounted for only 10% of new enrollees within the top fifth of US colleges, despite substantial considerations for race in the admissions process. This disparity is in turn driven by gaps in pre-college academic achievement: in that same year, median minority GPAs and SAT scores were slightly below the 25th percentile for Whites and Asians.

Changes to admissions rules have complex impacts on the productive and competitive channels of incentives, which in turn influence the students' endogenous choices of pre-college HC. Color-blind and proportional systems have opposite effects on the HC investments of minority students (relative to the status quo). In a color-blind scheme, minority students with high learning costs reduce their HC choices. This is because a high learning-cost minority student must out-compete a larger number of other students to significantly improve her college assignment. In other words, the competitive channel is weaker under a color-blind system for these students, which in turn reduces their HC choices. Under a color-blind system, minority students with low learning costs find that the quality of the available college seat improves rapidly as she out-competes her rivals, which strengthens the competitive channel and increases her HC investments. For nonminority students, the effects of shifting from the status quo to color-blind are the opposite, but small in magnitude. The changes in HC investment from a shift toward a

⁶We remain agnostic about whether diversity is an end itself or whether it is instrumental in providing a high quality college education (see *Regents of the University of California v. Bakke* [1]).

proportional quota from the status quo are larger in magnitude and reversed in sign. The reasoning is essentially symmetric to the color-blind case: the competitive channel is stronger (weaker) for high (low) learning cost minority students, which leads to more (less) HC accumulation.

We find that admissions would be quite different under a color-blind system. Under the status quo AA regime generating the data, minorities are under-represented at the best schools and over-represented at the worst ones, but a color-blind college admissions scheme would result in a shift of minority students into lower ranked schools (e.g., minority enrollment in the bottom quintile of colleges increases by 50% and decreases in the top quintile by 25%). Under a proportional quota, the opposite occurs: minorities enroll in higher quality programs with (mechanically) proportional representation in each college quality quintile.

The impact of AA programs on graduation rates are complicated by the joint effects of changing enrollment and altered HC investments. In a color-blind system, the lowest learning-cost minority students have higher graduation rates because of increased HC investments, despite enrolling in worse schools. In contrast, the highest learning-cost minority students both have less HC and enroll in worse colleges, both of which reduce graduation rates. In a proportional quota, low learning-cost students invest less, but enroll in better schools, yielding no net effect. The highest learning-cost students both enroll in better colleges and invest more in HC, which provides a boost to their graduation rates. Moreover, failing to account for the endogeneity of HC choices can cause predictions of the graduation-rate effect size to be off by more than 50%.

Since household income 10 years after graduation is driven by college quality and learning cost under our empirical estimates, the counterfactual results are simpler to explain. Color-blind systems place minority students in worse colleges, so their average household incomes drop by \$818 relative to the status quo. Under a proportional system, minority student household income increases by an average of \$1,936 due to better college placements. The changes for nonminority students are in the opposite direction and less than $1/5^{th}$ as large. In terms of equivalent variation, the net utility effect on minorities of a color-blind (proportional) system is a decrease (increase) of roughly 2%-3% of household income.

The remainder of this paper has the following structure: in Section 2, we describe the US college market structure and the data that will be used. Section 3 outlines the theoretical model on which the econometric exercise is based. In Section 4, we describe our semi-parametric identification results, our estimator, and an extensive series of robustness checks and model extensions. In Section 5, we discuss the results of estimation, and in 6 we present the counterfactual exercises. Section 7 concludes and briefly describes directions for future research.

1.1. Related Literature. Coate and Loury [22] and [23] were among the first to explore a theory of endogenous educational achievement as a function of AA rules. Both of these papers assume homogenous firms, binary HC investment, and no strategic interaction between the students/workers. Olszewski and Siegel [54] and Bodoh-Creed and Hickman [15] introduce contest models of college admissions that include heterogeneous school quality and endogenous HC choices. Olszewski and Siegel [54] does not include AA or noisy human capital, both of which

are crucial for our identification strategy, so we base our empirical model on Bodoh-Creed and Hickman [15]. To explore the empirical relevance of these theories, Cotton et. al. [25] replicates the Bodoh-Creed and Hickman [15] framework in a real-effort field-experimental classroom setting with junior-high-school-age students. The study involved human capital investment (math learning) and shifts in (short-term) rank-order monetary incentives that mimic the difference between color-blind college admissions and AA via representative quotas. They found strong empirical evidence that students are capable of responding to changes in competitive incentives in ways that closely conform to the predictions of Bodoh-Creed and Hickman [15].

Several other papers have examined observational data to see whether students respond to changes in the affirmative action system.⁷ For example, the “Texas Top 10%” (TTT) program allowed any Texan students in the top 10% of their class to enroll in the University of Texas campus of their choice, which was implemented following the 1996 *Hopwood v. Texas* decision that ruled race could not be used as a factor in university admissions. Cullen et al. [26] showed that at least 5% of students enroll in lower quality high schools in order to have a higher probability of participating in the TTT program, and the authors argue that this is a lower bound on the fraction of strategizing students. Cortes and Zhang [24] argue that achievement increases in low quality schools in response to the program. Kapor [44] estimates a model of student application and enrollment decisions to compare the TTT program with other enrollment systems, and he argues that the TTT program increased enrollment of minority students, partially by changing student application behavior. Since the model does not include private information or an endogenous HC decision, Kapor [44] focuses on questions other than those we pose regarding the determinants of post-college income, the effect of AA on pre-college HC accumulation, or the relative importance of the competitive and productive channels of incentives.⁸ Akhtari and Bau [2] analyzed data from Texas following the 2003 *Grutter v. Bollinger* Supreme Court decision that reinstated the legality of race-based affirmative action. The authors found that the minority-nonminority achievement gap narrowed, and survey data suggested that student behavior and expectations responded to the resulting change in admissions policy.⁹ We take from the above research that students are aware of and respond to the incentives provided by the college enrollment system.

There is a sizable literature examining the role of AA in college admissions, including Bowen and Bok [16], Kane [43], Chung and Espenshade [20], Chung, Espenshade, and Walling [21], Epple, Romano, and Sieg [29], Arcidiacono [7], and Howell [41]. These papers find that AA plays a significant role in shaping minority educational outcomes, especially at selective institutions. Some of these studies lack nation-wide data on college admissions, and others condense college heterogeneity to coarse terms for tractability. Moreover, exam scores are used as an exogenous proxy for student ability. If scores are jointly determined by both ability and market-based incentives for investment, then this is problematic.

⁷There is ample evidence that Universities respond to changes in the legality of different forms of affirmative action (e.g., Antonovics and Backes [4], Antonovics and Sander [5], Long and Tienda [47], Yagan [63], and others).

⁸The model includes an endogenous application decision and several frictions that our model abstracts from.

⁹Ferman and Assunção [31] also use a natural experiment in top Brazilian universities induced by a court ruling, though important differences in their application preclude a direct comparison to our results.

Our paper is one of the first to combine market-wide data, rich supply-side heterogeneity, post-college outcomes, and endogenous pre-college HC. Our model and methodology allow us to study pre-college achievement, enrollment, and post-college outcomes on a nation-wide basis, rather than just focusing on one section of the country or market. We are also able to measure the effect of affirmative action without making any a priori assumptions on the parametric form of the demographic admissions bonus for minority students. Since we find that affirmative action benefits minority students the least at the top and bottom of the college quality spectrum, it is plausible that the average effect over the quality spectrum is relatively weak. Therefore, it is not obvious how to compare our results to the prior literature that assumed a parametric form for the effect of affirmative action (e.g., Hinrichs [40]).¹⁰

Hickman [39] also studies college admissions using a contest framework, but this earlier work required several assumption that only allowed for less credible answers to a more restricted set of questions. First, Hickman [39] relied on college-level data, which required that the model assume perfectly assortative matching between students and colleges based on SAT score. Second, the model could not separate between the effects of HC and learning costs on household income, and so assumed that learning costs had no effect on income. Third, graduation probability was identified from parametric assumptions. These limitations made the counterfactual predictions about the effects of AA on outcomes such college assignment, graduation probability, and household income less credible than those provided in this paper.

Of course, the current paper is subject to its own limitations as well. We abstract from the intricacies of the admissions process and concentrate on the link between achievement and final college placement outcomes in terms of matriculation. We do not explicitly model “supply-side” concerns—e.g., decisions on how many students to admit and how much to charge them—but instead we treat college seats as fixed objects of known quality in order to concentrate on student HC investment. To the extent that supply-side competition plays a role, this work can be seen as complementary to models such as Epple et. al. [29]; Chade, Lewis, and Smith [19]; Fu [35]; Azevedo and Leshno [9]; and Fillmore [32] who treat these forces explicitly.

Our paper also contributes to an established literature on the returns to attending a more selective college.¹¹ One of the first such papers was Brewer, Eide, and Ehrenberg [17] who estimated large income gains to attending an elite college after modeling college choice as a function of net costs; they also found that these gains persist after attempts to control for selection on unobservables. Dale and Krueger [27] (DK) attempted to control for selection on unobserved student characteristics by using overlap across 1970s college applicants in the sets of schools that accepted and rejected them. DK found small average returns to a more selective college, but that for low-income applicants the economic gains were meaningful. Black and Smith [12] (BS04) re-examined this question with a propensity-score matching estimator that was designed to address

¹⁰The difference between our work and the prior literature is magnified by our focus on the effect of AA on final enrollment outcomes rather than college admissions that merely determine the enrollment choice set.

¹¹We discuss the mismatch hypothesis and the related literature at length in Appendix G.

potential failures of a common support condition used by linear-in-parameters models in previous studies. They found evidence from their alternative approach that all students—rich and poor—benefit from an elite education. Long [46] re-visited the methodologies of DK and BS04 with a more recent dataset (NELS, 2000 wave), and found robust evidence across methodologies for significant gains to college quality in terms of graduation rates and household income. Finally, Black and Smith [13] use multiple proxies for college quality, rather than just measuring quality by median exam scores of enrollees as done in previous studies, and they find evidence that the existing literature under-states the gains to attending a higher quality college.

Our paper draws insights from previous work and also contributes to the literature in some new ways. Like Black and Smith [13], we use a measure of college heterogeneity that incorporates information from multiple proxies for quality. Like Long [46], we focus on household income as our preferred post-college outcome, as it encapsulates the total economic gains from a more selective college, which may influence both one’s own earnings as well as earnings of one’s spouse. Like DK, our empirical strategy acknowledges the importance of selection on unobservables; however, unlike DK who control for student characteristics that are unobserved to the econometrician but observed to the college admissions officer, our approach allows for the unobservable student characteristics driving selection to be known only to the student.

Within the literature on returns to college quality, we are the first to explicitly model and quantify the role of unobserved student characteristics as a productive factor (separate from the error term) in generating post-college income.¹² Previous work has instrumented for unobservable characteristics or included them in the error structure in order to focus on the return to college quality, but counterfactual analysis of alternative college admissions scenarios requires knowing how student characteristics contribute to match quality as well. As an applied question of independent interest, little is known about the relative importance of students’ unobserved types and college quality in determining match utility, and we are the first to quantify an answer. We point-identify students’ learning costs at the individual level and show that they play a secondary, but important, role to college quality. We also show that our estimated costs are correlated with childhood home environment characteristics (see Appendix D) that one would expect to be important drivers of learning cost and life success.

2. DATA AND MARKET STRUCTURE

To conduct our analysis, we need enough data to effectively model the incentives facing college applicants when they make their HC accumulation choices. Our estimates require metrics for the HC of students that graduate from each college, measures of the quality and graduation rate of each college, and statistics on the post-college outcomes for each student. We begin by describing the combination of college-level and individual-level data we use.

¹²For example, Arcidiacono [7] allows for a finite set of unobserved student types that induce correlation in the error structure for the decisions to apply to different colleges, enroll in college, and choose a major. In contrast, we allow unobserved characteristics to exist on a continuum and include them explicitly in the income equation as a productive factor separate from the error term, while abstracting from the finer details of the applications and major-selection process.

We use US college data for academic year 1992-1993 for two main reasons. First, one can reasonably assume AA policies were stable and understood by decision-makers at that time. The only successful legal challenge prior to 1993 was in 1978, when the Supreme Court declared quotas unconstitutional in *University of California v. Bakke* [1]. The second reason for studying academic year 1992-1993 is that individual-level data on students graduating during this academic year are available from the Baccalaureate and Beyond (B&B) database linking college quality and HC choices to the household income of college graduates from that year. Later waves of the B&B data either did not collect income data or have not yet reached that point in the panel.

2.1. Colleges. For a sample $L = \{1, 2, \dots, L\}$ of 4-year colleges we have a vector \mathbf{Y}_l of school characteristics. The first is a quality measure derived from data and methodology by US News & World Report (USNWR) for their annual *America's Best Colleges* rankings (see Morse [51]).¹³ We adopt this measure as the college quality index p_l , and we argue that interpreting this index as a reflection of meaningful quality rankings is sensible for three reasons. First, USNWR removes information frictions by providing a wealth of data on many schools, along with advice on how to interpret the data. Consumers' response to this service has been large enough that rankings are now the primary focus of USNWR's business model. Second, the validity of USNWR rankings is reinforced in students' minds by the enthusiasm with which universities advertise their status in *America's Best Colleges*. Third, many previous studies have depicted college quality either with coarse, discrete measures (e.g., flagship state schools versus non-flagship schools) or with relatively simplistic ones such as mean student-body exam score alone. Our measure provides for a continuous transition from low quality to high, and it encompasses a host of factors influencing the college experience such as selectivity, per-student spending, and faculty quality.

The US postsecondary education industry is vast and diverse, with thousands of institutions offering students at least one type of 4-year degree, but many of these specialize in vocational training. Thus, we adopt the USNWR universe of schools as our definition of "the college market." This leaves us with 1,245 non-profit colleges and universities specializing primarily in liberal arts education leading up to a bachelor's degree.¹⁴ This set of schools accounts for the majority of 4-year degree production in the United States. Between the late 1980's and early 1990's the total number of incoming freshman each year was roughly 1.6 million students.

¹³USNWR computes its quality score as a weighted arithmetic mean of a school's quantile rank in 15 quality indicators, falling into 6 different categories: *Reputation Rank* (based on survey data from college presidents and deans; 25% weight), *Student Selectivity* (acceptance rate for 1992 freshman class, yield rate for 1992 freshman class, % of enrollees in top 25% of high school class, and mean SAT/ACT score among enrollees; 25% weight), *Faculty Resources* (comprising student-faculty ratios, % of full-time faculty with terminal degrees, % of faculty on part-time status, average salary and benefits for full-time faculty, and proportion of classes with fewer than 20 students; 20% weight), *Financial Resources* (per-student education expenditures and per-student other expenditures; 15% weight), *Graduation Rate* (% of students in 1983-1986 freshman classes who graduated within 6 years; 10% weight), and *Alumni Satisfaction* (% of living alumni who donated to AY1991-1992 fund drives; 5% weight). USNWR computes the quality metric separately within 4 different groups of schools (National universities, national liberal arts colleges, regional universities, and regional liberal arts colleges, see Morse [51]), but we modify their method slightly and define quality indicator quantile ranks across the entire universe of schools. This produces a single index ranking that applies to all colleges in the sample. We used the 1994 edition because several of its quality indicators are on a 2-year lag.

¹⁴We discuss the differences between our nation-wide market model and an alternative model allowing for regional markets in Section 3.5. Our main estimates are similar under both models.

TABLE 1. Racial Representation Within Academic Quality Quintiles

Group	Total Share	Tier I Share	Tier II Share	Tier III Share	Tier IV Share	Tier V Share
Minority:	15.81%	10.23%	13.09%	11.75%	22.16%	21.81%
Non-Minority:	84.19%	89.77%	86.91%	88.25%	77.84%	78.19%

The other college-level data are provided by the National Center for Education Statistics (NCES) through their Integrated Postsecondary Education Data System (IPEDS), which includes school-level enrollment for all first-time freshmen (including full-time and part-time) by race for Whites, Blacks, Hispanics, Asians or Pacific Islanders, and American Indians/Alaskan Natives. For the 1988 incoming class, we have freshman headcount, denoted M_l , for underrepresented minorities—Blacks, Hispanics, and Native Americans—and a headcount, denoted N_l , for all others—Whites and Asians. Aggregating this information across schools allows us to compute $m = \hat{a}_{l=1}^L M_l / \hat{a}_{l=1}^L (M_l + N_l)$. IPEDS also allows the researcher to compute race-specific, 6-year graduation rates for each college campus, which we denote by G_{jl} , where $j \in \{M, N\}$ refers to Minority and Nonminority students. In total then, the data representing schools $l = 1, \dots, L$ are denoted $\mathbf{Y}_l = \{p_l, M_l, N_l, G_{Ml}, G_{Nl}\}_{l=1}^L$.

Colleges are separated into five tiers with Tier I representing the top quality quintile of college seats. Entries in Table 1 represent the fraction of all students within a tier that are of a given demographic group. In a hypothetical world with no under-representation, each cell would be the same as the overall share of each race group. “Minorities” are demographic groups that are under-represented in the top three tiers and over-represented in the bottom two. Similar patterns hold as well for minority sub-categories: Blacks, Hispanics, and Native Americans are individually under-represented in top tiers. The reverse is true for Whites and Asians who comprise the “Nonminority” group. Minority and nonminority students are defined as they are because AA in college admissions specifically targets under-represented minority groups.

2.2. Students. Our individual-level data on the student population comes from the 1993 Baccalaureate and Beyond Survey (B&B), which randomly samples colleges and then samples students graduating in academic year 1992-1993 within each college. The data contain several variables pertaining to pre-college investment for student $i \in \{1, 2, \dots, Ig\}$. The first variable we draw from the B&B survey is annual household income after 10 years in the workforce, denoted w_i . Household income has two advantages relative to individual-level measures of income. First, it captures the effect of educational assortative mating, which enhances a student’s welfare by increasing the income of his or her spouse. Second, including both men and women increases the power of our estimates. Previous work has focused on employed men to avoid confounding the effects of gender on individual salary/wages with other features of the data, but the distribution of household income in our data is essentially identical for men and women.¹⁵ Appendix

¹⁵In addition, under our metric of HC (see Section 4), the average HC of women is only 0.08 standard deviations higher than the population mean.

H.2 provides a robustness check that repeats our analysis using only the individual-level income and salary measures of male students 10 years after graduation. We find that our main utility parameter estimates do not change significantly. The top panel of Table 2 provides summary statistics for college-level variables and the bottom panel summarizes student-level data.

Our exclusion restriction, Assumption 4.4, is that minority status does not affect income conditional on school quality, HC, and unobserved student characteristics, but this might fail if marriage is less assortative in terms of income across minority and nonminority groups. However, the correlation between self- and spousal-income is 0.176 and 0.142 for nonminorities and minorities respectively. If we construct a conservative test for the difference of those correlations at the 90% level, we fail to reject the null hypothesis of equality.

The second piece of information we require from the data is a measure of pre-college academic preparation. Two outcome variables which researchers and college admissions officers focus on most for assessing achievement are exam scores, denoted e_i , and academic record as measured by grade point average (GPA), denoted a_i . In the B&B data, e_i takes the form of either the ACT or the SAT, both of which are standardized college entrance exams.¹⁶ B&B does not contain high-school GPA directly, but it has other information including GPA from all non-major and non-minor courses.¹⁷ We adopt college non-major/non-minor GPA (NMGPA) as a proxy for high school GPA.¹⁸ One might worry, however, that NMGPA at institutions with different qualities might not be comparable,¹⁹ but there are several reasons to believe it is a meaningful reflection of pre-college academic preparation. First, one might worry that grade inflation is correlated with school quality, which would mechanically make any college GPA a predictor of college quality. However, we find that the correlation between major GPA and college quality is -0.006, suggesting grade inflation across schools at a given point in time is negligible.²⁰ Second, NMGPA primarily consists of coursework during the first one or two years of college, meaning the student has not had much time to accrue college-specific HC. Third, under the liberal arts model typical of American higher education, students take a wide variety of non-major courses focusing on general education, much as they did in high school. These courses are not within the student's specific area of interest or idiosyncratic talent, so the fact that better students are matched with better universities is of less importance. In addition, content covered in general education courses is largely independent of a student's major, with English and Economics students

¹⁶The companies which develop these exams produce concordance tables which allow one to relate ACT scores into the SAT scale and vice versa.

¹⁷We do not use the information on major in our household income regressions. Although major choice plays a role in some prior studies (e.g., Arcidiacono [6]), our main focus is on disentangling the roles of school quality, HC, and unobserved characteristics in determining college outcomes. For major choice to strongly affect our analysis, there would have to exist large variation in the quality of education between majors at a given college relative to the inter-college variation in quality that our aggregate college quality metric captures or large variation in the fraction of students majoring in high earning subjects that is also correlated with college quality.

¹⁸Specifically, we assume that exam scores plus high-school GPA together contain roughly the same information about a student's pre-college HC as exam scores plus NMGPA together.

¹⁹Previous research on high school achievement has frequently used high-school GPA, but the interpretation of GPAs from different high schools suffers the same problem due to different grading standards and course offerings.

²⁰One would expect all grades to be inflated if cross-sectional grade inflation trended with college quality.

being required to take many of the same non-major courses. Fourth, non-major courses are introductory in nature and use textbooks that are standardized across schools. For example, Lopus and Paringer [48] reported in 2011 that the two most popular Principles of Economics textbooks—Mankiw [49] and McConnell, Brue and Flynn [50]—had a combined 40% market share.²¹ Fifth, in another dataset containing similar information, the National Longitudinal Survey of Youth-1997 wave, or NLSY97, we find that NMGPA and high school GPA are highly correlated.²² Finally, we experimented with models that defined HC based solely on test scores, and the results are reported in Appendix H.1. Because the model provided a poor fit to the data and implausibly low levels of assortativity between HC and school quality, we rejected this model in favor of one that includes proxies for high school GPA in the HC metric.

3. THEORETICAL MODEL

Following Bodoh-Creed and Hickman [15], we model college admissions as a Bayesian game where high-school students are characterized by a privately-known type that governs the costliness of HC production and the payoff from enrolling in college. On the other side of the market are colleges that have preferences over students based on their HC and race. Colleges mechanically admit the best students from each demographic group that the school can attract, and the trade-off between diversity and pre-college academic achievement is captured by the colleges' AA plans. In this section we outline the model and summarize technical results proven in Bodoh-Creed and Hickman [15]. We summarize the technical assumptions required in Appendix A, although the empirical model we estimate satisfies all of these assumptions.

3.1. Agents. There are two demographic subgroups within the student population, minority students (M) and nonminorities (N), and the demographic class of each student is observable. There is a continuum of students of total mass 1 with mass m of them in the minority group. We often refer to our model with a continuum of students as a limit model to emphasize the fact that it can be viewed as a limit of a model with a finite number of students as their number approaches infinity. We elaborate on this point in Section 3.4.

²¹Lopus and Paringer [48] found a high degree of content overlap across 26 different principles textbooks.

²²This begs the question why we did not use the NLSY97: there are three reasons why we believe that B&B is superior for our purposes. The first is sample size: the NLSY97 sample contains less than 20% as many observations (875 total) as B&B with the requisite data for estimating our household income equation (i.e., college graduates with income reported several years after graduation). The second major problem is that NLSY97 respondents are spread across five separate cohorts of students. This means that payoffs would have to be estimated using an unbalanced panel of earnings data, which would require estimation of additional parameters to control for the potential effects of workforce experience and/or other unobservable differences across college graduation cohorts. By contrast, the B&B database contains only students who graduated college in the 1992-1993 academic year and entered the workforce at the same time under the same general market conditions. Third, the NLSY97 respondents had the misfortune of experiencing the worst economic downturn (unrelated to war) since the great depression. This further exacerbates the previous two problems: not only was the financial meltdown of 2008-2009 a significant shock for all respondents, but it likely affected the various waves of college graduates in different ways by interrupting career paths with involuntary unemployment spells at different points during the formative, early career years. Adequately controlling for these systematic differences would require allowing for the economic returns to college quality, pre-college HC, and unobserved student characteristics to vary by college graduation cohort, requiring many more parameters to be estimated with much less data.

Table 2. Descriptive Statistics

Variable	Obs	Mean	Median	Std. Dev.
IPEDS/USNWR: (school-level data)				
6-Yr Graduation Rate: M	961	0.3427	0.2812	0.2159
6-Yr Graduation Rate: N	961	0.4831	0.4568	0.2295
Freshman Cohort Size: M	1,245	145.37	46	254.56
Freshman Cohort Size: N	1,245	660.86	378	788.68
College Quality Index	1,245	0.4842	0.4598	0.2132
BACCALAUREATE AND BEYOND : (individual level data, college graduates only)				
SAT/SAT equivalent scores: M	500	820	820	220
SAT/SAT equivalent scores: N	4,980	990	980	190
Academic Record (GPA): M	500	2.7	2.7	0.6
Academic Record (GPA): N	4,980	3.0	3.0	0.6
College Quality: M	500	0.5174	0.5236	0.2140
College Quality: N	4,980	0.5846	0.6051	0.2024
10-Year Household Income: M	260	\$100,700	\$89,500	\$52,500
10-Year Household Income: N	2,800	\$108,300	\$92,700	\$77,300

As per USDOE data security requirements, in order to protect anonymity of B&B respondents, sample sizes have been rounded to the nearest 10, dollar figures have been rounded to the nearest \$100, and GPAs have been rounded to the nearest 0.1.

Each student has a privately-known learning cost type $q \in [q, \bar{q}]$, and the distribution of q in group $j \in \{M, N\}$ follows a cumulative distribution function (CDF) $F_j(q)$. We denote the unconditional type distribution by $F_K(q) = mF_M(q) + (1 - m)F_N(q)$. Each agent's strategy space, $S = [s, \bar{s}]$, is the set of attainable HC levels. We view the HC choice as the outcome of a long-run plan for the accumulation of HC over several years leading up to the college admissions process. These are observable (e.g., through standardized exam scores and high school GPAs) and \bar{s} is the minimum required to attend some college. In order to produce s units of HC, a student incurs cost $C(s; q)$ which is increasing in both s and q .²³ Investment costs can arise in various ways, such as from a leisure–investment tradeoff or psychic costs from difficult learning activities.

Mathematically, the difference between type q and HC s is that the former reflects exogenous factors affecting HC costs and the latter arises from an endogenous HC decision of the student. Learning-cost type q amounts to a student-specific effect that encapsulates cognitive and non-cognitive factors, and q may be influenced by forces internal to the individual (e.g., natural curiosity) and external (e.g., home environment, primary/secondary school quality, parental education and financial resources). For example, consider the case of a parent who may enrich his daughter's educational experience by spending time reading with her. In a low-income

²³One might have imagined a more complex model where the students can generate distinct kinds of HC through different costly activities. All that is essential about our modeling choice is that the schools reduce each student's portfolio of HC to a single index and that the schools do this in the same way. It is not necessary to understand exactly how a student created their HC to answer the questions we have posed.

household where the parent must work two jobs, his time may be more constrained than in an affluent household where the parent has a single, high-paying job. In this example, the parent's and child's choice of how much time to spend on learning activities is encapsulated in s , whereas the opportunity cost of time is reflected in q . Asymmetric type distributions reflect these factors, many of which are correlated with race. When we condition outcomes on unobserved types in our empirical framework, q controls for many environmental correlates of race even if they do not appear in our model explicitly. We remain agnostic on the exact interpretation of q . Appendix D demonstrates that our estimated q values are strongly correlated with childhood home environment characteristics that one would expect to be important drivers of learning cost. In contrast, s represents everything that is subject to the endogenous control of the student. This could include standardized exam scores, grades, or extracurricular activities like reading, playing an instrument, or leadership exercises. Although the measure of s we implement econometrically is a single index based on standardized exam scores and a proxy for high school grades, the ideal measure would include all of the relevant margins a student can adjust to accrue human capital.

One empirical regularity of American college data is a nontrivial degree of academic heterogeneity among students on a given college campus (even within demographic groups). To rationalize these deviations from perfect assortativity between school quality and student HC, our model includes market frictions in the form of a random matching shock to the colleges' perceptions of each student's HC choice. This shock is observable to colleges at the time the student applies, but not to the student while she is investing in her HC. The unobservability of the shock to the student is consistent with an interpretation of HC choice as the gradual accumulation of learning over years, while the shock is the result of events that occur within a short time before her application to college and that are out of the student's control. The matching shock is applied to each student's choice of HC, s , to generate a noisy HC (NHC) value, t , that is commonly observed by all of the colleges. We assume the noise enters additively, so if student i chooses s , the associated NHC is $t = s + \epsilon_i$, $\epsilon_i \sim F_\epsilon(\epsilon)$, and the draws of ϵ are independent across students. The assignment of students to colleges becomes more assortative as the variance of ϵ shrinks, and the market becomes a lottery as the variance grows.

3.2. Payoffs. On the other side of the market there is a continuum of colleges with total mass 1. Each college's quality is described by an index $p \in [\underline{p}, \bar{p}]$ that is distributed $P = F_p(P)$. By assuming the measure of students and college seats are the same, we abstract from the extensive margin of college attendance and focus on the competition for admission to the best colleges conditional on entering the market. Both college quality and HC are intrinsically valued: match utility $U(p, s, q)$ results from a student with type q having HC s and enrolling in a college with quality p . The ex post payoff to agent i in group $j \in \{M, N, G\}$ is the match utility minus the cost of achievement, $U(p, s, q) - C(s, q)$.

3.3. Allocation Mechanisms. In an admissions contest, the quality of the seat allocated to a student is a function of how her NHC realization compares to the distribution of NHC across the

population of college applicants. AA schemes cause the NHC realizations of minority and non-minority students to be compared to the total distribution of NHC differently, which makes the contest asymmetric between the two demographic groups. We consider color-blind (cb), proportional quota (q), and admissions preferences (ap) AA systems.

The allocation mechanism is an assignment mapping $P_j^r : R \rightarrow [p, \bar{p}]$, $j \in M, N$ and $r \in \{ap, cb, q\}$, that maps student NHC realizations into college assignments, where $r \in \{ap, cb, q\}$ indicates the allocation mechanism and $j \in M, N, K$ indicates a demographic group or the unconditional population.²⁴ P_j^r is an endogenous object that depends on the form of AA at work in the market, the distribution of college qualities, and the equilibrium distribution of NHC in the population, which is in turn determined by the endogenous choices of HC by the students. To facilitate discussion, we denote the equilibrium CDFs of HC and NHC as G_j^r and H_j^r , respectively. H_j^r is a convolution of G_j^r and $F_{\#}$ with density $h_j^r(t) = \int_{\#} f_{\#}(\#) g_j^r(t - \#) d\#$.

We define AA as the difference between the level of NHC required for a minority and non-minority applicant to be admitted to a school. The fraction of minority students on the campus of a particular college is the result of the interaction of the school's AA policy with the densities of the distributions of NHC for each group. The assignment mappings describe how admissions practices work at different points in the college quality spectrum.

In the context of our model, the problem facing the students is rather simple. Instead of requiring common knowledge of the strategies of the other agents, we require that the students understand the assignment mapping. In order to do this, the students merely need to understand the admissions standards of colleges of a quality similar to those to which they will eventually enroll. This information can be gleaned from public data about (for example) the typical SAT and high school GPA of freshmen enrolling at different colleges.²⁵ To the extent that college students make errors in this process, it would be reflected in a lack of assortativity in the match, an issue that is discussed further in Section 5.1 and Appendix H.9.

The color-blind allocation mechanism is the simplest to define since it assigns students to schools assortatively with higher NHC realizations leading to assignment at higher quality colleges. Since demographics do not affect the assignment, $P_M^{cb}(t) = P_N^{cb}(t) = P^{cb}(t) = F_P^{-1}(H_K^{cb}(t))$. In words, a student with NHC realization t at quantile rank $H_K^{cb}(t)$ of the NHC distribution (for all students) is placed at a school with the same quantile rank in the college quality distribution, $F_P^{-1}(H_K^{cb}(t))$. Since the assignment mapping is the same for both groups, marginal investment incentives, conditional on type q , are also identical across groups.

Under a quota scheme, students from group $j \in M, N$ are reserved sets of seats with qualities distributed $P = Q_j(P)$, where $mQ_M(p) + (1 - m)Q_N(p) = F_P(p)$ is required for feasibility.

²⁴To the extent that admissions is influenced by other factors correlated with race (e.g., legacy admissions), then our empirical model will attribute these to race, which we think of as “effective affirmative action.” In that case, AA counterfactuals should be re-interpreted as modifying explicit racial AA and implicit racial AA based on correlated factors (e.g., legacy status). We could study “affirmative action” explicitly in terms of these other factors, but we begin to encounter data sparsity issues. For an example, see Appendix H.3, which explores regional admissions differences.

²⁵See Bodoh-Creed and Hickman [15] for further discussion of the relationship between the continuum and finite-agent games

Members of each demographic group compete in disjoint contests for prizes only against the other members of their same group. The resulting assignment map is $P_j^{pq}(t) = Q_j^{-1} H_j^{pq}(t)$, which means a student with NHC t at quantile rank $H_j^q(t)$ is placed at a school with the same quantile rank in the distribution of colleges allocated to her demographic group, $Q_j^{-1} H_j^q(t)$. The most familiar member of this class is a proportional quota where $Q_M = Q_N = F_P$, meaning that a fraction m of seats at each point in the college quality spectrum are reserved for minorities.

Although proportional quotas violate the US constitution, they provide a useful benchmark relative to a color-blind mechanism. When racial asymmetries in learning costs exist—because race is correlated with childhood school quality, for example—proportional quotas are designed so that these differences are not reflected in the fraction of students from each demographic group present on each college campus. A color-blind mechanism allows asymmetries in learning costs between the demographic groups to be maximally reflected in minority student enrollment.

Finally, the only legally permissible form of AA in US college admissions is an admission preference system where race is taken as a “plus factor” among other considerations like grades and test scores. Under such a system, the NHC value of each student is compared to the distribution of the total student population, but the demographic of each student determines how this comparison is made. An admission preference is defined by a markup function $\Phi : \mathbb{R} \rightarrow \mathbb{R}$ that transforms the NHC levels of the minority students. The assignment mappings are

$$(1) \quad P_M^{ap}(t) = F_P^{-1} m H_M^{ap}(t) + (1 - m) H_N^{ap}(\Phi(t)), \quad P_N^{ap}(t) = F_P^{-1} m H_M^{ap}(\Phi^{-1}(t)) + (1 - m) H_N^{ap}(t)$$

In words, a minority student's NHC level t is compared to the raw NHC of other minority students and marked up when compared to nonminority students. Conversely, NHC for a non-minority student is compared to the raw NHC levels of other nonminority students and it is “de-subsidized” for comparisons to NHC levels of minority students.

Bodoh-Creed and Hickman [15, Theorem 4] proves that quota and admissions preference systems have the same set of equilibria. Our empirical model is based on a (non-proportional) quota system that is outcome equivalent to the real-world admissions preference. Since the allocation of students to schools is known ex ante under a quota, this simplifies estimation.

Theorem 3.1. [Theorem 4, Bodoh-Creed and Hickman [15]] $P_j^q(t) : \mathbb{R} \rightarrow \mathbb{P}$, $j \in \{M, N\}$, is the result of an equilibrium of some quota system with (Q_M, Q_N) if and only if there exists an admissions preference system with some Φ that has the same equilibrium assignment mappings and strategies.

Although we base our estimator on a quota-system model, it is straightforward to derive the equivalent admissions preference system. Note that when a quota system is equivalent to an admissions preference system, we have $P_j^{ap}(t) = P_j^q(t)$, $j \in \{M, N\}$. Given an estimate of $P_j^q(t)$, $j \in \{M, N\}$, we can then write

$$(2) \quad P_M^q(t) = P_M^{ap}(t) = F_{P_M}^{-1} [H_M(t)] = F_{P_N}^{-1} H_N(\Phi(t)) = P_N^{ap}(\Phi(t)) = P_j^q(\Phi(t)).$$

Manipulating this equation we find

$$(3) \quad \Phi(t) = H_N^{-1} F_{P_N}^h F_{P_M}^{-1} [H_M(s)]^i .$$

3.4. Equilibrium. The distribution of college seats, the form of the admission system, and the measures and distributions of student competitors are common knowledge prior to individual choices of HC investment. When combined with the equilibrium strategies of the minority and nonminority students, denoted $s_M(q)$ and $s_N(q)$ respectively, the students can forecast the form of the assignment mapping P_j^r arising in equilibrium. Each student solves the following optimization problem where $j \in \{M, N\}$ and $r \in \{c, p, a\}$

$$(4) \quad s_j(q) = \arg \max_s E_{\#} U_{P_j^r}(s + \#, s, q) - C(s, q) .$$

In equilibrium, students' beliefs about $s_M(q)$ and $s_N(q)$ must be consistent with the solution to equation (4). Bodoh-Creed and Hickman [15, Theorem 6] proves that an equilibrium exists.

The first-order condition can be written as

$$\underbrace{E_{\#} U_{P_j^r}(s + \#, s, q; a)}_{\text{Productive Channel}} + \underbrace{E_{\#} U_{P_j^0}(s + \#, s, q; a)}_{\text{Competitive Channel}} = \underbrace{C_s(s, q)}_{\text{Marginal Cost}} ,$$

with the left-hand side describing marginal benefits and the right-hand side reflecting marginal costs. The productive channel represents the direct effect of HC on utility, while the competitive channel captures the marginal effect of HC on utility through its role in college placement.

Actual college markets one might study empirically have only finitely many students and colleges, but on the other hand, the finite version of this model is computationally intractable. In this finite model, K_M minority students draw types from the distribution $F_M(q)$, K_N nonminority students draw types from the distribution F_N , and $K_M + K_N$ college seats draw their qualities from F_P . In the limit as $K_M + K_N \rightarrow \infty$ and $K_M / (K_M + K_N) \rightarrow m$ the primitives of the finite games approach those of the continuum model. Bodoh-Creed and Hickman [15] show that the model with a continuum of students approximates the finite game in the following sense.

Definition 3.2. Given $\epsilon > 0$, an ϵ -approximate equilibrium of the K -agent game is a K -tuple of strategies $s^\epsilon = (s_1^\epsilon, \dots, s_K^\epsilon)$ such that for all agents, almost all types and all HC choices we have

$$U_{P_j^r}(s_i^\epsilon(q), s_i^\epsilon(q)), s_i^\epsilon(q), q) - C(s_i^\epsilon(q), q) + \epsilon \geq U_{P_j^r}(s_i^0, s_i^\epsilon(q)), s_i^0, q - C(s_i^0, q)$$

Definition 3.2 describes an approximate equilibrium in terms of incentives: agents that follow an ϵ -approximate equilibrium can gain at most ϵ by deviating. Intuitively, students lose little utility if they follow the easy-to-compute limit game equilibrium. Bodoh-Creed and Hickman [15, Theorem 7] shows the equilibrium of our limit game is an ϵ -approximate equilibrium of the K -agent game and that we can choose ϵ arbitrarily small for K sufficiently large. This justifies our use of a continuum model to approximate the more realistic, but intractable, finite model.

3.5. Motivating Our Model. Our contest model assumes that students agree on the ranking of schools and that schools agree on the quality of each student within each demographic group. As a result, the outcome predicted by our contest model will be nearly assortative in terms of school quality and student HC choice. Violations of assortativity are caused both by AA policies and by an exogenous matching shock. We also assume students are incentivized by the expected household income premium obtained by college graduates. We now motivate these modeling choices by describing how well the raw data fit these assumptions.

First let us consider whether a contest structure fits the patterns we observe in our data. The Spearman rank correlation between SAT scores and college quality for college graduates in the B&B survey is 0.449.²⁶ In addition, we estimate a reduced-form single index for HC in the form of a polynomial in student exam scores and GPA (see Section 4). This single index of academic achievement has a correlation of 0.88 with our metric of college quality. These high correlations imply a high degree of assortativity, in line with a contest model. The fact that students do not seem willing to compromise significantly on college quality is not surprising given the stark quality heterogeneity among U.S. colleges and the range of post-college outcomes (as noted in the introduction).

Our model also assumes a national market for college admissions as has been done in previous work (e.g., Epple, Romano, and Sieg [30]; Chade, Lewis, and Smith [19]; and Fu [35]). The key requirement in our context is that all college applicants have the same access to the college quality spectrum regardless of where they attend high school. For example, we assume that a student from the Midwest can access the full range of school qualities without leaving the Midwest if she so chooses. The median college enrolls students from over one third of the states in the country, 76% of colleges in our sample enroll students from at least 10 different states, and 23% of them enroll students from 30 or more states. To check whether potential regional differences influenced our results, Appendix H.3 extends our analysis to allow for regional markets for higher education that are defined by the colleges in each of the four US Census Bureau regions and the students that eventually enroll in those colleges.²⁷ We found that our main parameter estimates do not significantly change.

As discussed in detail in Section 5.1, deviations from the contest structure would imply a low degree of assortativity (i.e., small correlation) between college quality and academic achievement, which would be expressed as a large estimate of the variance of the matching shock $\#$. In fact, we find a low variance for the matching shock. This suggests that (1) aspects of colleges other than quality (e.g., student demography) are not strong drivers of student preferences; (2) students have access to the full spectrum of college qualities; and (3) information frictions are not causing students to enroll in lower quality colleges than would admit them.²⁸

We also assume that there is a continuum of college seats, meaning no school enrolls a significant fraction of the students. In 1988 (four years prior to the graduating class of academic year

²⁶We find a similar correlation between college quality and median within-campus SAT score in the USNWR data.

²⁷Data sparsity prevents us from repeating our analysis at finer geographical resolution.

²⁸We test for demographic specific matching shock variances or trends in the matching shock variance with respect to HC and find no significant effects. See Appendix H.9.

1992-1993) there were a total of 1,644,340 freshman seats in the market with any single school having only a negligible market share. The largest college in 1988 (Ohio State) had a total market share of only 0.76% of new freshman seats. The mean, median, and standard deviation of market shares for individual universities were 0.091%, 0.047%, and 0.102%, respectively.

Finally, we assume that students are motivated by the expected college income premium. Although there are non-pecuniary benefits to college, it is reasonable to assume that the impact on lifetime income is the largest benefit and that many of the other benefits (e.g., social esteem, student lifestyle) are correlated with the pecuniary outcomes.

4. MODEL IDENTIFICATION AND ESTIMATION

Our basic identification challenge is to disentangle the influence of college quality, HC investment, and privately-known learning-cost on the returns to attending a higher quality college. The empirical auctions literature has developed a set of tools specifically designed to identify private information in game-theoretic models. This literature was pioneered by Paarsch [55] and then revolutionized by Guerre, Perrigne, and Vuong [36, GPV] who proposed a non-parametric estimator for mapping observed bids into underlying private valuations in first-price auctions. We combine this approach with techniques from labor econometrics to parse between the influences of student and school characteristics in producing post-college household income. Intuitively, because AA changes the marginal investment incentives across race groups before college, one can surmise that two students having the same GPA/exam scores but different race must have distinct underlying types. This fact breaks what would otherwise be perfect rank correlation between achievement and unobserved types, and the matching shock breaks what would otherwise be perfect rank correlation between college quality and HC. Together, these components of the empirical model provide a full-rank condition that allows us to separate the influences of HC, college quality, and learning cost in our household income regression.

4.1. Identification. We begin this section by coupling some additional assumptions with the theoretical framework above in order to complete the formal definition of our empirical model. Some of these merely serve the purpose of tractability, and some are crucial for model identification; we make these distinctions clear in our discussion. Overall our identification/estimation strategy is semi-parametric, although we do not require restrictions on the functional forms of the type distributions or the equilibrium college assignment mappings.

Assumption 4.1. (Single Index) The econometrician's measurement of human capital is a single index function of exam scores E and academic record A ,

$$(5) \quad \tilde{S}_i = \tilde{S}(E_i, A_i) = b_1^s E_i + b_2^s E_i^2 + b_3^s A_i + b_4^s A_i^2 + b_5^s E_i A_i + h = S_i + h,$$

where S is the underlying HC choice of the student and we assume $\tilde{S}_i(E_i, A_i) > 0$, $\tilde{S}_a(E_i, A_i) > 0$, $\partial \tilde{S}_i / \partial E_i > 0$, $\partial \tilde{S}_i / \partial A_i > 0$, $\partial^2 \tilde{S}_i / \partial E_i^2 < 0$, $\partial^2 \tilde{S}_i / \partial A_i^2 < 0$, $\partial^2 \tilde{S}_i / \partial E_i \partial A_i > 0$, and $\max_{(E,A) \in \mathbb{R}^2} \tilde{S}(E, A) = 1$. We assume $h \sim N(0, s_h^2)$.

Assumption 4.1 imposes a quadratic form on the single index equation for HC with regularity conditions and a scale normalization to \times the units of \tilde{S} .²⁹ This single index equation condenses the two margins of achievement—GPA and exam scores—into one measure of HC. Implicit here is the idea that individuals/households optimize their portfolio of investment activities (e, a) to generate the highest composite output s_i at the lowest possible cost. h is an error term that is not observed by either the student or the colleges.

Throughout the main text we assume $s_h^2 = 0$, so $\tilde{S}_i = s_i$ —in other words, our single index measures the students' choice of HC without error. With this in mind, we use S_i and $\tilde{S}(E_i, A_i)$ interchangeably. In Appendix H.4 we test whether $s_h^2 > 0$ by adding additional variables into the single index equation \tilde{S} that are potentially relevant to HC acquired during high school. We find that, conditional on observed GPA and exam scores, the additional variables play an insignificant role in \tilde{S} and do not lead to greater predictive power of HC for college placement.

Appendix H.4 also performs a sensitivity analysis to explore how different assumptions about the relative magnitudes of s_h^2 and $s_{\#}^2$ would affect our estimates. Although we can estimate $s_h^2 + s_{\#}^2 = \text{var}(\tilde{S}(E, A))$ from our data, there is no variation that allows us to estimate s_h^2 and $s_{\#}^2$ separately. Our estimates of the household income equation parameters (equation (12)) are essentially unchanged from our benchmark estimates if s_h^2 is less than half of the estimated value of $s_h^2 + s_{\#}^2$. Once $s_h^2 \geq s_{\#}^2$, there is insufficient matching shock remaining to separately identify the effects of p and s on household income. However, our estimate of the relative magnitudes of the effect of observables (p and s) on household income relative the effect of unobserved characteristics (q) is stable.

Assumption 4.2. (Separable Exponential Costs) $C(s; q) = qc(s)$ with $c(s) = \exp(s)$.

Assumption 4.2 imposes separability and a functional form on the cost function. The exponential form is primarily for computational tractability; we found that alternative functional forms do not appear to alter our empirical results in meaningful ways (see Appendix H.9). The assumption that costs are separable in q and $c(s)$ is more fundamental to identification.

Assumption 4.3. Match Utility

Cobb-Douglas Income Production:

$$u(P_i, S_i, q_i) = a_0 P_i^{a_p} S_i^{a_s} q_i^{a_q} z_i, \text{ where } 0 < a_0, a_p, a_s, a_q \leq (0, 1)$$

Graduation Probability:

$$\begin{aligned} r(P_i, S_i) &= \Pr(i \text{ graduates college} | P_i, S_i) \\ &= b_0^r + b_1^r P_i + b_2^r P_i^2 + b_3^r P_i^3 + b_4^r S_i + b_5^r S_i^2 + b_6^r S_i^3 + b_7^r P_i S_i + b_8^r P_i S_i^2 + b_9^r P_i^2 S_i \end{aligned}$$

Expected Log-Utility:

$$U(P_i, S_i, q_i) = r(P_i, S_i) E_z [\log [u(P_i, S_i, q)]] + [1 - r(P_i, S_i)] E_z [\log [ku(P_i, S_i, q)]] \text{ where } k \in (0, 1)$$

²⁹ We experimented with a cubic form $S(E_i, A_i) = b_0^s + b_1^s E_i + b_2^s E_i^2 + b_3^s E_i^3 + b_4^s A_i + b_5^s A_i^2 + b_6^s A_i^3 + b_7^s E_i A_i + b_8^s E_i^2 A_i + b_9^s E_i A_i^2$, but this did not produce a statistically or economically meaningful change in our estimates.

Assumption 4.3 imposes Cobb-Douglas household income production from a match as a function of college quality, pre-college HC, a student's type, and a transitory random shock z . We further assume that a college dropout's income is $k < 1$ times what it would have been had she graduated, and we adopt a flexible, cubic complete polynomial form for the graduation probability function. We adopt a log utility form for the student's preferences over household income, as it is a benchmark choice for lifetime consumption models.

There is a long history of estimating Mincerian regressions to identify the causal effect of college quality on earnings (Brewer, Eide, and Ehrenberg [17]; Dale and Krueger [27]; Black and Smith [13]; and Long [46]), and our household income equation lies firmly within this tradition. This approach abstracts away from the particulars of how/why college produces these effects. Rather, our empirical model views college quality, pre-college HC, and individual characteristics as the raw materials which, over the course of one's college career, give rise to post-college outcomes. One might wonder whether the estimated household income equation still applies under counterfactual equilibria. For example, sorting students into different colleges might alter college quality, although peer effects have proven small or hard to detect on college campuses (Winston and Zimmerman [62]). We now describe a straightforward microfoundation for $u(p, s, q)$ under which the household income production function is invariant across counterfactuals. Suppose students mechanically accrue HC during college through a combination of the resources of the college, P_i ; the student's pre-college preparation, S_i ; and effort during college as determined by the student's learning-cost, q_i . Post-college household incomes are then set by a competitive labor market based on observations of post-college HC, which provides household income $u(P_i, S_i, q_i)$.

The following exclusion restriction is central to our identification strategy.

Assumption 4.4. (Exclusion Restriction) Race does not directly affect match utility conditional on type q_i , investment S_i and college quality P_i if $D_{Mi} = 1$ ($i \in M$) is an indicator for minority status then $U(P_i, S_i, q_i, D_{Mi}) = U(P_i, S_i, q_i)$.

Assumption 4.4 is supported by the labor economics literature. Beginning with Neal and Johnson [53], a body of empirical work has emphasized pre-market factors as explaining the majority of racial wage differentials. Neal and Johnson [53] find that observables capturing skill endowment in late childhood accounted for all of the wage gap for Hispanics and Black females. For Black males they found that pre-market factors accounted for two-thirds of the wage gap, but they also found suggestive evidence that for college-educated Black males the remaining gap was eliminated, conditional on pre-market skill measures.³⁰ More recently, Fryer, Pager, and Spenkuch [34] find that controlling for pre-market factors reduces the black-white wage gap to only 8%, and they argue the remaining gap is best explained by statistical discrimination at hiring combined with post-hiring learning about a worker's productivity. Based on this, the authors estimate that the black-white wage gap is small and statistically insignificant after controlling for pre-market factors, education, and job experience (see Fryer et. al. [34, Table 12, columns 6 and

³⁰This finding is strongly consistent with empirical evidence in later research using other methods and more recent data by Black, et. al. [11] among others.

8]). We interpret evidence by Neal and Johnson [53]; Black, et. al. [11]; and Fryer, et. al. [34] as providing some con dence that 4.4 is a reasonable assumption for college-educated adults.

Although Assumption 4.4 rules out some forms of racial discrimination on the labor market (e.g., taste-based racial animus), our model is compatible with a world in which statistical discrimination occurs.³¹ In Appendix I we discuss different interpretations of the Mincer equation, including ones that admit a role for signaling to employers. In the equilibrium data generating process, statistical discrimination would appear as minority and non-minority students with the same realizations of (s, p) receiving different household incomes because of differing values of q that employers infer from the joint distributions of (p, s, q) in conjunction with any other information employers possess (e.g., HC mechanically acquired during college) conditional on race.

It is important to emphasize that our exclusion restriction does not rule out influences by a host of other important factors that are highly correlated with race and influence learning costs. Recall that q is a student-specific fixed effect that implicitly subsumes environmental factors such as parental education and income, availability of early-life developmental resources, and K-12 school quality, as well as cognitive and non-cognitive ability which are influenced by childhood experience. Therefore, our exclusion restriction requires only that racial affiliation plays no direct role in post-college household income conditional on HC, college quality, and the environmental and idiosyncratic factors contained in q .

If our exclusion restriction were violated, then one might expect the impact of animus on household income to be loaded onto our estimates of q and a_q , as these are the only components of the household income equation that can reflect systematic variation between the minority and nonminority groups conditional on fixed (p, s) . Therefore, if minority status predicts q , then Assumption 4.4 might be violated. Appendix D tests this idea by regressing our estimated q 's on a rich set of controls for childhood home environmental factors, a race dummy, and a full set of race interactions. We find that, while home background and childhood socioeconomic controls are highly significant predictors of one's estimated q , a joint F-test on the race dummy and race interaction terms fails to reject the null hypothesis that all the corresponding regression coefficients are zero. We interpret this result as indicating that there is nothing in our data which reveals a clear violation of the exclusion restriction. Of independent interest, Appendix D demonstrates a strong relationship between one's estimated type q and other factors that are believed to be important predictors of educational and life success, such as parents' income and

³¹In order for our identification strategy to break down, the racial animus would have to exactly offset the (non-linear) incentive effects of affirmative action (i.e., Φ). Since the AA markup function, Φ —which we recover without imposing equilibrium assumptions—is non-linear, a college-specific racial-animus penalty would be required to offset the AA incentive effects. If the effect of racial animus on household income was known ex ante and did not exactly offset the AA incentive effects, then our identification result would continue to hold. Appendix H.5 describes the result of estimating a model wherein we impose a uniform 8.3% animus-based household income penalty on minority students based on the estimates of the black-white wage gap studied by Fryer et al. [34, Table 4, column 5] before controlling for job tenure. Our estimates of the household income equation parameters did not qualitatively change, which suggests our model and identification strategy are robust to uniform deviations from our exclusion restriction.

education. This is true despite the fact that no such relationship is explicitly imposed on the empirical model *ex ante*.

At the end of the day, our estimator will rely on differences in the mapping between high-school effort and college placement across demographic groups in order to disentangle the influences of s and q in the income equation. As a final check on whether the estimates we get by relying on cross-race variation are credible, Appendix H.3 estimates a model where we divide the country into regional college markets. Because of inter-regional differences in the incentives facing nonminority students, we can identify the household income equation from nonminority outcomes alone if the marginal effects of (p, s, q) on income are identical across regions. In this alternative model specification we require no assumption on whether or not minority status directly influences household income. Our resulting parameter estimates are essentially unchanged, which suggests that the role we estimate for q in determining household income is not driven by a violation of Assumption 4.4.

Assumption 4.5. (Unique Investment) $U(P_j(s), s, q) - qc(s)$ is strictly concave in s .

Assumption 4.6. (Normal Shocks) Matching shocks $\epsilon_{ij} \sim N(0, s_{ij}^2)$ are normally distributed with zero mean and variance s_{ij}^2 and are independent of HC s and demographic status $\{M, N\}$.

Assumption 4.5 ensures that the agent's decision problem (equation 4) has a unique solution, a needed property for mapping s into a corresponding q . Note that this assumption is testable: the model requires strict supermodularity of the student's HC choice problem (see Appendix A), which implies the investment strategy is strictly monotone, so there would be jumps in the student's HC accumulation strategy if Assumption 4.5 failed. Assumption 4.6 assumes normal matching shocks with 0 mean. Since the NHC index is scale-free, the assumption of a 0 mean is without loss of generality, and the choice of a normal distribution was made solely for tractability.

The structural objects to identify are the type distributions, $F_j(q)$; the matching shock variance, s_{ij}^2 ; the assignment functions, P_j^{ap} ; and the match utility parameters, (b^r, a) . The single index parameters b^s and the joint distributions of (P, S) across race groups are intermediate model components to be identified along the way.

4.1.1. Identification: Single Index Parameters and Graduation Probabilities For simplicity of discussion, assume at first that the single index parameters b^s are known. The first hurdle to overcome is a problem of sample selection: because the B&B survey only contains information for college graduates, we do not observe (P, S) pairs for anyone who failed to graduate. Thus, at first we can only treat the school quality and HC distributions, conditional on graduation, $f_{PS}(p, sj | M, grad)$ and $f_{PS}(p, sj | N, grad)$, as observables. Given the graduation probability function $r(p, s)$, from Bayes' law we know that $f_{PS}(p, sj | M, grad)$ and $f_{PS}(p, sj | N, grad)$ relate to the unconditional densities as follows:

$$(6) \quad f_{PS}(p, sj) = \frac{f_{PS}(p, sj, grad)L_j}{r(p, s)}, \quad j \in \{M, N\},$$

where L_j is a constant that equals the total fraction of enrollees from group j who graduate college and normalizes the joint density to integrate to one. $r(p, s)$ is pinned down by the graduation rate of each demographic group at each college. The model-generated graduation rate at each college is computed by averaging over the graduation rates of the individual students at the respective college, whose achievement is distributed as $f_{sjp}(sjp, j)$. These model-generated race-college graduation rates are then matched to the moments of their observed counterparts, G_j . Formally, we estimate the regression equation

$$(7) \quad G_j = Z_{jl} b^r + e_{jl},$$

where $Z_{jl} = [1, p_l, p_l^2, p_l^3, \bar{S}_{jl}, \bar{S}_{jl}^2, \bar{S}_{jl}^3, p_l \bar{S}_{jl}, p_l \bar{S}_{jl}^2, p_l^2 \bar{S}_{jl}]$ contains the regressors for group j at school l , e_{jl} is random and arises from finite sampling within campus l ,

$$(8) \quad \bar{S}_{jl}^k = \int_{\underline{s}}^{\bar{s}} s^k f_{sjp}(sjj, p_l) ds = \int_{\underline{s}}^{\bar{s}} s^k \frac{f_{PS}(p_l, sjj)}{f_{P_j}(p_l)} ds$$

is the conditional expectation, across both graduates and non-graduates, of the k^{th} power of s given p_l , and

$$(9) \quad f_{P_j}(p_l) = \int_{\underline{s}}^{\bar{s}} f_{PS}(p_l, sjj) ds, \quad j \in \{M, N, G\}$$

is the unconditional marginal distribution of P for group j . Equations (6) – (9) provide a sample selection correction to identify graduation probability parameters b^r as long as the single index parameters b^s are known. One appealing characteristic of our selection correction is that it does not require parametric restrictions on the form of the unconditional joint distributions of (P, S) .

Now we require a further condition to pin down the single index parameters. Our model predicts that, subject to the limits imposed by the matching shock, the outcome will be assortative within each demographic group. Since the mapping between P and S arises from a rank-order contest, we adopt the Kendall's τ measure of rank correlation to formalize our notion of assortativity. Maximizing Kendall's τ is equivalent to choosing b^S to maximize the τ between our theory and the data.³² Recall that the HC index S represents all observable information about the student prior to the application process that predicts where he/she will place. Maximizing assortativity (i.e., the predictive power of the model) can be thought of as equivalent to college admissions officers preferring to enroll high-HC students, and using all of the information at their disposal to do so. With the joint distribution (P, E, A) known, we formalize the notion that college admissions officers, in their role as the gatekeepers to college quality P , extract maximal information content from (E, A) within a contest setting by assuming that b^S is the single index

³²Assuming that the matching shock is orthogonal to HC as per our model, the exercise of maximizing assortativity in the context of our contest is analogous to minimizing the sum of squared errors in an OLS model.

parameter vector that maximizes the Kendall's τ rank correlation between P and S :

$$(10) \quad b^s = \arg \max_{(b^s)} \left[m \Pr (P_1 - P_2)(S_1 - S_2) > 0 \mid M \quad \Pr (P_1 - P_2)(S_1 - S_2) < 0 \mid M \right. \\ \left. + (1 - m) \Pr (P_1 - P_2)(S_1 - S_2) > 0 \mid N \quad \Pr (P_1 - P_2)(S_1 - S_2) < 0 \mid N \right].$$

From the above arguments, the first part of our identification result follows:

Proposition 4.7. There is a unique configuration of the single index and graduation probability parameters (b^s, b^r) that is consistent with the joint distributions of the observables $\{Y_{j=1}^L, f_{p_i, e_i}, a_{g_{i=1}^L}\}$ and equations (6), (7) and (10).

4.1.2. Identification: Matching Shock Variance At this point several important equilibrium objects can be treated as known, including the unconditional joint distribution of HC and college quality. Identifying the matching shock variance parameter $s_{\#}$ is simple since it uniquely determines the degree to which the joint distribution of P and S deviates from full rank correlation within each race group. Intuitively, the larger is the variance of the matching shock, the more latitude there is for students with lower HC levels to place above students with more HC. Thus, it is easy to see that Kendall's τ within each race group is decreasing in $s_{\#}$.

For each group $j \in \{M, N\}$, let $t_{PS}(s_{\#} \mid j)$ denote the rank correlation between HC and college quality implied by shock variance parameter $s_{\#}$ holding $G_j(s)$ and $F_{P_j}(p)$ fixed. Since school assignment within each group is determined by the rank ordering of perturbed HC levels and the perturbations are independent, the following must be true: (i) $t_{PS}(0 \mid M) = t_{PS}(0 \mid N) = 1$; (ii) $t_{PS}^0(s_{\#} \mid j) < 0, j \in \{M, N\}$; and (iii) $\lim_{s_{\#} \rightarrow \infty} t_{PS}(s_{\#} \mid M) = \lim_{s_{\#} \rightarrow \infty} t_{PS}(s_{\#} \mid N) = 0$. The following result directly follows from these facts:

Proposition 4.8. There is a unique value $s_{\#}^2$ that is consistent with perturbed, rank-order allocations and the joint distributions $F_{PS}(p, s \mid M)$ and $F_{PS}(p, s \mid N)$.

One might worry that our estimate of $s_{\#}^2$ is biased downwards by the fact that $\tilde{S}(E, A)$ was chosen to maximize assortativity. To better understand the effects of under-estimating $s_{\#}^2$, in Appendix H.6 we re-estimated our model with $s_{\#}^2$ set to up to 16 times the estimate produced in our data. We find that our household income parameter estimates are essentially unchanged.

4.1.3. Identification: Utility Parameters and Cost Types An approach for estimating strategic models with private information was proposed by Guerre, Perrigne, and Vuong [36] for first-price auctions. Their idea was simple but powerful: since the equilibrium distributions of bids are observable, one can reverse engineer a bidder's private valuation as that which rationalizes her bid as a best response to competitors' bids. Our setting is similar in that each student's investment choice is a best response to the distribution of HC choices given her type. The first-order

condition for the student's problem is:

$$(11) \quad q = \frac{E_{\#} U_p P_j(s + \#, s, q; a) + E_{\#} U_s P_j(s + \#, s, q; a)}{c^{\alpha}(s)}, \quad j \in \{M, N\}, g,$$

where $a = (\log(a_0), a_p, a_s, a_q)$ is the vector of utility parameters governing household income production. Due to the strict supermodularity of our model (Appendix A), the equilibrium strategies must be strictly monotone, which in turn implies equation (11) uniquely defines the inverse strategy mapping, $q_j(s; a)$ if the utility parameter vector a is known. This in turn implies that the type distributions $F_j(q) = G_j(q_j^{-1}(q; a))$, $j \in \{M, N\}, g$, are known if a is known.

Our estimator will infer a student to have a high (low) value of q if the student chooses to accumulate a low (high) level of HC given the effect on future household income. Any factors that change the perceived marginal benefit/cost ratio of HC will be encapsulated in estimated \hat{q} . For example, high discount factors that depress the discounted value of future household income would generate a high \hat{q} . One concern is that there may be factors that drive high-school achievement but have no impact on household income, which could bias \hat{q} . If the high discount factor reduces the marginal benefit of accruing HC for high-school students, then it is plausible that it also reduces the marginal benefit of accruing college-level HC that generates post-college household income. Other effects, such as patience during a job search, also point towards an household income effect. This suggests a broad interpretation of q including more than just learning costs, rather than a problem of bias.

One might be concerned that low ability students may not understand the decision problem they are facing, either because of a lack of information or an inability to fully optimize. To address this, in Appendix H.7 we recompute our household income parameter estimates after dropping the students that occupy up to the bottom 50% of the quality distribution of college seats and the students that occupy them, which are taken by the lowest HC (highest q) students. If our results are driven by a failure of these low achieving students to optimize effectively, then one would expect our estimates of the household income equation parameters to change. To the contrary, we find that the household income parameters are insignificantly changed from those we find using the full sample.³³

Our last moment condition comes from the household income regression model:

$$(12) \quad \log(w_i) = \log(a_0) + a_p \log(p_i) + a_s \log(s_i) + a_q y(s_i, D_{Mi}; a) + e_{wi},$$

where $y(s_i, D_{Mi}; a) = \log[D_{Mi} q_M(s_i; a) + (1 - D_{Mi}) q_N(s_i; a)]$ is i 's log-learning-cost type and e_{wi} is a transitory shock to 10-year household income. We assume transitory shocks are exogenous.

Assumption 4.9. $E[\log(p_i), \log(s_i), y(s_i, D_{Mi}; a)]^{\#_{wi}} = 0$.

³³Cotton, et. al. [25] study a real-effort field experiment and found evidence that students in grades 5 through 8 responded optimally to changes to a contest structure. This suggests that older students are likely capable of sophisticated strategic behavior as actors within the college admissions market.

Essentially, $y(s_i, D_M; a)$, which is derived from our theory of investment behavior, serves as a nonlinear control function and allows the researcher to separate out the influence of unobserved student characteristics q from achievement s and school quality p . The intuition behind structural identification is as follows: in order for the parameters of the household income regression to be identified, we need an orthogonality condition and a full rank condition. Assumption 4.9 establishes orthogonality based on the idea that the same unobserved characteristics that govern a student's pre-college achievement also govern the HC accumulation process during college.

For the full-rank condition, there must be something present in the data-generating process that prevents the regressors from being perfectly colinear. Note that $s_{\#} > 0$ implies a non-degenerate distribution of HC types on each college campus under any college admissions rule, so that rank correlation between $\log(s)$ and $\log(p)$ must be less than one in absolute value. Second, denote the expected household income of minority (nonminority) students graduating from college p with HC level s as $U_M(p, s)$ ($U_N(p, s)$). Suppose we observe a positive measure of (p, s) such that $U_M(p, s) = U(p, s, q_M(s; a)) \neq U_N(p, s) = U(p, s, q_N(s; a))$. Our exclusion restriction implies that it must be the case that $q_M(s; a) \neq q_N(s; a)$, which insures that $\log(s)$ and $\log(q)$ are not perfectly correlated. Thus, in expectation the matrix of regressors

$$(13) \quad X(a) = \begin{bmatrix} 1 & \log(p_1) & \log(s_1) & y(s_1, D_{M1}; a) \\ 1 & \log(p_2) & \log(s_2) & y(s_2, D_{M2}; a) \\ \vdots & \ddots & \ddots & \vdots \\ 1 & \log(p_l) & \log(s_l) & y(s_l, D_{Ml}; a) \end{bmatrix}$$

has full rank for any value of a . This logic yields our final result on structural identification:

Proposition 4.10. Under Assumptions 4.1 – 4.9, household income parameters and cost type distributions $F_M(q)$, $F_N(q)$ are identified, provided that

- (i) $0 < m < 1$
- (ii) $s_{\#} > 0$,
- (iii) $\exists(p, s)$ such that $U_M(p, s) \neq U_N(p, s)$
- (iv) $r_a^2 Q(a)$ is positive definite on $(0, \infty) \times (0, 1)^3$, where

$$Q(a) = [D_w(W'X(a)a)]' [D_w(W'X(a)a)],$$

$W = [\log(w_1), \log(w_2), \dots, \log(w_l)]'$ is the vector of corresponding observed household incomes, and $D_w = [D_{w1}, D_{w2}, \dots, D_{wl}]$ is a vector of sampling weights.

4.2. A Two-Stage, Semiparametric Estimator. We now construct a two-stage GMM estimator to implement our identification strategy. First, we recover the preliminary model parameters that do not directly depend on our strategic investment model, b^s , b^r , $P_M(s)$, $P_N(s)$, and $s_{\#}$. Then we use these estimated values and the first-order conditions to recover the utility parameters a and the learning cost distributions, $F_M(q)$ and $F_N(q)$. The technical details of the stage I and II estimators are in Appendices B.1 and B.2 respectively.

4.2.1. Stage I Estimation. The first hurdle to overcome is to find a computationally tractable way of representing the joint distribution of (P, S) conditional on graduation. Multidimensional density estimation is a difficult problem both computationally and because of the rapid decay of optimal statistical convergence rates as the dimensionality of the underlying random variable grows.³⁴ Recent work by Hubbard, Li, and Paarsch [42] has employed parametric copula functions to solve this problem. We follow this dimension-reduction strategy by adopting the Gumbel-Hougaard copula to represent the correlation structure in the joint distribution between HC's and college quality p . For the marginal distributions of HC and college quality, we use a flexible approach based on B-splines. B-splines are a class of finite-dimensional functions that can be made arbitrarily flexible, and they are better behaved numerically than global polynomials (e.g., Chebyshev).³⁵

With the above functional representations of the marginal and joint distributions in equations (6) – (9) and (3), it is straightforward to understand how one could build a GMM-style estimator for all stage I parameters. The single index parameters, b^s , are chosen to maximize the rank-predictability (as measured by the Kendall's τ rank correlation) of HC's for the quality of a student's college placement p , given the empirical joint distribution of (E, A, P) .³⁶ The joint moments of graduation rates, HC, and college quality across different colleges pin down the graduation rate parameters b^r . The empirical joint distribution of P and S for graduates is used to pin down the B-spline-copula representation of the same distribution, and Bayes' Rule in conjunction with the graduation probability parameters is used to recover the unconditional joint distribution of (P, S) for enrollees. Finally, the matching shock parameter $s_{\#}$ is chosen to match the empirical rank correlation in the distribution of (P, S) (for enrollees) as closely as possible.

While it is straightforward to establish intuitive connections between the moments in the data and stage I model parameters, a formal definition of the estimator is notationally intense since most of these terms must be estimated simultaneously. Therefore, we leave a formal treatment of the GMM stage I estimator to Appendix B.1. In what follows, we represent the stage I parameters by $\hat{\rho} = \langle \hat{g}_M^p, \hat{g}_N^p, \hat{g}_M^q, \hat{g}_N^q, \hat{\eta}_M, \hat{\eta}_N, \hat{b}^s, \hat{b}^r \rangle$ and $\hat{s}_{\#}$, which includes estimates for the B-spline parameters for the race-specific marginal distributions of P conditional on graduation, (g_M^p, g_N^p) ; the B-spline parameters for the race-specific marginal quantile functions of S conditional on graduation, (g_M^q, g_N^q) ; the (unconditional) race-specific copula parameters (η_M, η_N) ; the single-index parameters b^s ; the graduation probability parameters b^r ; and the matching shock parameter $s_{\#}$.

4.2.2. Stage II Estimation. Our stage I estimator was based on intuitive moment conditions that were notationally intense to formalize. Stage II estimation is the reverse: notationally compact with considerable computational complexity under the surface. In this stage we simultaneously estimate the parameters of the household income equation and an inverse strategy that maps HC

³⁴See Silverman [59] and Campo, Perrigne, and Vuong [18] for a lengthy discussion on this concept.

³⁵For a brief primer on B-splines and their advantages in empirical auctions models, see Hickman, Hubbard, and Paarsch [38] and Bodoh-Creed, Boehnke, and Hickman [14].

³⁶Recall $\#$ is independent of pre-college HC and school quality. See footnote 32.

Table 3. ESTIMATES: Single Index Function and Matching Shock Variance

Parameter	Point Est.	Std. Err.	P-Value	95% Confidence Interval
$b_1^s(e)$	0.0506	(9.63 10^{-4})	< 0.001	[0.0485, 0.0524]
$b_2^s(e^2)$	0.1567	(0.0099)	0.006	[0.1399, 0.1662]
$b_3^s(a)$	0.3039	(0.0089)	< 0.001	[0.2912, 0.3208]
$b_4^s(a^2)$	0.2898	(0.0096)	< 0.001	[0.2738, 0.3067]
$b_5^s(e \cdot a)$	0.2296	(0.0090)	< 0.001	[0.2114, 0.2421]
$s_{\#}$	0.0275	(7.25 10^{-4})	—	[0.0260, 0.0289]
$t_{PS}(s_{\#} M)$	0.8723	(0.0072)	—	[0.8635, 0.8929]
$t_{PS}(s_{\#} N)$	0.8952	(0.0034)	—	[0.8893, 0.9049]

NOTE: Significance at the 10%, 5%, and 1% levels is denoted by *, **, and ***, respectively. e denotes normalized standardized exam scores, a denotes normalized nonmajor GPA.

In order to explore the role of sampling variability, we employ a block-bootstrap procedure which involves re-sampling with replacement 1000 times from the race-specific B&B subsamples. We separately re-sample I_M student observations from the minority subsample $\{w_i, p_i, e_i, a_i, g_{i=1}^M\}$ and I_N student observations from the non-minority subsample $\{w_i, p_i, e_i, a_i, g_{i=1}^N\}$. Because college-level variables represent the universe of four-year colleges and the students enrolled in these colleges, we hold these observables fixed during our bootstrap procedure. Moreover, we calibrate the minority mass m from the IPEDS college-level data and hold it fixed as well. Thus, all variation in our standard errors come from the finite sampling of the student-level B&B data. We re-estimate all model parameters on each bootstrap sample, and confidence intervals are then biased-corrected and accelerated in the standard way.

5. ESTIMATION RESULTS

We now discuss our model estimates. Goodness of fit metrics are provided in Appendix C.

5.1. ESTIMATES: Single Index Function $\tilde{S}(e, a)$ and Matching Shock Variance $s_{\#}^2$. The parameter estimates for the HC single index equation are contained in Table 3.³⁷ The HC index is convex in both e and a and admits significant complementarities between the arguments. The marginal effect of a one standard deviation change of e from the median is 0.0668, while the marginal effect of a one standard deviation change in a from the median is equal to 0.1520. In other words, a is 2.5 times as important as e for determining a median student's HC single index. Figure 8 in Appendix F contains a plot of the single index function with confidence bounds.

Table 3 includes estimates of the matching shock standard deviation, $s_{\#}$, and group-specific Kendall's τ at enrollment to provide different perspectives on the matching frictions. The noise-to-signal ratio of the matching shock (the ratio of $s_{\#}$ to a standard deviation of HC) is 18.2%. Since

³⁷Recall that the single index function is scale-free and normalized to attain a maximum value of one. Recall also that a more flexible cubic form did not improve our model fit (see footnote 29).

Table 4. Graduation Probability Estimates $b(p, s)$

Parameter	Point Est.	Std. Err.	P-Value	95% Confidence Interval
$b_1^r(p)$	0.0471	(0.0805)	0.4530	[-0.0901, 0.2283]
$b_2^r(p^2)$	0.1991	(0.1391)	0.0990	[-0.443, 0.0755]
$b_3^r(p^3)$	0.2160	(0.1305)	0.0550	[-0.0070, 0.4891]
$b_4^r(s)$	0.8447	(0.1658)	< 0.001	[0.6102, 1.3090]
$b_5^r(s^2)$	0.0179	(0.1421)	0.9640	[-0.0651, 0.0108]
$b_6^r(s^3)$	0.0029	(0.0129)	0.3020	[-0.0254, 0.0249]
$b_7^r(p \cdot s)$	0.0464	(0.3753)	0.9040	[-0.4627, 1.1309]
$b_8^r(p \cdot s^2)$	0.0093	(0.1443)	0.8450	[-0.1216, 0.0309]
$b_9^r(p^2 \cdot s)$	0.0975	(0.2739)	0.4800	[-0.3067, 0.9115]
$b_0^r(\text{const.})$	0.0532	(0.0236)	0.065	[-0.0844, -0.0086]

NOTE: Significance at the 10%, 5%, and 1% levels is denoted by *, **, and ***, respectively. p denotes college quality, s denotes pre-college human capital.

Kendall's τ for minority students is estimated at 0.872, this means for two randomly selected minority students there is a 93.6% chance that the student with higher HC will enroll in a higher quality college. This probability for the nonminority group is 94.8%. Thus, while matching shocks play a role, our empirical model suggests a high degree of assortativity in the market.

If students are motivated by concerns other than quality or make mistakes in the college application process, then this would be reflected in a large value $s_{\#}$. For example, if students were willing to make trade-offs between the geography or demography of a school and college quality, then the assortativity of our match would be lowered (i.e., $s_{\#}$ would be large). If students were unable to identify the highest quality of college that they could be admitted to because of (for example) information frictions, then the match would be less assortative. The fact that $s_{\#}$ is low leads us to conclude that, while these issues may play a role, they are not as powerful as the incentives our model captures. We also test whether assortativity fails for low HC and minority students in Appendix H.9 and find no significant effects, which lends confidence in our model.

5.2. ESTIMATES: Graduation Probability Function, $r(p, s)$. Point estimates and standard errors for the graduation probability parameters are displayed in Table 4. We compute the marginal effect of a one standard deviation change in p and s (from the median values) on $r(p, s)$ as our metric of the relative importance of these variables for determining the graduation probability. In this case, we compute the median and standard deviation statistics using the selection-corrected distributions of the respective variables. The marginal effect of a one standard deviation change in p is 0.026, while the marginal effect of a one standard deviation change in s is 0.143. This means that s is roughly 5.5 times as important as p for determining college graduation probabilities. Appendix F contains an illustrative plot of $r(p, s)$ with confidence bounds.

(a) HC Distribution (b) School Quality Distribution

Figure 1. HC and School Quality Distributions

5.3. ESTIMATES: Selection-Corrected Joint Distributions $f_j(p, s)$. Figure 1a displays the distributions of HC levels for each demographic group, including the selected sample of graduates from the raw data (dashed lines) and the selection-corrected distributions for all enrollees (solid lines). Figure 1b displays the distributions of college seats allocated to each group, including the selected raw samples (dashed lines) and the selection corrected distributions for all enrollees (solid lines). The CDF plots also include 95% confidence bounds at the deciles of the population-wide distributions. The first plot illustrates the achievement gap, a stochastic dominance relationship between minority HC and non-minority HC. The second plot illustrates the enrollment gap, a similar stochastic dominance relationship between minority and non-minority college quality.

5.4. ESTIMATES: Minority Markup Function, Ψ . Figure 2 describes the estimated markup function, $\Psi(t)$.³⁸ The horizontal axes of both panels display quantile ranks of NHC for nonminority students. The left panel describes the shape of Ψ . The vertical axis displays the quantile rank of subsidized NHC (within the non-minority NHC distribution) with 95% confidence bounds at the deciles of NHC. If a minority student has an NHC at the quantile rank marked on the horizontal axis, the student gets the same college assignment as a nonminority student with an NHC at the quantile rank denoted on the vertical axis. For example, a minority student with an NHC equal to the median of the nonminority population gets the same college assignment as a nonminority student at the 64th percentile of the nonminority population. The dashed line denotes the 45 line for reference.

The right panel of Figure 2 describes the markup function in terms of school quality. The vertical axis denotes the gap in the quantile rank of college quality between a minority student and a nonminority student at each NHC quantile with 95% confidence bounds at the deciles of NHC. For example, the plot shows that if two students from different groups both have an NHC value equal to the median of the nonminority population, then the minority student is assigned to a school whose quantile rank is 0.13 higher in the school quality distribution.

³⁸Figure 10 in Appendix F.3 describes the markup function that is equivalent to a proportional quota when used in an admissions preference system.

Figure 2. Noisy Human Capital Markup Function, μ

Our plot reveals that the effect of the status quo admissions preference scheme is insignificant at colleges in the bottom decile, but is statistically and economically significant across the rest of the college quality spectrum. Several papers have estimated a substantial impact of AA, including Bowen and Bok [16]; Chung and Espenshade [20]; and Chung, Espenshade, and Walling [21], but these studies used data from elite colleges, whereas ours provides a market-wide picture.

The most similar previous study is Kane [43], which also used a nationally representative sample (the High School and Beyond (HS&B) survey), but estimated a significant role for AA only in the top quintile of the market. Several differences exist between Kane [43] and our study. First, we use measures of national market allocations (enrollment data), whereas Kane [43] uses applications data which may not fully reflect national enrollment decisions. Second, the HS&B data contain potentially important sources of sample selection that could affect probit regression results in unpredictable ways. HS&B respondents were asked their two top choices (sample truncation) among the schools to which they applied (endogenous selection in student-school pairs) and whether they were accepted.

At the end of the day, our estimates of the markup function are the most directly data-driven component of the empirical model: they hinge on a stage I reduced-form sample-selection correction to map our observed set of college graduates into the original set of college enrollees. The intuition behind our result is that a color-blind world implies a very specific form for the distribution of (p, s) . Our stage I reduced-form data products deviate from this form, and in such a way that more generous admissions practices toward minorities must exist on the majority of the market in order to rationalize observed allocations from observed achievement.³⁹

5.5. ESTIMATES: Match Utility and Learning Cost Type Distributions. The type distributions are presented in Figure 3. The left panel displays the type CDFs $F_M(q)$ and $F_N(q)$ with 95% confidence bounds represented by the shaded areas. The type distribution of the minority students stochastically dominates the type distribution of nonminority students, which is consistent with minority students having higher learning costs than nonminority students. The right panel

³⁹Dillon and Smith [28] also studied over/under placement in college admissions. Methodological differences make direct comparisons hard, though their results are qualitatively similar (see Appendix G).

Figure 3. Agent Type Distribution

Table 5. Estimates of the Household Income Equation

Parameter	Point Est.	Std. Err.	95% Confidence Interval
a_p (College Quality)	0.1349	(0.0260)	[0.0894, 0.1902]
a_s (Pre-College Human Capital)	$1.181 \cdot 10^{-6}$	(0.0306)	[$1.000 \cdot 10^{-5}$, $2.412 \cdot 10^{-4}$]
a_q (Learning Cost)	0.0574	(0.0386)	[0.009, 0.1478]
a_0 (Intercept)	\$79,108	(\$4,735)	[\$69,105, \$87,499]

NOTE: Significance at the 10%, 5%, and 1% levels is denoted by *, **, and ***, respectively.

of Figure 3 plots the pointwise difference between the CDFs with 95% confidence bounds, and the difference between the type distribution is significant at all learning-cost quantiles.

This result conforms with a large body of empirical evidence on stark racial differences in access to resources that affect childhood development in the US. For example, Black and Hispanic children are nearly three times as likely to live below the poverty line (see Kena et. al. [45]). They are also much less likely to be covered by health insurance (see Smith and Medalia [60]) or to be raised by parents with bachelor degrees (see Fox, Kewal, and Ramani [33]). Holding income level fixed, Blacks and Hispanics are also much more likely to attend under-performing schools that serve poorer student bodies relative to their White and Asian counterparts of similar incomes (see Reardon, Townsend, and Fox [58], Reardon [56], and Reardon, Kalogrides, and Shores [57]). In our data we find direct evidence consistent with this view as well. In Appendix D we show that various factors relating to childhood family background—e.g., parents' education and wealth among others—are highly predictive of the learning cost types we estimate. With these empirical facts in mind, the estimated stochastic dominance relationship in learning costs would seem a natural, though unfortunate, consequence of resource stratification by race.

Table 5 presents our estimates of the returns to the three inputs of household income production. The first takeaway is that both school quality and the agent's unobservable type have economically and statistically significant effects on household income production. Though both

are important, college quality plays a dominant role: the marginal effect of a one standard deviation increase in $\log(p)$ is 3.44 times larger than the marginal effect of a one standard deviation reduction in $\log(q)$. The second takeaway is that pre-college HC (i.e., s), has little direct effect on household income, conditional on graduation. However, it does drive income indirectly because it causes the student to enroll in a higher quality college and graduate with higher probability (holding college quality fixed). A one standard deviation increase in HC from the median increases income conditional on graduation by 5.9% through its effect on college placement, and, holding placement fixed at the median college, the same change in HC increases expected income (unconditional of graduation) by 5.1% through an improved graduation rate.

It is also important to stress that our results do not imply that all forms of human capital are only indirectly relevant to post-college income. Rather, it only implies that the general human capital a student acquires during high school and early childhood are not directly generating post-college income. There are several likely interpretations for the importance of college in our wage equation estimates. First, it could be that more selective colleges offer higher wages to faculty and can attract instructors who are more effective at developing college-level human capital in their students in a specific field of expertise (e.g., major coursework in engineering, business, journalism, physics, etc.). It is also possible that, holding educational quality fixed, elite colleges offer better career placement services or access to more advantageous referral networks when their students graduate. In reality, the explanation is likely a mixture of these and perhaps other factors, and is left to future research. However, our results strongly suggest that that phenomena which occur during college—including post-secondary human capital accumulation—seem to be the primary, direct drivers of post-graduate income.

In general, the empirical evidence literature on the returns to college quality finds that more selective colleges benefit poorer students significantly, but there is some disagreement as to the magnitude of the return for more affluent students. Dale and Krueger [27] find negligible returns to a more selective college for affluent students, but Brewer, Eide, and Ehrenberg [17]; Black and Smith [13]; Long [46]; and Andrews, Li, and Lovenheim [3] estimate significant returns on the investment in a more selective college for all students. Since a_p is the largest of all the production weights, our results support the latter view that a higher quality college is a resource from which all students benefit.⁴⁰ Holding s and q fixed at their medians, a move from the 25th percentile college to the 75th percentile college induces an estimated shift of \$7,500 in annual household income. We are the first paper to directly include unobserved heterogeneity as an independent variable (separate from the error term) driving the returns to an education, which we find is significant: reducing the learning cost from the upper quartile of q to the lower quartile, holding p fixed at the median, induces an increase of roughly \$2,100 in annual household income.

⁴⁰Although our main model specification is monotone in college quality p by construction, we explore robustness checks in an online appendix which serve to probe for whether the B&B data show evidence of the returns to college being weaker for more affluent students. Appendix H.3 estimates a version of the model using only non-minority students and basing identification on regional variation in contest structure, rather than on variation due to AA. Excluding minorities systematically selected out of the sample many of the poorest students, but our point estimate on the returns to college quality is qualitatively and statistically unchanged.

While a student's pre-college HC choices do not significantly affect household income conditional on graduation, these investments still have an influence on the probability of graduating college. This is consistent with the view that a student's learning in middle school and high school helps her pass the curriculum of her college, but post-college household income is determined by the level of HC accrued during college, which depends primarily on college quality and learning cost.

5.6. Robustness Check Summary. We now summarize the main robustness checks mentioned throughout the paper, which are fully discussed in Appendices G and H.

First, we extended our model to include regional higher education markets, defined by the colleges in each U.S. Census Bureau region, in order to investigate whether our stage II household income regression results change. This expanded model could reflect inter-regional differences in the model primitives—e.g., regional differences in assignment mappings if geography (or students' geographic preferences) plays a role in access to college quality. Under this model, we can identify α from the differing incentives facing nonminority students in different regions in lieu of using exclusion restriction 4.4. In both cases—when we estimated the regional model using both racial and regional variation in incentives, and when we re-estimated the model using only regional variation on the sub-sample of nonminorities—our resulting household income parameter estimates were nearly the same as in the baseline, national market model.

Second, Appendix H.4 conducts a sensitivity analysis to explore how different assumptions about the relative magnitudes of s_h^2 (unobserved HC) and $s_\#^2$ (matching shock) would affect our estimates. (Although we can estimate the sum $s_h^2 + s_\#^2$ from our data, there is no variation that allows us to estimate s_h^2 and $s_\#^2$ separately.) We find that our estimates of the household income equation parameters (equation (12)) are essentially unchanged from our benchmark estimates if s_h^2 accounts for less than half of the estimated value of $s_h^2 + s_\#^2$. For $s_h^2 \geq s_\#^2$, there is insufficient matching shock variation remaining to separately identify the effects of p and s on household income. However, the relative magnitudes of the effect of observables (p and s) on household income relative the effect of unobserved characteristics (q) is stable. We also test whether $s_h^2 > 0$ by adding additional variables into the single index equation \tilde{S} that are relevant to mathematics and English preparation during high school. We find that, conditional on observed GPA and exam scores, this additional information plays a trivial role in the estimated single index and does not increase predictive power of HC for college placement. Recall that our point estimate $\hat{s}_\# = 0.0275$ (Table 3) is small, so there is not much college assignment uncertainty left for these variables to explain.

Third, we also re-estimated our model using standardized exam scores alone to measure HC. The model fit poorly and implied an implausibly low degree of assortativity between p and s .

Fourth, to test whether a failure of low achieving students to optimize, due to information frictions or to a lack of ability, influences our results, we repeated stage II using only the students at the highest quality colleges. Our household income parameter estimates were essentially unchanged.

Fifth, for a comparison with papers that focus on male students, we also repeated our stage II analysis on our male subsample. We found that the household income parameters were essentially unchanged.

Sixth, to address concerns that our estimator depresses the estimate of $s_{\#}$ by specifying assortativity maximization as our empirical criterion in stage I and that this may be driving our stage II results, we repeated our analysis with $s_{\#}$ fixed by up to 16 times our estimated value. We found that our parameter estimates didn't change.

Seventh, we checked for sensitivity to deviations from our exclusion restriction 4.4 by imposing a hypothetical 8.3% animus-based household income penalty on minority students and re-estimating our stage II household income equation.⁴¹ We found that our parameter estimates didn't change in a meaningful way. As a further probe of our exclusion restriction, if it were violated one might expect racial animus to influence our estimates of q . In Appendix D we regress estimated learning costs on demographics and childhood home characteristics and find that race plays no role, conditional on factors such as parental education and income.

Eighth, we conducted a variety of robustness checks of our functional form assumptions including different specifications of the cost function; allowing $s_{\#}$ to vary by race group; allowing for heteroskedastic matching shocks (i.e., letting $s_{\#}$ be a linear function of s); allowing estimated matching shocks to enter the stage II household income equation as a fourth productive factor; including quadratic terms in p and s in the household income function; and allowing the HC single index function $\tilde{S}(E, A)$ to include cubic interactions of e and a . None of these changes qualitatively altered our results.

Finally, in Appendix G we estimate an alternative model that allows for explicit mismatch, which is the idea that some students may be worse off at a higher quality college. While our baseline model allows for some forms of mismatch (discussed in the next section), we estimate an alternative specification of the income equation with an additional term that measures the effect of being far behind the mean achievement level of one's classmates in college. We find no evidence of mismatch in this sense, and our model estimates remain essentially unchanged.

6. COUNTERFACTUAL POLICY EXPERIMENTS

In this section we explore the economic implications of the college admissions contest. First we consider the amount of HC accumulation that is “wasted” by the rat-race for admissions to a high quality college. We then explore the effect of changes to the AA system. For each admissions scheme (color-blind and proportional quota), we solved for the equilibrium holding our structural point estimates fixed (for details see Appendix E), which reveals how HC accumulation strategies change under each regime. We also discuss the extent to which controlling for the changing HC incentives is necessary to predict the effects of AA. Note that we hold the set of college applicants fixed throughout these analyses.

⁴¹The 8.3% quantity was taken from estimates by Fryer, et. al. [34]; (see discussion of Assumption 4.4).

6.1. The Relative Force of the Competitive and Productive Channels of Investment. The incentives underlying the college admission contest and the effects of changing the AA system can be understood by examining the first-order condition. We use the equivalence of all three systems to a quota to simplify the first-order condition, which is

(16)

$$E_{\#} \left[\underbrace{U_p P_j(s_{\#}), s, q; a \frac{h_j(s_{\#})}{q_j(P_j(s_{\#}))}}_{\text{Competitive Channel}} \right] + E_{\#} \left[\underbrace{U_s P_j(s_{\#}), s, q; a}_{\text{Productive Channel}} \right] = \underbrace{c_j^0(s)}_{\text{Marginal Cost}}, \quad j \in M, N, G.$$

h_j is the endogenous density of the NHC distribution for group j , and q_j is the density of seats assigned to group j in the equivalent quota.⁴² The effect of the AA system on the competitive channel of investment incentives is captured by the student-to-seat ratio, $h_j(s_{\#})/q_j(P_j(s_{\#}))$. The fraction of competitors surpassed given a small increase in HC is equal to the change in the enrolled school quality percentile. When $h_j(s_{\#})$ is large, small increases in HC result in the student surpassing a large fraction of her competitors. When $q_j(P_j(s_{\#}))$ is low, small changes in the fraction of competitors surpassed generate large improvements in college placement. Both effects increase the incentive to accrue HC.

The only endogenous component of equation (16) is h_j , which is determined by the primitive distribution of learning-costs ($f_j(q)$) and matching shock variance ($s_{\#}$) as well as the endogenous strategy adopted by the agents. If $f_j(q)$ is high, then (ceteris paribus) $h_j(s_{\#})$ will be large, meaning that the student faces a high degree of competition. If $s_{\#}$ increases, then h_j becomes more “uniform;” that is, h_j becomes lower (higher) in regions where $f_j(q)$ is higher (lower).⁴³ However, h_j eventually falls as $s_{\#} \rightarrow \infty$ and college assignment becomes random. Of course, $s_{\#}$ also represents productive human capital bearing a direct marginal return, as depicted in the second term on the left-hand side of (16). A full discussion of comparative statics related to changes in the contest primitives can be found in Bodoh-Creed and Hickman [15, Appendix C].

Equation (16) decomposes the marginal benefit of HC investments into two components (left-hand side), which are equated to marginal cost (right-hand side) in equilibrium. The competitive channel is the indirect, “Spencerian” incentive to invest in order to obtain a seat at a better college (Spence [61]). Investments undertaken purely due to the competitive channel represent wasteful over-investment (from the perspective of the students). The productive channel can be thought of as the “Beckerian” incentives that represent direct marginal benefits of holding an additional unit of productive HC. In a first-best world of complete information, students would be assortatively assigned to colleges according to their types, and students would then accrue the ideal amount of HC given their assignment (i.e., the amount dictated by the productive channel).

To get a sense for the magnitudes of these forces as a function of the student's type, Figure 4 displays the ratio of the competitive channel to the total marginal benefit. The horizontal axis describes the quantile rank of the cost type in the group-specific type distribution. Whenever the

⁴²The quotas equivalent to the endogenous outcome of each AA system are presented in Table 6.

⁴³Recall that h_j is the convolution of $G_j(s)$ and the distribution of the matching shock $\#$.

Figure 4. Relative Force of Competitive and Productive Channels

line is above 0.5 the competitive channel is dominant, and whenever it is below, the productive channel is most important. The competitive channel is stronger than the productive channel for all but the lowest-cost agents. Most students' academic achievement levels would be much lower in a complete-information world where the competitive channel is turned off.⁴⁴ Of course, this discussion ignores some issues in uencing optimal policy—for example, if HC spillovers are important in the economy at large, the social planner may wish to use any means to maximize HC output—but it casts light on the relative strengths of the incentives driving HC choices.

6.2. Effects of AA on HC Investment. It is theoretically ambiguous whether the type-specific minority HC choices will increase or decrease under a particular AA scheme. Figure 5 presents the change in minority HC investment under our two counterfactual admissions schemes, so positive (negative) values indicate increases (decreases) relative to the status quo.⁴⁵ We describe the student's type in terms of quantiles of the minority cost distribution, and changes in HC are displayed as a fraction of a standard deviation in the status quo level.

The decisions of high and low learning-cost minority students are pushed in opposite directions, so the effect on average HC is relatively small. The HC choices of high and low learning-cost students are pushed farther apart in a color-blind world. The lowest 25% of the cost distribution increase their HC by more than 0.1 standard deviations, but the highest 25% of the learning cost distribution decrease their HC by about 0.05 standard deviations. In contrast, the HC choices of the students are brought together under a proportional quota system. A proportional quota drives down the HC choices of the bottom 10% of the learning cost distribution, while the upper 50% of the learning cost distribution increases their HC by 0.05-0.1 standard deviations.

⁴⁴HC affects utility only via graduation probability, and this is estimated as nearly linear in HC. There are only slight differences between the magnitude of the productive channel in the status quo and under the first-best.

⁴⁵Akhtari and Bau [2] estimate reduced-form behavioral responses to a 2003 court decision that induced a shift from color-blind admissions to race-based AA in 2003 in Texas. Using data from a large urban school district in Texas, they find that the race gap in test scores closed by 0.17 standard deviation units, and the gap in grades closed by 0.07 standard deviation units. Since the policy shift is more localized and their sample contains only urban minority students, their results are not directly comparable to ours. However, their estimates are of similar magnitude as what we see in the upper third of the cost distribution in Figure 5 (non-minority achievement changes little, so the minority change is roughly the amount the gap closes).

Figure 5. Minority Investment Strategy

Table 6. Enrollment by Group and School Quality Quintile

MINORITIES				
College Quality Tier	Status Quo	Proportional Quota	Color-Blind	Color-Blind, Fixed HC
Top College Quintile	0.130	0.200	0.101	0.082
Second College Quintile	0.166	0.200	0.134	0.122
Third College Quintile	0.149	0.200	0.156	0.168
Fourth College Quintile	0.281	0.200	0.191	0.227
Bottom College Quintile	0.275	0.200	0.419	0.402
NON-MINORITIES				
College Quality Tier	Status Quo	Proportional Quota	Color-Blind	Color-Blind, Fixed HC
Top College Quintile	0.213	0.200	0.219	0.223
Second College Quintile	0.206	0.200	0.212	0.214
Third College Quintile	0.210	0.200	0.209	0.206
Fourth College Quintile	0.185	0.200	0.202	0.195
Bottom College Quintile	0.186	0.200	0.159	0.162

Under a color-blind scheme, the lowest-cost minority students face more competitors (i.e., h_M / q_M increases) and the marginal benefit from HC rises, while the highest-cost minority students compete for a glut of seats at low quality colleges (i.e., h_M / q_M drops) and the marginal benefit to HC accumulation drops. Symmetrically, under a proportional quota, the lowest learning-cost students face a lower student-to-seat ratio, which depresses the marginal benefit of HC. The highest learning-cost students face a higher student-to-seat ratio, which increases the marginal benefit of HC. The effect on nonminority students is more limited since the distributions of college seats and competitor types closely resembles the corresponding status quo distributions.

6.3. Effects of AA on Minority Enrollment. Table 6 describes counterfactual enrollment in terms of the fraction of each demographic group enrolled in each college quality quintile. Equal representation of each group in each quintile, 20%, is mechanically achieved by a proportional quota, regardless of the strategies used by students. Numbers below this imply under-representation and vice versa. For example, 13% of minority students enroll in colleges in the top quintile under the status-quo AA scheme, but only 10% enroll in top colleges under a color-blind scheme. While the status-quo AA is less generous to minorities at top colleges than a proportional quota would be, a ban of race-based AA reduces minority enrollment in the top tier by nearly one quarter.

The status quo AA scheme has the intended result in that there is a first-order stochastic dominance (FOSD) shift in the distribution of the quality of colleges in which minority students enroll, relative to the color-blind case. Interestingly, the largest shift is from the lowest quintile of college quality into the second lowest quintile: in a color-blind world, minority enrollment in the bottom quintile would increase by almost one half. A proportional quota yields an even stronger FOSD shift toward minority enrollment in better colleges, relative to the status quo. For completeness, Table 6 also provides the effects of AA on nonminorities. Changing the AA scheme has the opposite effect on nonminority students, but the magnitudes are small since the non-minority mass is five times larger.

The final column of Table 6, denoted Fixed HC, describes the outcome if we compute the enrollment under a color-blind system while holding fixed the status quo HC choices. Comparing the color-blind equilibrium outcomes with the fixed HC results, we see that accounting for endogenous HC yields more (less) diversity at high (low) quality colleges than a fixed HC analysis would predict, which is due to the fact that HC investment is higher (lower) for low (high) learning-cost students in the equilibrium of the color-blind system. Moreover, these differences are large. For example, the naive, fixed-HC counterfactual predicts a 65% bigger drop in enrollment of minority students in the best quintile, while predicting a 35% larger increase in enrollment of minority students in the bottom two quintiles.

6.4. Effects of AA on Household Income, Graduation Probability, and Welfare. In order to give the reader a sense for how AA shapes minority outcomes, we present graduation rate and income changes in two ways. The first separates students by learning-cost quintile, each of which represents the same set of individuals under each counterfactual scenario. We also present the outcomes in terms of college quality quintiles to depict the effect on aggregate, college-level statistics, but note that individuals enrolling in each quality quintile change across counterfactual scenarios. We do not display the effects for nonminority students as they are much smaller.

The top panel of Table 7 displays graduation probability changes by learning cost quintile. Two forces govern the results. First, changes in the AA system alter investment incentives. The second force is the change in college assignments. Low learning-cost minority students have a stronger (weaker) incentive to make HC investments under a color-blind (proportional quota) system, and the changed investment is reflected in graduation rates, though mitigated by countervailing college placement shifts for these students. The lower (higher) HC incentives and worse (better)

Table 7. Counterfactual Minority Graduation Probability

...BY LEARNING-COST QUINTILE			
Learning-Cost Type Tier	Status Quo	Proportional Quota	Color-Blind
Top Achiever Quintile	0.562	0.562	0.580
Second Achiever Quintile	0.389	0.402	0.399
Third Achiever Quintile	0.299	0.311	0.302
Fourth Achiever Quintile	0.236	0.249	0.229
Bottom Achiever Quintile	0.159	0.165	0.151
...BY COLLEGE QUALITY QUINTILE			
College Quality Tier	Status Quo	Proportional Quota	Color-Blind
Top College Quintile	0.612	0.560	0.661
Second College Quintile	0.439	0.401	0.481
Third College Quintile	0.353	0.310	0.391
Fourth College Quintile	0.270	0.251	0.307
Bottom College Quintile	0.178	0.166	0.195

enrollment outcomes both hurt (help) the highest learning-cost minority students under a color-blind (proportional) system. The average effect of a change in AA systems on graduation rates across the student population is under 0.5%. In addition, because of the countervailing effects of the HC and enrollment changes, it need not be the case that the status-quo outcome is “in between” the color-blind and proportional quota outcomes.

The bottom panel of Table 7 breaks out the graduation rate by college quality quintile. Graduation rates increase for minority students who enroll at top colleges under a color-blind scheme, due mostly to composition effects: a smaller number of minority students with lower cost types enroll in the best colleges, relative to the status quo. Due to their lower learning costs and stronger HC accrual incentives in a color-blind system (Figure 5), the graduation probability of minority students at the best schools rises. Interestingly, the average graduation rate of minority students at low quality colleges also rises. Since minority students are assigned to low quality colleges at a higher rate under a color-blind system, the average learning-cost of minority students at these colleges falls. This in turn causes the level of HC accrued by minority students at low quality colleges to rise, which raises the average minority graduation rate. A proportional quota has the opposite effect. For example, under a proportional scheme minority students enrolling at the top quintile of schools have a lower average rate of graduation because on average they have higher learning costs and the investment incentives for top minority achievers are weakened.

The contrast between the upper and lower panels of Table 7 shows the importance of taking into account the composition of minority students enrolling in each college quality quintile when assessing the effects of AA. The differences also highlight the importance of taking a market-wide perspective when investigating the impact of AA on the rates at which minorities

Table 8. Counterfactual Minority Graduation Probability

Learning-Cost Type Tier	...BY LEARNING-COST QUINTILE			
	Proportional Quota	Proportional Quota, Fixed HC	Color-Blind	Color-Blind, Fixed HC
Top Achiever Quintile	0.562	0.573	0.580	0.549
Second Achiever Quintile	0.402	0.399	0.399	0.383
Third Achiever Quintile	0.311	0.304	0.302	0.297
Fourth Achiever Quintile	0.249	0.237	0.229	0.235
Bottom Achiever Quintile	0.165	0.160	0.151	0.159

graduate college: flows of heterogeneous students to alternative segments of the market can create a misleading picture if one focuses only on minorities who enroll within a narrow band of the quality spectrum under alternative admissions systems. For example, a naive reading of the bottom portion of of Table 7 would suggest that the mean minority graduation rate is highest under a color-blind system. However, this reasoning ignores compositional shifts in where minority students enroll, which generates the contrast between the top and bottom portions of Table 7.

Table 8 describes the graduation rate by learning-cost quintile when we hold fixed the status quo HC choices of the students. For comparison, the equilibrium graduation rates from Table 7 are presented as well. The difference between the equilibrium and fixed-HC analyses for the proportional quota are small, which is due to the fact that the effect on the investment incentives is modest for most student types and enrollment is (mechanically) identical in the two cases. The exception is the top achievers, students that significantly reduce their equilibrium HC investment under a proportional quota. The reduced HC explains why the graduation rate for this group under a fixed-HC analysis is higher than in an equilibrium analysis.

A naive, fixed-HC analysis of a color-blind AA system would be incorrect in both the quantitative and the qualitative effects on minority students. For high achievers, a fixed-HC analysis predicts lower graduation rates as the students shift into lower quality schools, whereas an analysis that controls for the increased incentives for low learning-cost minority students to accumulate HC predicts increased graduation rates for these students. On the other hand, a fixed-HC analysis will underestimate the decrease in the graduation rates of high learning cost minority students since it would not account for the lower incentive for these students to accumulate HC.

Table 9 depicts the effect of AA on minority household income by learning-cost quintiles. Since HC investment has negligible influence on household income conditional on graduation, the counterfactual impact of AA is mediated by college assignment shifts. An AA ban reduces household income across all quintiles of the type distribution, whereas a proportional quota would increase it across all quintiles. These effects are weakest for the highest achievers, where the difference between a color-blind system and a quota is around \$1,200/year. The strengths of the effects increase until the fourth quintile, where the difference between the two peaks at almost

Table 9. Minority HHI Conditional on Graduation by Learning-Cost Quintile

Learning Cost Type Tier	Status Quo	Proportional Quota	Color-Blind	Color-Blind, Fixed HC
Top Achiever Quintile	\$80,873	\$81,560	\$80,336	\$79,825
Second Achiever Quintile	\$75,709	\$77,223	\$74,714	\$74,236
Third Achiever Quintile	\$71,308	\$73,350	\$69,856	\$69,846
Fourth Achiever Quintile	\$67,360	\$69,143	\$65,264	\$65,723
Bottom Achiever Quintile	\$61,606	\$62,673	\$60,010	\$61,006

Table 10. Household Income (HHI) 10 Years After Graduation

	Average HHI of Minority Students	... Relative to Status Quo	Average HHI of Nonminority Student	... Relative to Status Quo
Status Quo	\$70,854	—	\$74,130	—
Color-Blind	\$70,036	\$818	\$74,288	\$158
Proportional	\$72,790	\$1,936	\$73,768	\$363

\$3,900/year of household income. The changes for nonminorities tend to have the opposite sign but are small and relatively inconsequential.

The last column of Table 9 compares the household income predictions of equilibrium and fixed-HC analyses in the color-blind case.⁴⁶ Low (high) learning-cost students make more (less) money under an equilibrium analysis because an equilibrium analysis predicts higher (lower) enrollment by minority students in the best (worst) colleges. For the lowest and highest learning-cost quintiles, where the HC incentive changes are most pronounced, not accounting for these incentive changes result in predictions that are off by \$500 and \$1,000 in annual household income. For high learning-cost students, the fixed-HC analysis reflects less than 40% of the equilibrium loss in household income.

Table 10 provides the average effect for both groups. The effect of changes in AA policy averaged across the entire population is on the order of \$20, implying a small loss of total household income production. In terms of the inequality in household income between the groups, a proportional quota results in the smallest gap and a color-blind system causes the largest gap.

Figure 6 evaluates the net effect of changes to graduation rate, household income, and learning cost in terms of equivalent variation, which is the proportional increase in status quo household income that would make the student as well off as under a counterfactual AA system. Values above 1 indicate utility increases, and the effect will be different for different types of students. For example, if we multiply the median minority student's status quo household income by a factor of 1.027, then his utility would rise to the same level as in equilibrium under a proportional

⁴⁶In the proportional case the effect is under \$30 for all quintiles because the enrollment and learning-costs are identical under both analyses and the differences in HC have only a small effect. See Appendix F.

Figure 6. Minority Equivalent Variations

quota regime. Minority equivalent variation for a proportional quota ranges from 0% to 2.8% of household income. The equivalent variation for a color-blind system is around 2.5%, and the effect remains significant for all but the lowest-cost minority students. The welfare effect for non-minorities is less than one-fifth as large.

7. CONCLUSION

This paper separates the effects of school quality, pre-college human capital, and unobserved student characteristics on the probability of graduating from college and household income 10 years after graduation. The key analytical tool we develop is a control function approach derived from the incentives provided by affirmative action programs. By using individual-level data from the B&B survey, rather than focusing only on elite private colleges, we can provide a market-wide analysis of how admissions rules impact incentives and how changes to one's college placement impact the returns to a college education, conditional on individual characteristics. We find that household income conditional on graduation is determined almost entirely by college quality and the student's unobservable type, but pre-college HC plays an important indirect role through its large effect on college placement and graduation probability. We also find that a large fraction of the accumulated pre-college HC is "wasted" in the rat-race to enroll in a high quality college.

AA is a prominent feature of the college market and plays a significant role in investment and redistribution. A strong AA regime such as a proportional quota results in minority students enrolling at better colleges, while color-blind admissions results in minority students predominantly enrolling in the bottom two quintiles of the college quality distribution. Interestingly, the effect on HC investment incentives is more ambiguous. A color-blind (quota) rule results in the lowest learning-cost minority students increasing (reducing) their HC investments, while high learning cost minority students reduce (increase) them. We also find that failing to account for changes to HC investment incentives causes significant errors in the estimates of the effect of AA.

Finally, we analyzed the strength of the incentive to accrue HC solely for its productive value relative to the incentive to accumulate HC to compete for access to a better college. We find that the competitive channel of incentives is stronger for all but the best students. Moreover, there is a stark contrast in the strength of these two incentives for most students: the competitive channel is twice as strong as the productive channel for half of college-bound students.

There remain many unanswered questions. E.g., are income effects of college quality, HC, and learning costs different for students in STEM fields relative to those in the humanities? Can one design a better AA scheme than the examples we study? However, the biggest question is whether one could use our model to say anything about the long-run impacts of different college admissions systems on the evolution of distributional inequalities over time. Our analysis, which was static by design, can only be the first step in such a research agenda.

References

- [1] Regents of the University of California v. Bakke. U.S. 438 265, 1978.
- [2] Mitra Akhtari and Natalie Bau. Race-based affirmative action and student effort. working paper, UCLA Economics Department 2018.
- [3] Rodney J. Andrews, Jing Li, and Michael F. Lovenheim. Quantile treatment effects of college quality on earnings. *Journal of Human Resources* 51(1):200–238, 2012.
- [4] Kate L. Antonovics and Ben Backes. The effect of banning affirmative action on college admissions policies and student quality. *Journal of Human Resources* 49(2):295–322, 2014.
- [5] Kate L. Antonovics and Richard H. Sander. Affirmative action bans and the “chilling effect”. *American Law and Economics Review* 15(1):252–299, 2013.
- [6] Peter Arcidiacono. Ability sorting and the returns to college major. *Journal of Econometrics* 121:343–375, 2004.
- [7] Peter Arcidiacono. Affirmative Action in Higher Education: How do Admission and Financial Aid Rules Affect Future Earnings? *Econometrica* 73(5):1477–1524, 2005.
- [8] Susan Athey and Phillip A. Haile. Identification of standard auction models. *Econometrica* 70(6):2107–2140, 2002.
- [9] Eduardo Azevedo and Jacob Leshno. A Supply and Demand Framework for Two-Sided Matching Markets. *Journal of Political Economy* 124(5):1235–1268, 2016.
- [10] Gary S. Becker. A Theory of Marriage: Part I. *Journal of Political Economy* 81:813–846, 1973.
- [11] Dan Black, Amelia Haviland, Seth Sanders, and Lowell Taylor. Why do minority men earn less? a study of wage differentials among the highly educated. *The Review of Economics and Statistics* 88(1):300–313, 2006.
- [12] Dan A. Black and Jeffrey A. Smith. How Robust is the Evidence on the Effects of College Quality? Evidence from Matching. *Journal of Econometrics* 121:99–124, 2004.
- [13] Dan A. Black and Jeffrey A. Smith. Estimating the Returns to College Quality Using Multiple Proxies for Quality. *Journal of Labor Economics* 24(3):701–728, 2006.
- [14] Aaron L. Bodoh-Creed, Joern Boehnke, and Brent R. Hickman. How Efficient are Decentralized Auction Platforms? Typescript, University of Chicago Department of Economics 2017.
- [15] Aaron L. Bodoh-Creed and Brent R. Hickman. College Assignment as a Large Contest. *Journal of Economic Theory* 175:88–126, 2018.
- [16] 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.
- [17] Dominic J. Brewer, Eric R. Eide, and Ronald G. Ehrenberg. Does It Pay to Attend an Elite Private College? Cross-Cohort Evidence on the Effects of College Type on Earnings. *The Journal of Human Resources* 24(1):104–123, 1999.
- [18] Sandra Campo, Isabelle Perrigne, and Quang Vuong. Asymmetry in First-Price Auctions with Affiliated Values. *Journal of Applied Econometrics* 18(2):179–207, 2003.

- [19] Hector Chade, Gregory Lewis, and Lones Smith. Student Portfolios and the College Admissions Problem. *Review of Economic Studies* 81(3):971–1002, 2014.
- [20] Chang Y. Chung and Thomas J. Espenshade. The opportunity cost of admission preferences at elite universities. *Social Science Quarterly* 86(2):293–305, 2005.
- [21] 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.
- [22] Stephen Coate and Glenn Loury. Anti-discrimination enforcement and the problem of patronization. *The American Economic Review* 83(2):92–98, 1993.
- [23] Stephen Coate and Glenn Loury. Will affirmative-action policies eliminate negative stereotypes. *The American Economic Review* 83(5):1220–1240, 1993.
- [24] Kalena E. Cortes and Lei Zhang. The incentive effects of the top 10% plan. Mimeo, 2012.
- [25] Christopher S. Cotton, Brent R. Hickman, and Joseph P. Price. Incentive Provision in Investment Contests: Theory and Evidence. Working Paper, University of Chicago Economics Department, 2017.
- [26] Julie B. Cullen, Mark C. Long, and Randall Reback. Jockeying for position: Strategic high school choice under texas' top ten percent plan. *Journal of Public Economics* 97:32–48, 2013.
- [27] Stacy Berg Dale and Alan B. Krueger. Estimating the Payoff to Attending a More Selective College: An Application of Selection on Observables and Unobservables. *Quarterly Journal of Economics* 117:1491–1528, 2002.
- [28] Elanor W. Dillon and Jeffrey A. Smith. Determinants of the match between student ability and college quality. *Journal of Labor Economics* 35(1):45–66, 2017.
- [29] Dennis Epple, Richard Romano, and Holger Sieg. On the Demographic Composition of Colleges and Universities in Market Equilibrium. *The American Economic Review* 92(2):310–314, 2002.
- [30] Dennis Epple, Richard Romano, and Holger Sieg. Admission, Tuition, and Financial Aid Policies in the Market for Higher Education. *Econometrica* 74(4):885–928, 2006.
- [31] Bruno Ferman and Juliano Assunção. Does Affirmative Action Enhance or Undercut Investment Incentives? Evidence from Quotas in Brazilian Public Universities. Typescript, Massachusetts Institute of Technology Department of Economics, 2011.
- [32] Ian Fillmore. Price Discrimination and Public Policy in the US College Market. Typescript, Washington University in St. Louis Department of Economics, 2017.
- [33] Aud S. Fox and A. KewalRamani. Status and Trends in the Education of Racial and Ethnic Groups. US Department of Education, National Center for Education Statistics Report NCES 2010-015: Retrieved 6/30/2017 from <http://nces.ed.gov/pubsearch>, 2010.
- [34] Roland G. Fryer, Devah Pager, and Jorg Spenkuch. Racial Disparities in Job Finding and Offered Wages. *Journal of Law and Economics* 56(3):633–689, 2013.
- [35] Chao Fu. Equilibrium Tuition, Applications, Admissions, and Enrollment in the College Market. *Journal of Political Economy* 122(2):225–281, 2014.
- [36] Emmanuel Guerre, Isabelle Perrigne, and Quang Vuong. Optimal Nonparametric Estimation of First-Price Auctions. *Econometrica* 68(3):525–574, 2000.
- [37] Fumio Hayashi. *Econometrics* Princeton University Press, 2000.
- [38] Brent Hickman, Timothy P. Hubbard, and Harry J. Paarsch. Identification and Estimation of a Bidding Model for Electronic Auctions. *Quantitative Economics* 8, 2017.
- [39] Brent R. Hickman. Using auction theory to study human capital investment in frictionless matching markets: A look at affirmative action in college admissions. Typescript, University of Chicago Department of Economics, 2012.
- [40] Brent R. Hickman and Timothy P. Hubbard. Replacing sample trimming with boundary correction in nonparametric estimation of first-price auctions. Typescript, University of Chicago Department of Economics, 2012.
- [41] Jessica S. Howell. Assessing the Impact of Eliminating Affirmative Action in Higher Education. *Journal of Labor Economics* 28(1):113–66, 2010.
- [42] Timothy P. Hubbard, Tong Li, , and Harry J. Paarsch. Semiparametric Estimation in Models of First-Price , Sealed-Bid Auctions with Affiliation. *Journal of Econometrics* 168(1):4–16, 2012.

- [43] Thomas J. Kane. Racial and Ethnic Preference in College Admissions. In Christopher Jencks and Meredith Phillips, editors, *The Black-White Test Score Gap*, pages 431–456. Brookings Institution Press, 1998.
- [44] Adam Kapor. Distributional Effects of Race-Blind Affirmative Action. Princeton University Economics Department 2016.
- [45] G. Kena, L. Musu-Gillette, J. Robinson, X. Wang, A. Rathbun, J. Zhang, S. Wilkinson-Flicker, A. Barmer, and E. Dunlop Velez. *The Condition of Education 2015*. US Department of Education, National Center for Education Statistics Reports NCES 2015-144: Retrieved 6/30/2017 from <http://nces.ed.gov/pubsearch>, 2015.
- [46] Mark C. Long. College Quality and Early Adult Outcomes. *Economics of Education Review* 27:588–602, 2008.
- [47] Mark C. Long and Marta Tienda. Winners and losers: Changes in Texas university admissions post-Hopwood. *Educational Evaluation and Policy Analysis* 30(3):255–280, 2008.
- [48] Jane S. Lopus and Lynn Paringer. *The principles of economics textbook*. Working paper, 2011.
- [49] N. Gregory Mankiw. *Principles of Economics*, 5th Edition. Mason, OH: Cengage Learning, 2009.
- [50] Campbell R. McConnell, Stanley L. Brue, and Sean M. Flynn. *Economics*, 18th Edition. New York, NY: McGraw-Hill Irwin, 2009.
- [51] Robert J. Morse, editor. *America's Best Colleges*. Washington D.C.: US News & World Report, 1994.
- [52] S. Navarro. Control functions. In S.N. Durlauf and L.E. Blume, editors, *Microeconometrics*. Palgrave Macmillan, London, 2010.
- [53] 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.
- [54] Wojciech Olszewski and Ron Siegel. Large contests. *Econometrica* 84(2):835–854, 2016.
- [55] Harry J. Paarsch. Deciding between the common and private value paradigms in empirical models of auctions. *Journal of Econometrics* 51:191–215, 1992.
- [56] Sean F. Reardon. School Segregation and Racial Academic Achievement Gaps. Stanford Center For Education Policy Analysis Working Paper 15-12, 2016.
- [57] Sean F. Reardon, Demetra Kalogrides, and Ken Shores. The Geography of Racial/Ethnic Test Score Gaps. Stanford Center For Education Policy Analysis Working Paper 16-10, 2017.
- [58] Sean F. Reardon, Joseph Townsend, and Lindsay Fox. Characteristics of the Joint Distribution of Race and Income Among Neighborhoods. Stanford Center For Education Policy Analysis Working Paper 15-01, 2015.
- [59] Bernard W. Silverman. *Density Estimation for Statistics and Data Analysis*. London: Chapman and Hall, 1986.
- [60] Jessica C. Smith and Carla Medalia. Health Insurance Coverage in the United States: 2014. US Census Bureau, Current Population Reports P60-253: Retrieved 6/30/2017 from <https://www.census.gov/content/dam/Census/library/publications/2015/demo/p60-253.pdf>, 2015.
- [61] Michael Spence. Job Market Signaling. *The Quarterly Journal of Economics* 87(2):355–374, 1973.
- [62] Gordon C. Winston and David J. Zimmerman. Peer Effects in Higher Education. NBER Working Paper 9501, 2003.
- [63] Danny Yagan. Law school admissions under the affirmative action ban. *Mimeo*, 2012.

SUPPLEMENTAL ONLINE APPENDIX

to accompany

Pre-College Human Capital Investment and Affirmative Action: A Structural Policy Analysis of US
College Admissions

by Aaron L. Bodoh-Creed and Brent R. Hickman

Appendix A. Theoretical Model Assumptions

We do not get existence of a monotone, pure-strategy equilibrium without some assumptions on the underlying theoretical model of human capital investment. Although not all these assumptions are used directly by our estimator, it implicitly relies on them as its behavioral foundation, as is common in all causal econometric analyses. For the sake of completeness, we briefly discuss here the various assumptions from Bodoh-Creed and Hickman [4, henceforth BCH18] that underlay our model of academic competition for college seats. These assumptions belong to two different groups, the first having to do with behavior of random variables in the model and the second pertaining to behavior of utility functions in the model.

First is the assumption that each of the central distributions in the model—i.e., the CDFs of student characteristics ($F_M(q)$, $F_N(q)$) and the CDF of college quality $F_P(p)$ —are well-behaved (see BCH18 Assumptions 7, 9). Specifically, we require that they have differentiable densities and finite supports. Intuitively, since Harvard is our estimated upper bound on the quality distribution for the sample year, we essentially assume there exist no other colleges that are thousands of times more effective than Harvard at helping their graduates generate income, and there are no students who are thousands of times more effective at acquiring productive human capital than the best students Harvard enrolls. For the distributions of student characteristics, we also assume a common support. In other words, both the minority and non-minority groups contain the same wide swath of humanity from the most talented and most well-resourced students to the least, and we allow only the relative frequencies at various levels of q to be different across groups. Finally, our assumption of zero-mean, normal matching shocks ϵ , which is explicitly part of our estimator, is consistent with requirements of the foundational theory (BCH18 Assumptions 10, 11, Section 6), being uniformly continuous and having full support.

Second is the assumption that the match utility function $U(p, s, q)$ and cost function $C(s, q)$ are well-behaved (BCH18 Assumptions 1–6). The first and most obvious aspect of good behavior is that these functions must be monotone in each argument, holding the others fixed. For costs, achieving at a higher level always costs more and having higher q is always worse. For match utility, increases in human capital s are weakly valued, college quality p increases are strictly valued, and the unobserved traits of the student q that drive up her investment costs do not somehow increase her match utility. The second aspect of good behavior is that match utility is weakly concave in human capital and costs are strictly convex in human capital (holding other factors fixed). The role of utility concavity is fairly ubiquitous and well-understood in economics. More central to the idea of a monotone equilibrium in our contest model is the third aspect of good behavior, complementarity. In the match utility function, p , s , and q are assumed to be

weak pairwise complements in the sense that the marginal utility of each is weakly increasing in the levels of the other two. For investment costs, the marginal dis-utility of human capital output s is strictly increasing in q , an idea commonly referred to as “supermodularity” in auctions and contests. If one student has a lower q than another, then it is not only less costly for her to achieve a given s , but it is also less costly on the margin for her to increase achievement to $s^0 > s$. This assumption ensures that competitors' best responses are increasing in their rivals' choices. The natural aspect of good behavior for utility is that there exists some level of investment that is so high it would not be profitable for anyone, ever. I.e., we assume that no high-school student (or her parents) would be willing to forego all her leisure time (including sleep) continuously for 4 years, just to improve her chances of being admitted to Harvard.

Appendix B. Estimation Technical Details

B.1. Stage I GMM Estimator. The first hurdle to overcome is to find a computationally tractable way of representing the selected joint distribution of (P, S) conditional on graduation. High-dimensional density estimation is a difficult problem both computationally and because of the rapid decay of optimal statistical convergence rates as the dimensionality of the underlying random variable grows.⁴⁷ Recent work in the auctions literature by Hubbard, Li, and Paarsch [15] has employed parametric copula functions to solve this problem. Sklar's Theorem states that any absolutely continuous joint distribution can be represented as a composition $F_{P,S}(p, s | \text{grad}) = C_j(F_{P_j}(p | \text{grad}), G_j(s | \text{grad}))$, $j = 2, \dots, M, N$, where $C_j(\cdot, \cdot | \text{grad})$ is a unique copula function. This implies that the rapidly increasing computational cost and data-hungriness of nonparametric estimators come from the complexity of the correlation structure C_j since the complexity of the marginal distributions does not increase with the dimension of the joint distribution. Hubbard, Li, and Paarsch [15] therefore propose a flexible approach to estimating the marginal distributions, while simplifying the copula with parametric assumptions for tractability. This allows the econometrician to maintain the familiar \sqrt{n} convergence rate when estimating a multi-dimensional joint distribution. We follow this dimension reduction strategy by adopting the Gumbel-Hougaard copula, $C(r, q; n) = \exp(-((-\log(r))^n + (-\log(q))^n)^{1/n})$, $n \geq 1$.⁴⁸ One advantage of the Gumbel-Hougaard copula is that it implies a closed-form expression for the Kendall's τ rank correlation index: $\tau_{P,S}^j = \frac{n_j - 1}{n_j}$, $j = M, N$.

For the selected marginal distributions we propose a flexible, semi-nonparametric approach based on B-splines. Like orthogonal polynomials, B-splines are defined as a linear combination of

⁴⁷See Silverman [20] and Campo, Perrigne, and Vuong [5] for a lengthy discussion on this concept.

⁴⁸Wang also experimented with several other copula functions including the Frank copula, $C(r, q; n) = \frac{1}{n} \log \left(1 + \frac{(\exp(-nr) - 1)(\exp(-nq) - 1)}{(\exp(-n) - 1)} \right)$, $n \geq 2$; the Clayton copula, $C(r, q; n) = \max\{r^{-n} + q^{-n} - 1, 0\}^{-1/n}$, $n \geq 1$; and the Gaussian copula, $C(r, q; n) = G_R(G^{-1}(r), G^{-1}(q))$ where G is a standard normal CDF and G_R is a bivariate normal CDF with correlation matrix $R = \begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix}$, $\rho \in [-1, 1]$. All produced very similar results, which is consistent with the assumption that our parametric restriction of the copula function provides a robust approximation to the nonparametric correlation structure.

global basis functions, and B-splines can be made arbitrarily flexible while remaining much better behaved than global polynomials.⁴⁹ For the selected marginal distributions of school assignment, we begin by specifying knot vectors $k_j^p = \underline{p} = k_{j1}^p < k_{j2}^p < \dots < k_{j, K_j^p+1}^p = \bar{p}$ that uniquely define a set of $K_j^p + 3$ cubic B-spline basis functions $B_{jk}^p(p) : [\underline{p}, \bar{p}] \rightarrow \mathbb{R}$, $k = 1, \dots, K_j^p + 3$, which in turn define our parameterization of the CDFs:

$$F_{P_j}(p; \text{grad}; g_j^p) = \sum_{k=1}^{K_j^p+3} \hat{a}_{jk}^p B_{jk}^p(p), \quad j \in \mathcal{M}, N, g.$$

For the marginal distribution of the HC index, S , we have an additional challenge: since its units (and therefore the relevant domain to span) are unknown ex ante, we instead parameterize the selected marginal quantile functions, whose domain is always $[0, 1]$. Let the knot vectors and basis functions for selected HC quantile functions be $k_j^q = 0 = k_{j1}^q < k_{j2}^q < \dots < k_{j, K_j^q+1}^q = 1$ and $B_{jk}^q(r) : [0, 1] \rightarrow \mathbb{R}$, $k = 1, \dots, K_j^q + 3$, respectively, with B-spline marginal quantile functions parameterized similarly as above by $Q_{S_j}(r; \text{grad}; g_j^q) = \sum_{k=1}^{K_j^q+3} \hat{a}_{jk}^q B_{jk}^q(r)$, $j \in \mathcal{M}, N, g$.

The parameterized, selected joint distributions are given by

$$F_{PS}(p, s; \text{grad}; g_j^p, g_j^q, \eta_j) = C \prod_{j=1}^h F_{P_j}(p; \text{grad}; g_j^p) Q_{S_j}^{-1}(s; \text{grad}; g_j^q; \eta_j), \quad j \in \mathcal{M}, N, g.$$

Going forward, one important detail to note is that the parameters η_M and η_N reflect the empirical correlation structure for the selected joint distribution of P and S for college graduates only. Below we will define other notation for separate copula parameters that apply to the selection-corrected joint distribution for all college enrollees (i.e., including dropouts).

In order to complete our GMM estimator, we also need to construct empirical analogs to the joint and marginal distributions of (p, s) . In the case of CDFs, we use the standard Kaplan-Meier empirical distribution functions

$$\hat{F}_{P_j}(p; \text{grad}) = \frac{\sum_{i=1}^l \mathbf{1}(p_i \leq p) \mathbf{1}(i \leq j)}{\sum_{i=1}^l \mathbf{1}(i \leq j)}, \quad \text{and}$$

$$\hat{F}_{PS}(p, s; \text{grad}) = \frac{\sum_{i=1}^l \mathbf{1}(p_i \leq p) \mathbf{1}(s_i \leq s) \mathbf{1}(i \leq j)}{\sum_{i=1}^l \mathbf{1}(i \leq j)}, \quad j \in \mathcal{M}, N, g.$$

For the empirical marginal quantiles of S , we use a new method developed by Hedblom, Hickman, and List [13] for smooth nonparametric quantile estimation. For a random sample $S_j = \{S_{ji}\}_{i=1}^{l_j}$ of size $l_j = \sum_{i=1}^l \mathbf{1}(i \leq j)$, this estimator exists as a weighted average of the ordered data $f(S_j(1) \dots S_j(2) \dots S_j(l_j))g$. Specifically, for $k \in \{1, 2, \dots, l_j\}$, the $(k/l_j)^{\text{th}}$ empirical quantile

⁴⁹For a brief primer on B-splines and their advantages in empirical auctions models, see Hickman, Hubbard, and Paarsch [14] and Bodo-Creed, Boehnke, and Hickman [3].

is estimated as

$$\hat{Q}_{S_j}(k/l_j) = \hat{a} \sum_{i=1}^{l_j} P_{ik}^{l_j} S_j(i),$$

where the weights $P_{ik}^{l_j}$ are known and mimic the limiting behavior of a re-sampled quantile estimator as the number of simulated samples approaches infinity. Specifically, $P_{ik}^{l_j}$ gives the probability that the i th order statistic of S_j will occupy the k th position in an ordered, randomly generated bootstrap sample from the raw data.⁵⁰ However, an important constraint for this estimator is that, for a fixed sample size l_j , it can only be evaluated at quantile ranks on the discrete grid $\frac{1}{l_j}, \dots, \frac{l_j}{l_j}$. Therefore, assuming $l_j = 101, j \in \{M, Ng\}$, we also define $r_j = (r_{j0}, r_{j1}, r_{j2}, \dots, r_{j100})$ as a vector of quantile ranks at which the empirical quantile function is to be evaluated, where $r_{j0} = 1/l_j$ and $r_{jk} = \text{round}(kl_j/100)/l_j, 1 \leq k$.

Finally, in order to estimate $s_{\#}$ we need to introduce four additional ancillary parameters. First, for notational ease let $p = (g_M^p, g_N^p, g_M^q, g_N^q, n_M, n_N, b^s, b^r)$ summarize all primary Stage I parameters except for the shock variance. Note that if p is known, Equation (6) defines the selection-corrected joint distribution of (P, S) , $F_{PS}(p, s; j; p)$, for group $j \in \{M, Ng\}$ as a function of the entire parameter vector p along with its two marginal distributions, $F_{P_j}(p; p)$ and $G_j(s; p)$. Let η_j denote the best-t copula parameter of the selection-corrected joint distribution:

$$\eta_j = \arg \min_{n \in \{1, \dots, n\}} \int_{S^h} \int_{P^h} C(F_{P_j}(p; p), G_j(s; p); \eta_j) F_{PS}(p, s; j; p)^{i_2} dp ds.$$

Implicitly, η_j is a function of b^s since these directly control how well the ranks of $S(e_i, a_i; b^s)$ predict the ranks of i 's placement in P space.

Now, let $F_{PS}(p, s; j; p, s_{\#})$ denote the joint distribution of (P, S) implied by the marginal distributions $F_{P_j}(p; p)$ and $G_j(s; p)$, but assuming rank-order allocations with respect to NHC are generated by mean-zero, normal shocks with variance $s_{\#}^2$. Let $\eta_j(s_{\#})$ denote the best-t copula

⁵⁰A benefit to this method is that it is differentiable and more efficient than the traditional empirical quantile estimator, $\inf_{s: \frac{1}{l_j} \sum_{i=1}^{l_j} 1(S_{ji} \leq s)} (k/l_j)$. Intuitively, the empirical quantile estimator, being the nearest-neighbor inverse of the Kaplan-Meier empirical CDF, incorporates cardinal information from only a single datum, with all other data providing only ordinal information. In contrast, $\hat{Q}_{S_j}(k/l_j)$ as defined above uses both ordinal and cardinal information from the entire sample. See Hedblom, Hickman, and List [13] for additional details.

parameter for that joint distribution, in the following sense:

$$\begin{aligned} \eta_j(s_{\#}) &= \arg \min_{n_2[1, \infty)} \int_S \int_P C_{F_{P_j}(p; p), G_j(s; p)}; n_{F_{PS}(p, s_{jj}; p, s_{\#})}^2 dp ds \\ \text{Subject to :} \\ F_{PS}(p, s_{jj}; p, s_{\#}) &= \int_s \int_p f_{P_j S}(p, s, j; p, s_{\#}) g_j(s; p) dp ds \\ &= \int_s \int_p f_{\#} P_j^{-1}(p; p, s_{\#}) \frac{dP_j^{-1}(p; p, s_{\#})}{dp} g_j(s; p) dp ds \\ P_j^{-1}(p; p, s_{\#}) &= H_j^{-1} F_{P_j}(p; p); p, s_{\#} \\ H_j(t; p, s_{\#}) &= \int_{\mathcal{X}} f_{\#}(s_{\#}) g_j(x_{\#}; p, s_{\#}) dx_{\#}, j = M, N. \end{aligned}$$

The ancillary parameters n_M , n_N , $n_M(s_{\#})$, and $n_N(s_{\#})$ are used below to define moment conditions for estimation of b^s and $s_{\#}$. Intuitively, η_j is the copula parameter that best reflects the correlation structure between P and S implied by the data (post-selection-correction), and $\eta_j(s_{\#})$ is the copula parameter that best reflects the correlation structure between P and S generated endogenously by our structural model of a noisy, rank-order college admissions contest given $s_{\#}$ and the empirical marginal distributions of P and S . With the above definitions, we can now formalize our Stage I GMM estimator:

$$\begin{aligned} \hat{\rho}_{s_{\#}} &= \arg \min_{\hat{\rho}_{s_{\#}}} \left(\sum_{i=1}^I D_{Mi} \int_P \int_S F_{P_M}(p_i; \text{grad}; g_M^p) \hat{F}_{P_M}(p_i; \text{grad})^2 \right. \\ &\quad \left. + (1 - D_{Mi}) \int_P \int_S F_{P_N}(p_i; \text{grad}; g_N^p) \hat{F}_{P_N}(p_i; \text{grad})^2 \right. \\ &\quad \left. + \sum_{k=0}^{100} \int_S Q_{SM}(r_{Mk}; \text{grad}; g_M^q) \hat{Q}_{SM}(r_{Mk})^2 + \int_S Q_{SN}(r_{Nk}; \text{grad}; g_N^q) \hat{Q}_{SN}(r_{Nk})^2 \right. \\ &\quad \left. + \sum_{i=1}^I D_{Mi} \int_P \int_S F_{PS}(p_i, s_{jM}, \text{grad}; g_M^p, g_M^q, n_M) \hat{F}_{PS}(p_i, s_{jM}, \text{grad})^2 \right. \\ &\quad \left. + (1 - D_{Mi}) \int_P \int_S F_{PS}(p_i, s_{jN}, \text{grad}; g_N^p, g_N^q, n_N) \hat{F}_{PS}(p_i, s_{jN}, \text{grad})^2 \right. \\ &\quad \left. + \sum_{l=1}^L [G_{Ml} \int_{Z_M} b^r]^2 + [G_{Nl} \int_{Z_N} b^r]^2 \right. \\ &\quad \left. + \frac{n_M}{n_M} \frac{1}{n_M} \int_{\mathcal{X}} 1^2 + \frac{n_N}{n_N} \frac{1}{n_N} \int_{\mathcal{X}} 1^2 \right. \\ &\quad \left. + \frac{n_M}{n_M} \frac{1}{n_M(s_{\#})} \int_{\mathcal{X}} 1^2 + \frac{n_N}{n_N} \frac{1}{n_N(s_{\#})} \int_{\mathcal{X}} 1^2 \right), \end{aligned}$$

Subject to :

$$s_i = S(e, a; b^s), i = 1, \dots, l$$

$$\frac{\partial S(e, a; b^s)}{\partial e} > 0, \frac{\partial S(e, a; b^s)}{\partial a} > 0 \quad 8(e, a; b^s), \text{ and } \max_{(e,a) \in 2R^2} f S(e, a; b^s) g = 1$$

$$g_{j,k}^v \quad 1 \quad g_{jk}^v, \quad k = 2, \dots, K_j^v + 3, \quad v = p, s, \quad j \in 2 fM, Ng$$

$$\min f G_{M1}, \dots, G_{ML}, G_{N1}, \dots, G_{NL}, g \quad r(p, s; b^r) \quad 1 \quad 8(p, s)$$

$$\frac{\partial r(p, s; b^r)}{\partial s} > 0, \frac{\partial^2 r(p, s; b^r)}{\partial s^2} < 0, \quad 8(p, s)$$

$$Z_{jl} = [1, p_l, p_l^2, p_l^3, \bar{S}_{jl}, \bar{S}_{jl}^2, \bar{S}_{jl}^3, p_l \bar{S}_{jl}, p_l \bar{S}_{jl}^2, p_l^2 \bar{S}_{jl}], \quad j \in 2 fM, Ng, \quad l = 1, \dots, L$$

$$\bar{S}_{jl}^k \text{ agrees with equations (6) (8) (9), } k = 1, 2, 3, \quad j \in 2 fM, Ng, \quad l = 1, \dots, L$$

$$n_M, n_N, n_M, n_N, n_M(s_{\#}), n_N(s_{\#}) \in [1, \infty); \text{ and } s_{\#} > 0.$$

The first two summations in the objective function are the moment conditions for the selected marginal distributions of P and the selected marginal quantile functions of S. The third summation contains moment conditions for the copula of the selected joint distribution of (P, S). The fourth summation contains the selection-corrected regression equations for graduation probabilities. The second to last line of the objective function contains moment conditions for the single-index parameters: they minimize the distance between the selection-corrected rank correlations, t_{PS}^M, t_{PS}^N , and their theoretical maxima of one. Recall from the discussion in Section 4.1.1 that we define the parameters b^s within our model as those which maximize the Kendall's τ rank correlation between the implied single index S and P, which is why the estimate \hat{b}^s is chosen in this way. The final line of the objective function contains moment conditions for the matching shock variance: it is chosen to minimize the distance between the empirical rank correlations, t_{PS}^M and t_{PS}^N , and their model-generated analogs, $t_{PS}^M(s_{\#})$ and $t_{PS}^N(s_{\#})$.

As for the constraints, the last line imposes natural bounds on the shock variance and copula parameters. The two lines above this establish the selection correction procedure for the graduation probability regressions, and the two lines above that impose regularity conditions on the graduation probability parameters. The third constraint from the top imposes monotonicity on the B-spline CDFs and quantile functions, and the first two lines define the s_i 's as a single index in (e, a) , and they impose monotonicity and a scale normalization to fix its units.

B.2. Stage II Type Distribution GMM Estimator. Amending notation somewhat, for computational convenience we first parameterize the assignment functions and inverse equilibrium strategies for group $j \in 2 fM, Ng$ as flexible B-splines, similarly as we did for other functionals in stage I. B-spline functions are defined as linear combinations of basis functions $B_{jk}^t(t), k = 1, \dots, K_j^t + 3$ and $B_{jk}^s(s), k = 1, \dots, K_j^s + 3$.⁵¹ When we combine these with weights $l_j^t \in 2 R^{K_j^t+3}$

⁵¹As explained in the appendix, the fundamental building block of B-spline functions are knot vectors $k_j^t = f \underline{t} = k_{j1}^t < \dots < k_{j,K_j^t+1}^t = \bar{t}g$ for the assignment functions, and $k_j^s = f \underline{s} = k_{j1}^s < \dots < k_{j,K_j^s+1}^s = \bar{s}g$, for equilibrium strategies. These knot vectors partition the domains of the assignment functions and inverse strategies into K_j^t and

and $l_j^s \in \mathbb{R}^{K_j^s+3}$ our parameterized B-spline functions have the form $P_j(t; l_j^t) = \sum_{k=1}^{K_j^t+3} l_{jk}^t B_{jk}^t(s)$, and $q_j(s; l_j^s) = \sum_{k=1}^{K_j^s+3} l_{jk}^s B_{jk}^s(s)$. Given a pre-specified grid of points $\{t_1, \dots, t_{K^T}\}$ spanning $[t, \bar{t}]$, the assignment function weights are chosen to satisfy

$$\hat{l}_j^t = \arg \min_{l \in \mathbb{R}^{K_j^t}} \left(\sum_{k=1}^{K^T} P_j(t_k; l) - F_{P_j}^{-1} H_j(t_k; \hat{p}, \hat{s}_\#); \hat{p} \right)^2$$

Subject to $l_k < l_{k+1}, k = 1, \dots, K_j^t + 2$.

The assignment mappings are a function of Stage I parameters, which are taken as fixed in Stage II, which is why we express the relevant B-spline weights using hat notation. On the other hand, the inverse equilibrium strategies solve

$$(17) \quad \hat{l}_j^s(a) = \arg \min_{l \in \mathbb{R}^{K_j^s}} \sum_{i=1}^I D_{wi} q_j(s_i; l) - \ddot{q}_{ji}^2$$

Subject to :

$$\ddot{q}_{ji} c^0(s_i) = E_\#^h U_p \left(P_j(s_i + \#; l_j^t), s_i, \ddot{q}_{ji}; a, \hat{p} \right) - P_j^0(s_i + \#; l_j^t) + E_\#^h U_s \left(P_j(s_i + \#; l_j^t), s_i, \ddot{q}_{ji}; a, \hat{p} \right)$$

$l_k > l_{k+1}, k = 1, \dots, K_j^s + 2, j \in \{M, N\}$,

and are not fixed during Stage II. The B-spline strategies are a function of the Stage I parameters as well as the Stage II utility parameters a , and therefore must be adjusted each time these are updated during estimator runtime.

The control function, takes the form

$$(18) \quad y(s, D_{Mi}; a, \hat{p}, \hat{s}_\#, l_M^s, l_N^s) = \log [D_{Mi} q_M(s; l_M^s(a)) + (1 - D_{Mi}) q_N(s; l_N^s(a))],$$

with the extra parameter arguments emphasizing its implicit dependence on Stage I objects.

Moving forward we suppress the additional pre-determined parameter arguments for notational simplicity. We can now re-express the matrix of explanatory variables as

$$X(a) = \begin{pmatrix} 1 & \log(p_1) & \log(s_1) & y(s_1, D_{M1}; a, l_M^s, l_N^s) \\ 1 & \log(p_2) & \log(s_2) & y(s_2, D_{M2}; a, l_M^s, l_N^s) \\ \vdots & \vdots & \vdots & \vdots \\ 1 & \log(p_I) & \log(s_I) & y(s_I, D_{MI}; a, l_M^s, l_N^s) \end{pmatrix}$$

from which our utility parameter estimator is defined by

$$(19) \quad \hat{a} = \arg \min_{a \in \mathbb{R}_+} [D_w(W - X(a))a]^\top [D_w(W - X(a))a]$$

[‡] K_j^s smaller segments, respectively. On each of these sub-intervals, the function t is governed by a strict subset of the overall basis functions; this locality property is what makes B-splines better behaved than Chebyshev polynomials. The knot vectors also uniquely determine the shapes of the basis functions themselves (see appendix for further detail).

In words, Stage II consists of estimating the HHI regression, where the third explanatory variable is a nonlinear function of all model parameters with a form that is derived from the equilibrium conditions of our theory of HC investment. Also, note that problems 17 and 19 are solved simultaneously.

The final step of estimation is to recover the type distributions. We recover the distribution of types by convolving the inverse strategy function for each demographic group with the CDFs of the distribution of HC choices for the respective group.

B.3. Practical Issues for Implementation of the Estimator. In the implementation of our GMM estimator of Stage I parameters, we adopt a simplification for computational convenience. First, during solver runtime we normalize $b_1^s = 1$ instead of constraining the objective function so that the maximal single index value is one. This simplifies the problem by reducing the number of parameters to choose and constraints to satisfy. After the estimator has run we re-scale the HC single index so that its maximal possible value is one, and we accordingly make adjustments to the graduation probability parameters and HC distributions to reflect the re-scaling.

There are several model tuning parameters that we must specify, among which are knot vectors k_j^p , k_j^q , k_j^t , and k_j^s . We adopt the convention that knots are to be chosen uniformly in empirical quantile space, as this evenly spreads the statistical power of the data across all basis functions and simplifies the decision to a choice of the number of knots. Specifically, we chose $K_M^p = K_N^p = K_M^t = K_N^t = K_M^s = K_N^s = 5$ (i.e., knots at the quintiles with 8 total B-spline basis functions) for the selected school quality distributions, assignment functions, and inverse strategies, respectively; and $K_M^q = K_N^q = 10$ (i.e., knots at the deciles with 13 basis functions) for the selected HC quantile functions. Additional knots did not appreciably improve model fit. When approximating the assignment mappings, we imposed a truncation $\underline{t} = \min_i \{f_{sq} - 5\hat{s}_\#\}$, $\bar{t} = \max_i \{f_{sq} + 5\hat{s}_\#\}$ on the support of NHC, and we chose a set of evaluation points that included the modes of the B-spline basis functions and the midpoints between the modes.

We calibrated k from the U.S. Census Bureau's Current Population Survey (CPS), which is consistent with the structural estimates of Heckman, Humphries, and Veramendi [12]. Recall that the students in our sample were initially surveyed after graduation in 1993, and the HHI data we use was collected in the 2003 follow-up survey. To get a benchmark for the fraction of college graduate incomes garnered by dropouts, we computed the ratio of the average HHI of 33-year-olds in the 2003 CPS survey with some college to the average HHI of 33-year-old college graduates (with no additional post-graduate education). The result is a value of $k = 0.714$. As a robustness check we also repeated our estimates assuming k values ranging from 0.5 to 0.9. Only two aspects of our analysis change. First, the choice of k affects the estimated values of a_q and the distribution of types, $F_j(q; I_j^q)$, $j \in M, N, g$. However, the overall role of q in the wealth equation, which we measure as a_q times the standard deviation of $\log(q)$, is stable to changes in k . The second thing that is affected is the productive channel of incentives, which we estimate to be weaker as k rises. This is intuitive since, as we see in Section 5.5, k only affects the graduation

probability, so anything that makes the utility gap between completing college and dropping out shrink will weaken the direct, productive benefit of HC.

Finally, two sources of sampling weights were used in our empirical implementation, but in order to avoid further complicating notation we left them out of the formal definition above. The first is cross-sectional sampling weights contained in the B&B data. In Stage I these were used for group $j \in \{M, N, G\}$ to calculate the empirical analogs of the joint distributions $F_{P,S}(p, s; j, \text{grad})$, the marginal quantile functions $Q_{S_j}(r)$, and the marginal distributions $F_{P_j}(p; j, \text{grad})$. In simple terms, each of these functions at a point (p, s) is a sample mean of indicator functions evaluated at each datapoint, and we converted them into weighted sample means (in the usual way) using the B&B cross-sectional weights. In Stage II they were also used to weight each component of the HHI regression (in the usual way that weighted regressions are constructed) and to calculate the empirical analog of the type distributions $F_j(q)$. The second source of sampling weights came from IPEDS. In Stage I, the graduation probability regression (the fourth summation in the definition of $[\hat{p}, \hat{s}_\#]$ above) is converted into a weighted regression in the usual way by using the number of individuals in each school-race freshman cohort in 1988 as weights.

Appendix C. Goodness of Fit Metrics

C.1. Stage I Estimates. Table 11 contains goodness-of-fit metrics for stage I estimates, which include the graduation rate equation, the single-index equation (that condenses exam scores E and $NMGPA$ into a single HC index S for predicting college placement), and the fit for the parametric representation of the joint distribution between P and S . Since each of these three things was estimated via a least-squares moment-matching routine, in all three cases it is possible to compute an R^2 or something like it to gauge model fit. The estimated graduation rate model is able to account for 57% of the variation in observed college-race group graduation rates. Using a parametric Gumbel-Hougaard copula for the correlation structure between P and S and B-spline forms for the marginal CDFs does not appear to unduly restrict the data-generating process. The resulting joint distribution achieves an R^2 of 0.935 for predicting the values of the empirical joint CDF of P and S .

For the single index equation, there are two alternative measures of fit to consider. First, recall that the single index parameters are chosen so as to maximize the the power of one's rank in S to predict one's rank in P . The pooled Kendall's τ measure in the table—a weighted average of the Kendall's τ 's for the minority and non-minority joint distributions—demonstrates that the single index parameters achieve a high degree of predictive power for college placement, since the probability of concordance minus the probability of discordance is over 92%. An alternative measure of fit for the single HC index model that more closely resembles an R^2 measure is 1 minus the noise-signal ratio, or in other words, one minus the standard deviation of the matching shock, s , divided by the standard deviation of the single index S , roughly 0.836.

C.2. Stage II Estimates. Note that stage II of estimation uses an additional piece of student-level information which did not enter stage I: household income. Unlike the other variables, household

Table 11. Stage I Goodness-of-Fit

Model Component	Fit Metric
Graduation Rate Fit	$R^2 = 0.565$
F_{PS} Fit	$R^2 = 0.935$
Single-Index Fit	Pooled Kendall's $\tau = 0.921$ 1 noise/signal-ratio under $\hat{\tau} = 0.836$

Table 12. HHI Equation Fit

College Quality Tier	Quality Upper Bound	Median Household Income	Model Predicted Median Household Income
Top College Quartile	0.955	\$80,000	\$80,917
Second College Quartile	0.660	\$77,500	\$77,212
Third College Quartile	0.473	\$75,932	\$73,129
Bottom College Quartile	0.312	\$65,000	\$68,272

income is available for a random subsample of the B&B data.⁵² Now we turn to the goodness of the fit of the HHI equation. Since the HHI equation includes a component identified through the structure of our model, q , we think the goodness of fit can be best assessed by comparing how well our model predicts the HHI of students from different quartiles of the college quality distribution 10 years after college. Because of a significant positive skew in the HHI distribution, we think the a comparison of the empirical and model generated medians is most revealing. For reference, the median HHI income in the data is \$78,000, and the median model generated HHI is \$77,449.

In order to get a more granular look at the fit of our HHI model, Table 12 provides both the median income within each quartile as well as the model generated median income. For reference we also include the upper bound of each quartile to give a sense for how much the college quality is changing. What we see is that the model overestimates the HHI of graduates from the lowest quality schools, and slightly overestimates the HHIs of graduates from the best colleges.

Now we consider the fit of the inverse strategy functions. We use the group-specific inverse strategy functions to directly impute a value of $\log(q)$ for each HC choice. When this is done exactly from the first order condition, there is no issue of goodness of fit. However, we parameterize the inverse strategy functions with B-splines for computational convenience. When we compare the B-spline and direct imputations of type, we find that the difference is less than or

⁵²To verify that data loss on the income measure is truly random, we ran two-sample Kolmogorov-Smirnov tests on the other three variables in the data, P , E , and A , to see whether their distributions were different across the subsample which included HHI versus the one that did not. All three tests fail to reject the null hypothesis of distributional equality at a conservative 90% confidence level. P-values for KS tests on the variables P , E , and A were 0.21, 0.48, and 0.23 respectively. The size of the subsample which included w_i was (rounding to the nearest 10) 2,800, and the size of the subsample which did not include w_i was 2,180.

Figure 7. Goodness of Fit of the Type Distributions

equal to 0.0055, which is roughly 1% of the standard deviation of $\log(q)$. The resulting differences in the type distributions are shown in Figure 7. The fact that the CDFs derived from the B-spline (e.g., Minority Type, B-Spline) lie directly on top of those derived from direct imputation (e.g., Minority Type, Direct) verifies the exceptionally good fit provided by the B-splines.

Appendix D. Determinants of q

If our exclusion restriction is violated and race directly influences HHI (i.e., racial animus has a significant effect), then our analysis will load the effects of this animus onto the estimated values of q for the minority students. If this is the case, then it may be possible to predict q using minority status in combination with other demographic traits. Of course, understanding the determinants of q is also of independent interest: if it is related to other factors known to play an important role in educational decisions and success, that would lend credibility to the interpretation of types as real learning costs. One method for assessing the drivers of q is to regress it on other plausibly exogenous demographic and socioeconomic traits contained in the B&B data.

In this section, we work with $\ln(q)$ as the outcome variable of interest. Note that our estimated \hat{q} is not a mechanical proxy for grades and SAT scores. The reason why stems from the idea behind the empirical auctions techniques we use to recover them. Simply stated, the equilibrium of a contest gives rise to a map between unobserved type q and observed HC choices s . Since we can observe the aggregate distribution of competitors' HC choices, the empirical model reverse-engineers one's unobserved learning cost q as that which rationalizes s as a best-response to competitor achievement levels (a related idea in the auctions literature was first proposed by GPV [11]). The mapping between s and q (i.e., our control functions) varies across race groups due

Table 13. Descriptive Statistics

Variable	Mean	Median	Std. Dev.
BACCALAUREATE AND BEYOND			
Minority Dummy: M	0.118	0	0.322
Gender Dummy:	0.578	1	0.494
Risk Index (7 Point Scale):	0.3139	0	0.536
Parent Adj. Gross Income (\$K) :	42.4	39.1	24.7
Mother College Grad.:	0.244	0	0.429
Father College Grad.:	0.399	0	0.490
Parents Married:	0.760	1	0.4273
# Members of Family:	3.97	4	1.27
Parental Cash Savings (\$K):	5.81	0.83	22.1

to AA and other factors—e.g., the quality distribution of college seats and the shock variance—that shape the college admissions contest. Grades and SAT scores are the result of a mixture of exogenous student characteristics (potentially unobserved by the econometrician) and features of the contest that incentivize hard work. The auctions tools filter out the latter (contest-based incentives) so that only the former remains.

Since q was inferred from the first-order condition facing the agents, the meaning of q and what determines it can be difficult to see at this abstract level. Our goal here is to regress $\ln(q)$ on plausibly exogenous drivers of learning costs to better understand the real-world meaning and causes of the student-specific fixed effect we recover from our empirical model. Given our interpretation of q as student characteristics that are influenced by forces internal to the individual (e.g., innate ability) and external conditions (e.g., primary/secondary school quality, parental education and financial resources), the most plausibly exogenous explanatory variables describe properties of the students' parents or family. The variables we consider are summarized in Table 13. The data set includes the state of residence of the student at enrollment, and this has been condensed into regional dummy variables. We consider only students that have nonmissing values for all of these variables.

The variable risk index is the sum of seven binary risk factors for becoming a future college dropout such as receiving a GED as opposed to a traditional diploma. Mother college graduate and father college graduate are dummies that indicate whether the respective parent graduated from college. Parents married is a dummy variable indicating whether both parents are married (but perhaps not to each other). Parental cash savings refers to the amount of cash and savings reported on financial aid forms (submitted at college enrollment).

We include the variables in Table 13 as well as a dummy for minority status and a full set of interactions with the minority status dummy in our regressions. Our regression results are contained in Table 14. Conditional on the various family background and socioeconomic controls, the point estimates on the minority variables are all individually insignificant in a statistical sense. Moreover, an F-test at a conservative 90% confidence level fails to reject the null hypothesis that the minority dummy and the minority interaction terms play no role as a group. On the

Table 14. Estimates of $\log(q)$

Parameter	Point Est.	Std. Err.
Minority Dummy:	0.0894	0.419
Gender Dummy:	-0.0855	0.0240
Gender Dummy x Minority Dummy:	-0.0224	0.0787
Risk Index > 0:	0.113	0.0273
Risk Index x Minority Dummy:	0.0792	0.829
Parent Adj. Gross Income :	$5.42 \cdot 10^4$	$5.62 \cdot 10^4$
Parent Adj. Gross Income x Minority Dummy :	$1.39 \cdot 10^4$	$2.20 \cdot 10^3$
Mother College Grad.:	$7.75 \cdot 10^3$	0.0314
Mother College Grad. x Minority Dummy:	$7.60 \cdot 10^3$	0.102
Father College Grad.:	-0.152	0.028
Father College Grad. x Minority Dummy:	-0.0407	0.0841
Parents Married:	-0.0803	0.0341
Parents Married x Minority Dummy:	0.144	0.0932
# Members of Family:	$2.65 \cdot 10^3$	0.0104
# Members of Family x Minority Dummy:	-0.0258	0.0824
Parental Cash Savings > 0:	-0.0894	0.0382
Parental Cash Savings > 0 x Minority Dummy:	-0.0144	0.0816
Region Dummies	Included	
Region Dummies x Minority Dummies	Included	
N:	1196	
R ² :	0.451	

NOTE: Significance at the 10%, 5%, and 1% levels is denoted by *, **, and ***, respectively.

other hand, several other variables that one would expect to influence learning costs are highly predictive of q and have the sign one would expect, including the student's risk index, the parents' marital status, whether the father has graduated from college, and whether the parents have liquid savings. Overall, our set of family background and socioeconomic controls explain 45% of overall variation in our estimated student types q .

It is interesting to note that in Section 5.5 we estimate a substantial race gap in learning costs, whereas through the current exercise we can conclude that the race gap in costs disappears when home environment controls are used. Our analysis implies that the black-white HHI gap in our data is due to the role of learning-costs in determining HHI combined with a black-white gap in the distribution of learning-costs. If racial animus played a significant role in the HHI gap (contrary to Assumption 4.4 and Fryer et al. [10]), then one might expect to find that the minority dummy is a significant predictor of learning-costs. However, our analysis suggests that inter-group differences in learning costs are not driven by minority status once one controls for other demographic traits. In short, we do not find evidence that our exclusion restriction is violated and causing the effect of racial animus to be loaded onto our estimates of q .

Appendix E. Solving the Model Numerically

Counterfactuals that compare different AA regimes necessarily require solving the model to determine the student behavior (i.e., human capital investments), which in turn determines college enrollments and final outcomes. There are two endogenous objects that need to be computed for each demographic group: the student strategies and the assignment functions that map NHC to college assignment.

Although we prove an equilibrium exists, we do not have a result that the equilibrium is unique. We used multiple restarts to test for the existence of multiple equilibria, but we never found multiplicity. This is perhaps not terribly surprising since Bodoh-Creed and Hickman [4] prove that the equilibrium of an analogous color-blind or quota model without a matching shock has a unique equilibrium since the FOC defines an ODE that has a unique solution. The difficulty of proving uniqueness for models with a matching shock is that the ODE defined by the FOC is nonlocal.⁵³ To the best of our knowledge, the study of nonlocal ODEs and proofs of properties like uniqueness are limited to special cases that our model does not fall under. Please see Bodoh-Creed and Hickman [4] for details.

We numerically approximate all four objects: two group-specific strategies and two group-specific inverse assignment functions. B-splines were used for the numerical approximations because these functions allow for accurate approximations of both a function and its derivatives using relatively few parameters. We chose to approximate the inverse assignment function because while the range of these functions (the colleges) are known ex ante, the domain (the noisy human capitals) is endogenous. We used six knots in our approximation, but found that using up to 20 knots had a negligible effect on our results. We insist that the strategies be consistent with the first order conditions for the student's decision problem in the AA regime of interest (Equation 4). The assignment functions (i.e., the inverse of the inverse assignment function approximated by our B-splines) are required to be consistent with the human capital choices of the students.⁵⁴

We solved the model using an optimization algorithm that tries to minimize the inconsistencies between the approximated objects and the theoretical analogs described above. The optimization algorithm adjusts the variables describing all four of the numerical approximations simultaneously. Our first step within each iteration is to compute the induced inverse assignment function for each group, which is simply the inverse of the assignment function generated by the approximate student strategies and the contest structure generated by the form of AA we are studying. We then use an L_2 penalty function for the distance between the B-spline fit of the inverse assignment function and the induced inverse assignment function, which is computed at 50 points evenly

⁵³By nonlocal, we mean that the derivative of the solution at a given point depends on the value of the solution at many other points. This occurs in our setting because of the expectations in the FOC. The derivative of a local ODE, the most familiar form, only depends on the value of the function at that point.

⁵⁴We solved for the status-quo system as a quality assurance check on our algorithms. Because quota and admissions preference schemes are outcome equivalent (Bodoh-Creed and Hickman [4]), we treated the status quo admissions preference scheme as quota wherein each group competed separately for the distribution of seats allocated to members of that group in the status quo.

Figure 8. Single Index Function

spaced across the support of the distribution of college qualities. The measure applied in the L_2 norm is the estimated CDF of the distribution of college qualities available to that demographic group.⁵⁵

We now turn to our metric of the inconsistency between the approximate and exact equilibrium strategies. Our third step within each iteration is to compute the assignment function implied by our B-spline of the inverse assignment function. The fourth step of our algorithm is to calculate the first order condition given the assignment function for each group implied by the associated B-spline of the inverse assignment function and the B-spline of the strategy for that group. Our error function was an L_2 penalty function enforced at 50 evenly gridded values over the support of q . The measure used in the L_2 -norm is the type distribution of the respective group.

The complete objective function for the optimization problem is the sum of the penalty functions for the inverse assignment mapping and the first order condition, which we chose for simplicity and the fact that the weighting did not seem to significantly affect the optimization results. If an optimal value of 0 is found, then the approximated inverse assignment functions and the approximate strategies are consistent with the first order conditions of the decision problem. In addition, an optimal value of 0 implies that the approximation strategies and inverse assignment functions are also consistent with each other as required by equilibrium. We used multiple restarts to test for multiple equilibria and never found multiplicity.

Appendix F. Additional Tables and Figures

F.1. Point Estimates: HC Single Index Function. Figure 8 illustrates our single index equation. Each line depicts the effect of academic record a on the HC index while holding exam score e fixed at one of its quartiles. The lesser importance of e relative to a is reflected in the fact that the difference between the 75th and 25th percentile lines is less than 0.1, while the difference between the 75th and 25th quintiles on the line describing $S(e_{\text{median}}|a)$ is close to 0.2. The upward

⁵⁵In the color-blind and proportional quota counterfactuals, this is the total measure of college qualities. In the status-quo model, the distribution of school available to students in a given demographic group is equal to the distribution of college qualities in which those students enroll.

Figure 9. Graduation Probability Function

Table 15. Counterfactual Minority HHI by Achievement Quintile

Learning-Cost Type Tier	Proportional Quota	Proportional Quota, Partial Eqm	Color-Blind	Color-Blind, Partial Eqm
Top Achiever Quintile	\$81,560	\$81,563	\$80,336	\$79,896
Second Achiever Quintile	\$77,223	\$77,212	\$74,714	\$74,351
Third Achiever Quintile	\$73,350	\$73,378	\$69,856	\$70,258
Fourth Achiever Quintile	\$69,143	\$69,143	\$65,264	\$66,154
Bottom Achiever Quintile	\$62,673	\$62,654	\$60,010	\$61,209

curve of the lines is a result of the convexity of the single index with respect to the student's GPA. We include 95% confidence intervals on the line describing $S(e_{\text{median}})$ at each decile of the distribution of a . The distance between the line representing the 75th percentile of e and the maximums of e is due to the convexity of the single index in that variable. The large distance between the line representing the 25th percentile of e and the minimums of e is due to a long tail of low value of e within the data.

F.2. Point Estimates: Graduation Probability Function. Figure 9 illustrates the relative impact of p and s on graduation probabilities. Each line of the plot depicts the effect of HC on graduation probability holding that college quality fixed at one of its quartiles. We have also plotted 95% confidence bounds on graduation rates at a college of median quality at the deciles of the distribution of HC. The convexity of $r(p, s)$ is evident from the increasing difference between the lines as the college quality improves, meaning that college quality matters most above the 75th percentile of college quality. The complementarity between p and s results in an increasing spread between the lines as s increases, meaning that college quality is more important for higher achieving students.

F.3. Noisy Human Capital Markup Function Equivalent to a Proportional Quota, Φ . Figure 10 describes the estimated markup function, $\Phi(t)$, that is equivalent to a proportional quota. In

Figure 10. Noisy Human Capital Markup Function Equivalent to a Proportional Quota, Ψ

other words, an admissions preference system using the markup function would yield 20% of the minority students in each quintile of the college quality spectrum. The horizontal axes of both panels display quantile ranks of NHC for nonminority students. The markup of minority student NHC required to mimic the effects of a proportional quota is, as intuition would suggest, larger than in the status quo admissions preference system (see Figure 2).

The left panel describes the shape of Ψ . The vertical axis displays the quantile rank of subsidized NHC within the non-minority NHC distribution. If a minority student has an NHC at the quantile rank marked on the horizontal axis, the student gets the same college assignment as a nonminority student with an NHC at the quantile rank denoted on the vertical axis. For example, a minority student with an NHC equal to the median of the nonminority population gets the same college assignment as a nonminority student at the 73^d percentile of the nonminority population. The dashed line denotes the 45° line for reference. The right pane of Figure 2 describes the markup function in terms of school quality. The vertical axis denotes the gap in the quantile rank of college quality between a minority student and a nonminority student at each NHC quantile. For example, the plot shows that if two students from different groups both have an NHC value equal to the median of the nonminority population, then the minority student is assigned to a school whose quantile rank is 0.23 higher in the school quality distribution.

F.4. Raw Strengths of Productive and Competitive Channels. Figure 11 displays the marginal effect of the productive and competitive channels on agent utility (as opposed to the ratio to the total effect as per Section 6.1). The analogous plot for nonminorities is very similar.

In the first-best world we hold the enrollments fixed at the status-quo level and let the students choose their optimal amount of HC, given their assigned college.⁵⁶ Because enrollments are fixed,

⁵⁶Note that in holding enrollments fixed, we are essentially assuming that affirmative action policies and matching shock realizations are fixed. A natural intuition is that the first-best matching is perfectly assortative between student-types and college qualities, but this would imply that the affirmative action policies are not welfare enhancing for colleges and that the matching shock does not represent some (exogenous) aspect of the applicant that is valued by

Figure 11. Raw Marginal Effect of Each Channel

the competitive channel vanishes in the first best. To understand why the productive channel is almost identical in the status quo and first best, recall that under our point estimates HC influences $U(p, s, q)$ essentially only through the graduation probability function and that the graduation probabilities are almost linear in HC (see Figure 9). As a result, the marginal effect of HC on utility in the first best (i.e., the productive channel of incentives in the first best) is essentially identical to what we find under the status quo, despite the fact that the level of HC investments in the first-best are significantly lower than in the status quo.

Appendix G. The Mismatch Hypothesis

A prominent vein of empirical literature on college admissions focuses on phenomena where college candidates place at a school either above or below their quantile rank in HC space. Perhaps the most closely related paper to our stage I estimator is Dillon and Smith [8, DS17], who studied the degree to which over-match—a student placing at a college above his quantile rank—and under-match—a student placing at a college below his quantile rank—occurs in US college admissions using the National Longitudinal Survey of Youth 1997 cohort. Like our study DS17 uses multiple proxies for college quality and condenses them to a single index in order to produce a data-driven vertical ordering of colleges.

However, three key differences exist between their study and ours which preclude a tight comparison of their empirical estimates to ours. First, they included the universe of enrollees at either

colleges. Since our main purpose here is to further investigate the role of competition in human capital investment, we avoid taking a stand on the value of affirmative action or the interpretation of the matching shock.

2-year or 4-year colleges, whereas the B&B includes only graduates at 4-year colleges. This makes comparisons hard since some of the students who begin at 2-year colleges later graduate from 4-year colleges and some never attain bachelor's degrees. Second, they base their human capital measure on an exam—the Armed Services Vocational Aptitude Battery (ASVAB)—whereas we base ours on an exam score—the SAT—and proxies for high school coursework and grades—non-major GPA and two binary variables, REMEVER and MATHPREP described in Appendix H.4.1 below. Third, their analysis is descriptive in nature—they condense the multi-part ASVAB to a single index using principal component analysis and then compare students' quantile ranks of HC and their college quality index. In contrast, our stage I estimator is structural in nature, deriving a single index s of HC by using a formal model of its role within the contest structure and the graduation probability function. This is because our stage I estimates must reflect information about effort incentives that we use in stage II to filter out the exogenous student fixed-effect q from the endogenous HC index s (see discussion on this point in Appendix D above). This process allows us to compute counterfactuals that control for endogenous changes to HC under alternative college admissions schemes. Another benefit of our structural approach is that it allows us to explicitly decompose over- and under-match into matching shocks versus race-based AA. Despite the differences between our stage I estimator and DS17, both studies broadly agree that there is a significant level of over- and under-match in the US college market.

Another important strain of the empirical literature on AA focuses on the mismatch hypothesis or the idea that AA may cause minority students to be placed in higher quality colleges, but then graduate with lower probability. The core idea behind the mismatch hypothesis is that students with a level of HC lower than their peers might have worse outcomes than if they had been matched to a lower quality school. Since higher quality schools enroll students with higher levels of HC, this would manifest empirically as a negative interaction between the minority student's HC and the school's quality. To the best of our knowledge, all of these papers assume student HC is exogenous and fixed regardless of the AA scheme at work. Some empirical studies have found evidence of mismatch (e.g., Loury and Garman [17] and Arcidiacono, Aucejo, and Hotz [1]), while other empirical work has found evidence that the mismatch problem is small and likely outweighed by other benefits of higher-quality placement (e.g., Long [16], Rothstein and Yoon [19], and Chambers, Clydesdale, Kidder, and Lempert [6]). Other evidence suggests that all students generally benefit from attending higher-quality schools (e.g., Dillon and Smith [9] and Badge, Epple, and Taylor [2]).

Our model contributes to this literature by studying traditional forms of mismatch effects in terms of graduation probabilities and HHI, but our model allows for a more exotic variety of mismatch: the idea that minorities could also be hurt on net by AA through changes to their endogenous choice of HC. We find that, relative to the standard metric of graduation rates, there is no clear-cut ranking between enrollment schemes in general since the sign and magnitude of the graduation rate change varies with endogenous HC choices and college assignment. While the status quo AA leads to a small reduction in average minority graduation rates relative to color-blind admissions (about 0.3%), proportional quotas (a more generous form of AA) would

lead to an increase in minority graduation (about 0.9%).⁵⁷ On the other hand, if we adopt a more holistic metric we find that the total economic gains to minorities from AA, including graduation rates and post-college HHI, are unambiguously positive.

In our context, the mismatch hypothesis would be verified if either the graduation probability function or the HHI equation were single peaked in school quality p for one or more value of human capital s . For example, if the graduation probability function exhibited such a single-peaked shape, then students with human capital s would maximize their graduation probability by enrolling in a college of sub-maximal quality p . Since the parametric form of the graduation probability function is a third-order polynomial, it is entirely possible for a single-peaked shape to be recovered from the data. We checked for the presence of non-monotonicities in the graduation probability function across a finely-gridded matrix of school quality and HC values and did not find any points where the estimated graduation probability was decreasing in school quality (see Figure 9).

Our HHI equation is, by construction, monotone in school quality. To detect mismatch in terms of HHI, we repeated our stage II estimates using the alternative HHI regression equation: (20)

$$\log(w_i) = \log(a_0) + a_p \log(p_i) + a_s \log(s_i) - a_q y(S, D_{M i}; a) + a_M 1f s_i < \bar{s}(p_i) g(\bar{s}(p_i) - s_i) + \#_{w_i},$$

where $\bar{s}(p)$ is the average HC across all demographic groups at colleges of quality p and a_M is a coefficient determining the importance of mismatch in the HHI equation. As before, $q_i = y(S, D_{M i}; a)$ is the learning-cost inferred from the student's demographics and choice of HC.

If the mismatch hypothesis is correct in terms of HHI, one would expect to estimate a significantly negative value for a_M . Instead, what we find is that a_0 , a_p , a_s , and a_q are essentially unchanged from the benchmark point-estimates of Table 5, while a_M is 1.432×10^{-6} . Since a_M is economically insignificant, we conclude that the mismatch hypothesis is not supported by our model estimates.⁵⁸

Appendix H. Robustness Checks

The robustness checks are presented in the order that they are referenced in the main paper. They appear in the following order, where the numbers refer to subsections:

- H.1 Estimation without high school grade proxy
- H.2 Estimation using only the subsample of male salaries
- H.3 Model extension and estimation with regional higher-education markets
- H.4 Estimation assuming a noisy measurement of human capital
- H.5 Estimation under an animus-based penalty on minority HHI
- H.6 Estimation given higher matching shock variance
- H.7 Estimation dropping low-quality colleges

⁵⁷Because of the endogeneity of HC, it is not obvious that the status quo system generates outcomes “in between” the color-blind and proportional quota systems.

⁵⁸We also experimented with a mismatch term of the form $a_M 1f s_i < \bar{s}(p_i) g(\bar{s}(p_i) - s_i)^2$ and obtained essentially identical results.

H.8 Partial adaptation to AA rule changes.

H.9 Other robustness checks

H.1. Estimating Without High School GPA Proxy. There is broad agreement in the literature that both standardized exam scores and high school GPA together are significant predictors of college performance, with each containing some information that the other does not. Most debate within the literature focuses on the relative importances of these two measures for predicting college success (e.g., Rothstein [18] and Cohn et. al. [7]). In Section 2.2 we argued that nonmajor grades are a reasonable proxy for the high school grades of the students in our data sample. In addition, we estimated that the correlation between school quality and major grades is -0.006 , which strongly suggests there is no relationship between grades and college quality (e.g., grade inflation) that would mechanically replicate the assortativity between HC and college quality predicted by our contest model. Finally, the standard deviation of within-campus SAT scores averaged across colleges is 194 points (i.e., roughly 80-90% of the standard deviation of all SAT scores), and so one should not expect this variable to explain college assignment very well on its own. Despite all of this, one could conduct stage I of our estimation using standardized exam scores as the only measure of pre-college HC.

As a first attempt, we consider a model where equation (5) is set to $s(E) = E$.⁵⁹ One advantage of this model is that since we have assumed the value of b^s , we can estimate the remaining parts of stage I separately. We rejected this model specification because the R^2 of the graduation probability function is only 0.477, whereas our benchmark model yields an $R^2 = 0.566$. We also found that the model with $s(E) = E$ produced a matching shock variance of $s_{\#} = 0.215$, which is an order of magnitude larger than our benchmark estimate. To put these parameter estimates in context, they imply that the noise to signal ratio inherent in college assignment is an implausibly high value of 1.47, which means that luck is nearly 50% more important than effort for determining college placement. We view this as an implausibly low level of assortativity.

Since our baseline model specification is a quadratic polynomial in E and our high school GPA proxy A (with no intercept term), we also estimated a quadratic polynomial $s(E) = a_1E + a_2E^2$ for the single index equation (5). This model again requires estimating b^r and b^s simultaneously, which we did using the GMM estimator described in equation (17). The resulting value of the matching shock variance we recover is $s_{\#} = 0.213$, which again implies very low levels of assortativity. However, the R^2 of the graduation probability is only 0.509, which still falls short of our benchmark model's ability to predict variation in graduation rates. These parameter estimates imply that the noise to signal ratio inherent in college assignment is still an implausibly high value of 1.48, which again means that luck is 50% more important than effort for determining college placement.

Hoping to give the SAT-only model the best possible demonstration, we experimented with re-tuning to force the solver to place more emphasis on the assortativity component of our objective

⁵⁹Our linear specification did not include a slope parameter because any slope coefficient would be offset by compensating scale shifts in the variance of the matching shock.

function. We find that $s_{\#} = 0.025$, which is closer to our benchmark level, but the graduation probability t is now $R^2 = 0.476$, which is on par with the predictive power of the linear model.

The large gap in R^2 of the graduation probability function t between the benchmark model and the SAT-only models estimated here is problematic. Our estimated single index is the most important driver of graduation probability, which suggests the single index actually does carry information about underlying HC. The natural interpretation of the gap in R^2 is that using the SAT-only model leaves an important fraction of HC unobserved. Moreover, the small estimated value of $s_{\#}$ in the benchmark model suggests there is little unobserved HC left to detect, a conclusion that is also supported by our failed attempt to find additional informative variables relative to the benchmark model specification (see Appendix H.4.1).

One methodological concern about the SAT-only estimator is that it is very sensitive to the tuning parameters of our solver algorithm, whereas our benchmark stage I estimator is not. However, the sensitivity of the SAT-only model is not surprising since the estimator is trying to achieve two goals (predict graduation rates and maximize the fit of the contest model to the data) with only one tool, the SAT score. These dual objectives are in tension with one another, and the tuning parameters play an outsized role in which of the two goals the estimator achieves. In contrast, the benchmark estimator can use both SAT scores and our GPA proxy to achieve these two goals without any need to tune the relative importances of the objective function. In other words, the data and not the tuning parameters dictate the parameter estimates.

H.2. Estimating Stage II on the Male Student Subsample. Our benchmark estimates of stage II are based on the HHI of male and female students. Although we believe that HHI provides a more complete picture of the returns to college than individual-level income measures, much of the literature focuses on the salaries of men as the basis for measuring returns to college. In this appendix, we re-estimate our model using an outcome measure comparable to previous literature, and focus our analysis (where we can) on individual income for male students in the B&B sample.

Stage I estimates various parameters of the nation-wide contest for college admissions. We cannot simply omit female students from this analysis as the choices of women play a crucial role in determining the outcomes of this contest. Instead, we take the stage I estimates as given, and then estimate stage II using only the male students. In such an analysis, the incentives facing the male students (and hence the first-order condition that defines the value of q we attribute to each male student) are shaped by the choices made by both male and female students in vying for college seats. However, the estimates of the parameters of the HHI equation are based solely on the individual-level income of the male students.

Table 16 displays the values of the HHI equation (i.e., a) that we recover if we restrict attention at stage II to male students and use their individual-level total income 10 years after graduation as the dependent variable in the regression. All of the parameters are within the confidence bounds provided for our benchmark estimates (see Table 5). We obtain similar results if we use salary 10

Table 16. Estimates of the Wealth Equation for Male Students

Parameter	Point Est.
a_p , College Quality	0.1868
a_s , Human Capital	$3.115 \cdot 10^{-6}$
a_q , Learning Cost	0.1085
a_0 , Intercept	\$72,162

years after graduation as the dependent variable. Therefore, we conclude that our results are not driven by our focus on HHI or the inclusion of both men and women in our sample.

H.3. Regional Higher Education Markets. We also experimented with models that divide the United States into geographically distinct regional markets. Of course, the size of our data set places bounds on how narrowly we can define the catchment areas for each market, particularly given that we need to continue to treat minority and nonminority students as distinct groups. In the end, we divided the country according to the four US Census Bureau regions: Northeast, South, Midwest, and West. We refer to each of these geographic areas as a region and index them by g . As we discuss below, the results of stage I necessarily change somewhat, but the results of stage II are almost unchanged.

We define the market for each of the four regions by the colleges within that market. This means that students participate in the market where they enroll, not the market they resided in prior to matriculating. We divide the students attending college within each region into four groups based on their minority status and whether they resided in the college's region prior to matriculating. Therefore, the groups are comprised of minorities that attend college within their region of prior residence (M_{In}), minorities that attend college outside their region of prior residence (M_{Out}), nonminorities that attend college within their region of prior residence (N_{In}), and nonminorities that attend college outside their region of prior residence (N_{Out}).⁶⁰ We do not seek to explain why a particular student enrolls in a university outside of his or her home region and merely take the choice as given. We refer to a generic demographic group (e.g., N_{Out}) as a group and to a generic group in a generic region as a region-group. The size of group j in region g is $m_{g,j}$ where $\sum_j m_{g,j} = 1$.

We provide a formal description of our estimators in Sections H.3.2 and H.3.3, but first we provide an informal description. We repeat our stage I analysis within each region, which includes fitting the distributions of school quality and HC for each group within each region, which yields 16 distinct distributions for each of school quality and HC (as opposed to only two in our benchmark model). We assume that the single index function for HC, the graduation probability function, and the matching shock variance are the same across regions. Implicitly, we are assuming that HC is not specific to a particular region, colleges are similar enough that the

⁶⁰Obviously this definition is not perfect. For example, students that reside in the periphery of a region might attend a school that is technically outside of that region and yet relatively close to their home.

Figure 12. Single Index Function, Regional Model

“graduation production technologies” used in different regions are roughly the same, and the source(s) of the matching shock are similar across regions. The graduation probability function is chosen to match the college-campus-level IPEDS race- and college-specific six-year graduation rates (as per the benchmark model), and the single index is chosen to maximize the assortativity of the match. We compute the Kendall's τ in each of the 16 region-groups separately, and the assortativity moment is the average of the region-group Kendall's τ measures weighted by $m_{g,j}$. Since the IPEDS data does not allow us to distinguish students simultaneously according to their state of residence prior to college and their minority status, we estimate $m_{g,j}$ from the B&B data.

Recall that we estimate a quota model in which members of each region-group participate in a region-group specific contest, which is equivalent to the admissions preference systems used in practice. In principle, we estimate three AA policies per region, all of which are defined relative to how the members of N_{In} are treated within that region. As per our benchmark model, we compute a markup function $\Phi_{M_{In}}$ that defines how members of N_{In} are treated relative to M_{In} . In addition (and unlike the benchmark model), we also estimate $\Phi_{N_{Out}}$ and $\Phi_{M_{Out}}$, which describe how members of N_{Out} and M_{Out} are treated relative to N_{In} .

We make two modifications to our stage II estimator to account for regional differences. First, we infer the learning-cost for each student in the B&B data using region-group-specific strategy functions, meaning that the control function in our Mincer equation is now region-group specific. The region-group specific strategies reflect the distribution of college quality available to that region-group and the distribution of HC of the other students within that region-group. Second, we allowed the constant total-factor-productivity term in α to vary by region in order to account for region-specific differences in the employment markets faced by the students.

Figures 12 and 13 provide plots of the single index function and the graduation probability function (see appendix F for a comparison with the equivalent images for the benchmark model).

Figure 13. Graduation Probability Function, Regional Model

Table 17. Estimates of the Wealth Equation from Regional Model

Parameter	Point Est.
a_p , College Quality	0.1601
a_s , Human Capital	$2.864 \cdot 10^{-7}$
a_q , Learning Cost	0.0319
$a_{0,Northeast}$ Intercept for Northeast Region	\$105,271
$a_{0,Midwest}$ Intercept for Midwest Region	\$96,568
$a_{0,South}$ Intercept for South Region	\$96,059
$a_{0,West}$ Intercept for West Region	\$104,738

The single index function is very similar to that estimated under our benchmark model. The graduation probability still displays a dominant role for s in determining graduation rates, but the role of p is now somewhat larger.⁶¹ Table 17 displays the results of the stage II estimates for the regional model. a_p is slightly larger than the benchmark estimate, a_q is slightly smaller than the benchmark estimate, and both are within the 95% confidence bounds of the benchmark estimates. Moreover, p and q are still the primary drivers of HHI, while the role of s remains insignificant.

We now describe the pros and cons of the regional model relative to our benchmark. Because of the different strategies in each region, we obtain identification of $u(p, s, q)$ both from the differences in the strategies adopted by a particular demographic group across regions as well

⁶¹The regional model does not improve the fit to the IPEDS graduation probabilities relative to the benchmark nation-wide model.

as the difference in strategies between minorities and nonminorities within each region.⁶² For example, the nonminority students that reside in each region (i.e., the members of N_{In} within each region) face four different contests based on the different distributions of college qualities and contestant HC choices in each of their four distinct contests. This means that we can estimate four region-specific strategies mapping learning-cost to HC based only on the students in N_{In} within each region. Doing this requires an exclusion restriction. If we are willing to assume that the HHI parameters a_p , a_s , and a_q are constant across regions, then this suffices.

In Appendix H.3.1, we conduct a further robustness check by redoing this estimation exercise using only inter-regional variation of the strategies adopted by the nonminority student subsample. The fact that our estimates from stage II of the regional model using only nonminority students and under the benchmark model without regions are quite similar lends credibility to our benchmark results. In particular, the fact that our estimates are insignificantly different when conducted using only nonminority students implies our results are not driven by a violation of our benchmark exclusion restriction (Assumption 4.4).

However, there are two significant downsides to the regional model that in our opinion overwhelm the benefits. First, there is very limited data coverage in the B&B survey for minority students in some regions. For example, only 80 minority students in the B&B data attend college in the Midwest, and 11 of these students resided outside the Midwest prior to college. This paucity of data on minority students is typical outside of the South and would lead to nontrivial increases in the standard errors of our estimates. Second, there is ambiguity about how to define the AA counterfactuals. In principal, students might move from one geographic region to another if the AA policy change, making it unclear how to appropriately compute counterfactuals. Note, however, that this concern does not affect our estimates of the structural parameters.

H.3.1. Identification from Inter-regional Incentive Differences Our benchmark model was identified from the different strategies utilized by the minority and nonminority students, which broke the otherwise perfect correlation between HC choices and type q . Under a regional model, the nonminority students in different regions face differing incentives based on the regional differences in school quality available to nonminority students in that region and the different distributions of HC chosen by nonminority students in different regions.

As mentioned above, we estimate eight different strategies, one for each of N_{In} and N_{Out} in each of the four regions. Each of these strategies provides a different mapping between learning cost and HC that we can use to break the otherwise perfect correlation between these two variables. The identification argument is essentially identical to that used in Proposition 4.10, but the inter-regional incentive differences play the role of the incentive differences generated by AA in our benchmark model. The key exclusion restriction is that the HHI ten years after graduation conditional on (p, s, q) can only vary in terms of the level parameter (i.e., $a_{0,West} \neq a_{0,South}$ is

⁶²We also obtain identification power from the differing strategies of the students that reside within that region prior to university relative to those that do not.

Figure 14. Nonminority Agent Strategy by Region

Table 18. Estimates Using Nonminority Students and Inter-Regional Variation Only

Parameter	Point Est.
a_p , College Quality	0.1613
a_s , Human Capital	$5.396 \cdot 10^{-10}$
a_q , Learning Cost	0.0294
$a_{0,Northeast}$ Intercept for Northeast Region	\$105,654
$a_{0,Midwest}$ Intercept for Midwest Region	\$97,234
$a_{0,South}$ Intercept for South Region	\$96,352
$a_{0,West}$ Intercept for West Region	\$104,131

allowed, but a_p , a_s , and a_q are identical across regions). The resulting strategies for the nonminority students enrolling in a college located within their region of residence before college is displayed in Figure 14.

When we include only nonminority students in our HHI estimator (equation 19) after the modifications to the control function described below in Appendix H.3.3, we find the estimates of a shown in Table 18. They retain the qualitative features of Tables 5 and 17 that the meaningful drivers of HHI are school quality and learning cost. More importantly, since these estimates are based solely on the nonminority students, this analysis implies that the importance of a_q in our benchmark analysis is not the result of a violation of the exclusion restriction used in the benchmark model.

H.3.2. Technical Details of the Stage I Estimator. We do not wish to rehash the entirety of Appendix B, so we focus our discussion on the differences between the estimation of the regional model and the benchmark model discussed in that appendix.

We use a combination of copulas and B-splines to model the joint and marginal distributions of college graduates within each region-group in a fashion that is identical to that described in Appendix B. The primary changes occur in the definition of the GMM objective function, as some moment conditions become more complicated in the regional model extension. In our benchmark estimator, D_{M_i} denotes a dummy variable indicating that student i described a minority student. We now use subscripts g, j to denote a variable refers to a student from demographic group j that resided in geographic region g prior to matriculating. For example $D_{g,j,i}$ is a dummy variable equal to one if student i is a member of group j attending college in region g . In the equations below, g ranges over the four geographic areas and $j \in \{M_{In}, M_{Out}, N_{In}, N_{Out}\}$. With this notation in hand, we can describe the objective function as:

$$\begin{aligned} \hat{p}_{\hat{s}\#} = \arg \min & \left(\sum_{i=1}^h \sum_{g,j} \hat{a} \hat{a} D_{g,j,i} F_{P_{g,j}}(p_{ij} \text{grad}; g_{g,j}^p) \hat{P}_{P_{g,j}}(p_{ij} \text{grad}) \right)^{i_2} \\ & + \sum_{k=0}^{100h} \hat{a} Q_{S_{g,j}}(r_{g,j,k} \text{grad}; g_{g,j}^q) \hat{Q}_{S_{g,j}}(r_{g,j,k})^{i_2} \\ & + \sum_{i=1}^h \sum_{g,j} \hat{a} \hat{a} D_{g,j,i} m_{g,j} F_{P_S}(p_i, s_{ij} g, j, \text{grad}; g_{g,j}^p, g_{g,j}^q, n_{g,j}) \hat{P}_{P_S}(p_i, s_{ij} g, j, \text{grad})^{i_2} \\ & + \sum_{l=1}^L [\hat{G}_{M_l} - Z_{M_l} b^r]^2 + [\hat{G}_{N_l} - Z_{N_l} b^r]^2 \\ & + \sum_{g,j} \hat{a} m_{g,j} \frac{n_{g,j} - 1}{n_{g,j}} \left(\frac{n_{g,j}(s\#) - 1}{n_{g,j}(s\#)} \right)^{i_2} \end{aligned}$$

Subject to :

$$s_i = S(e, a_i; b^s), \quad i = 1, \dots, l$$

$$\frac{\partial S(e, a; b^s)}{\partial e} > 0, \quad \frac{\partial S(e, a; b^s)}{\partial a} > 0 \quad 8(e, a; b^s), \quad \text{and} \quad \max_{(e,a) \in \mathbb{R}^2} f S(e, a; b^s) g = 1$$

$$g_{j,k}^v = g_{j,k}^v, \quad k = 2, \dots, K_j^v + 3, \quad v = p, s$$

$$\min f G_1, \dots, G_L g \quad r(p, s; b^r) \quad 1 \quad 8(p, s)$$

$$\frac{\partial r(p, s; b^r)}{\partial p} > 0, \quad \frac{\partial r(p, s; b^r)}{\partial s} > 0, \quad 8(p, s)$$

$$Z_{dl} = [1, p_l, p_l^2, p_l^3, \bar{S}_{dl}, \bar{S}_{dl}^2, \bar{S}_{dl}^3, p_l \bar{S}_{dl}, p_l \bar{S}_{dl}^2, p_l^2 \bar{S}_{dl}], \quad d \in \{M, N\}, \quad l = 1, \dots, L$$

$$\bar{S}_{dl}^k \text{ agrees with equations (6) (8) (9), } k = 1, 2, 3, \quad d \in \{M, N\}, \quad l = 1, \dots, L$$

$$n_{g,j}, n_{g,j}, n_{g,j}(s\#), 2 \in [1, \infty); \text{ and } s\# > 0.$$

Note that the average HC levels at each school (e.g., \bar{S}_{M_l}) refer an average over all of the minority students at that college irrespective of where the students resided prior to college. We do not

separate \bar{S}_{dij} based on pre-college residence due to an inability to identify students simultaneously by minority status and state of residence in the publicly available IPEDS data.

Note that the t of the joint PDF (third line), the assortativity moment (4th line), and the t of $s_{\#}$ (5th line) are all weighted by the size of the respective region-group, $m_{g,j}$, whereas these terms are unweighted in our benchmark estimator. We weight these terms in the regional model to reflect the large variation in region-group size. We experimented with an equivalent weighting in the benchmark estimator, but found that it made only a trivial difference in the results.

Finally, we need to compute $m_{g,j}$. Our benchmark analysis was able to compute this variable exactly from the IPEDS data, but we cannot do so as our analysis separates students by minority status and region of pre-college residence using the B&B data. To estimate $m_{g,j,Grad}$ we first compute the sample weight of the college graduates in the B&B data in each (g,j) category (i.e., each region-group). We can then compute the fraction of college enrollees from each (g,j) category, $m_{g,j}$, by correcting for the average graduation probability in that category:

$$m_{g,j} = \frac{\int_{\underline{p}}^{\bar{p}} \int_{\underline{s}}^{\bar{s}} m_{g,j,Grad} f_{PS}(p, s; g, j) ds dp}{r(p, s)}$$

H.3.3. Technical Details of the Stage II Estimator The changes to the stage II estimator are easier to describe due to the simplicity of the notation. Equation (17) describes the process of fitting the first-order condition with B-splines. This is unchanged except for the fact that we now estimate one B-spline for each region-group pair for a total of 16 B-spline functions. The control function now takes the form

$$(21) \quad y_{s,j}; a, \hat{p}, \hat{s}_{\#}, f_{g,j}^s = \log \int_{g,j} \hat{a} D_{g,j,i} q_{g,j} s; l_{g,j}^s(a) .$$

Informally, the control function is simply the log of the q value dictated by the B-spline of the inverse strategy for the student's region g and demographic group j . Given this new definition of the control function, the definition of the regressors for the Mincer equation, $X(a)$, is unchanged. Finally, the form of the Mincer equation we use is identical to that described in equation (19).

H.3.4. Computational Issues, Regional Model Each stage of our estimator has to optimize over many more variables given the 16 region-groups, which generates a substantial increase in the computational burden. In each iteration of our stage I estimator objective function, we need to re-fit the marginal distributions of college quality and HC for each region-group, which requires eight times as many operations. In addition, due to the higher complexity of the parameter space, the optimizer requires several times as many steps to complete the stage I optimization routine for the b^S and b^r parameters. This means the entire runtime is roughly 20 to 30 times longer than the algorithms used to estimate the benchmark model. The second step of the stage I algorithm estimates $s_{\#}$. The code to estimate this variable in the regional model again takes roughly 10 times as long as the algorithms that estimate the benchmark model, since it is now optimizing over model-generated rank correlations for eight times as many joint distributions.

The more difficult computational problems occur in the stage II estimator. Because of the 16 region-groups, we need to use eight times as many B-Splines to characterize the control functions and the assignment mappings, and we also estimate three additional total factor productivity parameters in the HHI equation. Since the 16 region-groups are connected through common HHI parameters, this means the variables describing the equilibrium outcomes for each region-group interact with each other. As a result, the runtime for each objective function execution requires roughly 10 times more computations than the benchmark model, and roughly 40 times as many steps are required for the optimizer to converge.

H.4. Unobserved Human Capital. As noted in Section 3.1, the econometrician's estimate of the single index of HC is $\tilde{s} = s + h$ where s is the true underlying HC and h is the error in the econometrician's observation. Suppose the college's observed $t = s + \#$, where $\#$ reflects aspects of an applicant's profile other than HC that might be attractive to colleges. We can link the econometrician's and the college's observations through the equation $t = \tilde{s} + \# - h = \tilde{s} + f$. Given an estimate of the assignment mapping $P_j^r(t)$, for any pair (p, s) it is straightforward to compute $f = (P_j^r)^{-1}(p) - \tilde{s}$, but we do not have any variation in the data to determine what fraction of f is generated by h or $\#$. In our benchmark model, we have assumed that our single index captures all relevant aspects of the students' HC choice, which means that the only error observed by the econometrician is the matching shock (i.e., $f = \#$). At the other extreme, one might have assumed that the matching shock entirely consists of HC that is endogenously chosen by the student and observed by the college, but not the econometrician (i.e., $f = h$).

In this section, we first test whether we can detect additional HC unaccounted for by our baseline single index model by including additional relevant variables in equation (5). The fact we cannot would suggest that f is not unobserved HC. We then test various assumptions about the relative magnitudes of the standard deviations $s_{\#}$ and s_h ranging from our benchmark model ($s_{\#} = s_h, s_f = 0$) to the other extreme ($s_{\#} = 0, s_f = s_h$).

The key difference between the analysis presented here and that in the main text is that if $s_h > 0$, then some fraction of the observed f represents HC that should be included in the graduation probability formulas used in stage I and the Mincer regression in stage II.⁶³ Throughout we assume that $\#$ and h are normally distributed random variables with standard deviations $s_{\#}$ and s_h , although this assumption is made solely so that closed form solutions are available for the distribution of h conditional on an observed realization of f .⁶⁴

H.4.1. Detecting Additional HC Omitted from the Baseline Single Index Model. Our first test includes two additional variables in our HC index equation (5) that describe the students' pre-college coursework. The first variable, MATHPREP, is a dummy variable that is equal to one if the student did not take any basic math courses (either remedial or college-level algebra) during

⁶³We also experimented with including $\#$ as a fourth productive input in the HHI equation in stage II, but found it was economically insignificant. See Appendix H.9.

⁶⁴In principle, given any assumption about the distribution of either $\#$ or h , one could find the distribution of the other variable by deconvolving the empirical distribution of $f = \# - h$.

Table 19. ESTIMATES: Single Index Function and Matching Shock Variance

Parameter	Point Est.
b_1^s (e)	0.050
b_2^s (e ²)	0.1608
b_3^s (a)	0.3099
b_4^s (a ²)	0.2743
b_5^s (e a)	0.2365
b_6^s (MATHPREP)	1.8 10 ⁻⁵
b_7^s (REMEVER)	6.7 10 ⁻⁵
b_8^s (MATHPREP a)	6.0 10 ⁻⁴
b_9^s (REMEVER a)	5.6 10 ⁻⁵
b_{10}^s (MATHPREP a ²)	6.7 10 ⁻⁴
b_{10}^s (REMEVER a ²)	6.1 10 ⁻⁴

college. Since all of the colleges required at least a college-level algebra class at the time our data were collected, and since college-level algebra courses were commonly offered in most high schools as an optional elective, this dummy indicates a high level of pre-college math preparation and should be expected to enter with a positive coefficient in the single index. This variable has a mean of 0.485 for the B&B population. The second variable, REMEVER, is a dummy that is set to one if the student ever took a remedial class in any subject during college. Intuitively, this indicates a deficiency in pre-college preparation and should be expected to enter the single index with a negative coefficient. This variable has a mean of 0.085 for the B&B population.

The correlation between college quality and MATHPREP is 0.193, while the correlation between REMEVER and school quality is -0.110. This is what one might expect if better prepared students attend superior schools. In addition to the dummy variables, we included first and second order interactions of these variables with nonmajor GPA, our proxy for high school academic performance. If these variables play an insignificant role in the resulting HC index, then this would suggest that the matching shock, $\#$, actually reflects underlying HC.

Table 19 displays the point estimates of the single index when MATHPREP, REMEVER, and the associated interactions are included. The coefficients on the original variables are essentially unchanged, and the coefficients on the new variables are three or four orders of magnitude smaller and have an insignificant effect on the index realizations. Therefore, we conclude that, conditional on our standardized exam score and high school GPA proxy, no additional information is conveyed by the remedial variables, and that the analysis does not suggest the unexplained component of the HC index, $\#$, reflects underlying HC.

H.4.2. Robustness with respect to σ_h and $\sigma_{\#}$. Of course, it could be that if we had access to other information about the students' preparedness, then we might be able to reduce the variance of h , which would imply that the β estimated under our benchmark model is contaminated by

endogenous HC. This leads to our second test, which probes the robustness of our results to different assumptions about the fraction of f that is endogenous HC. In our benchmark results, we have assumed $s_h = 0$ and $s_f = s_{\#}$ —in other words, that we perfectly observe the students' HC choices. While s_f is pinned down by the distribution of $t = \tilde{s}$ required to rationalize the degree of assortativity we observe in the data, we can make different assumptions about the relative magnitudes of the value of s_h and $s_{\#}$ that generate s_f . We now demonstrate that our qualitative results (e.g., the relative importance of observable traits such as p and s versus q in the HHI equation) are robust to these alternative assumptions.

Given an estimate of the assignment mapping $P_j^r(t)$, for any pair (p, s) it is straightforward to compute $f = (P_j^r)^{-1}(p) = \tilde{s} = t = \tilde{s}$. Throughout this section we assume that h represents measurement error and is orthogonal to the matching shock $\#$, which accords with our interpretation of $\#$ as a shock that is exogenous from the student's perspective. Given this orthogonality assumption, it is then a standard signal extraction problem to apply Bayes's rule to compute the distribution of h conditional on f . Since $h \sim N(0, s_h^2)$ and $\# \sim N(0, s_{\#}^2)$, we then have that h conditional on f is normally distributed with mean $f s_h^2 / (s_h^2 + s_{\#}^2)$ and variance $s_h^2 s_{\#}^2 / (s_h^2 + s_{\#}^2)$. Of course this means that given an observation of \tilde{s} , the econometrician's beliefs about the student's underlying choice of HC $s = \tilde{s} - h$ is a non-degenerate distribution.

Given the distribution of h conditional on f , there are two primary changes that need to be made to our estimator. First, the correction applied to the B&B data to account for the selection into graduation, equation (6), now includes an integral over h , where $n_h(z|f)$ is the distribution of h conditional on f .

$$(22) \quad f_{p\tilde{s}}(p, \tilde{s}|j) = \int_{-\infty}^{\infty} \frac{f_{p\tilde{s}}(p, \tilde{s}|j, \text{grad}) G_{\tilde{s}}}{r(p, \tilde{s} + z)} n_h(z|f = P_j^{-1}(p) = \tilde{s}), \quad j \geq 2 \text{ fM}, N g,$$

Note that h does not appear in the numerator since $f_{p\tilde{s}}(p, \tilde{s}|j)$ is the joint distribution of our estimate of the single index, \tilde{s} , and college quality, p . Analogous changes to equations (8) and (9) are also required to account for the new selection correction. Note that this does affect the moment conditions described for stage I in Appendix B.1. Specifically, the definition of \bar{S}_{dl}^k must be consistent with equation (22) and the updated versions of equations (8) and (9).

We also need to alter the Mincer equation defining the stage II moment. We must replace equation (13) with

$$(23) \quad X(a) = \begin{pmatrix} 2 \\ 1 \\ 6 \\ 4 \\ \vdots \\ 1 \end{pmatrix} \begin{pmatrix} \log(p_1) \\ \log(p_2) \\ \vdots \\ \log(p_l) \end{pmatrix} \begin{pmatrix} R_{\#} \\ R_{\#} \\ \vdots \\ R_{\#} \end{pmatrix} \begin{pmatrix} \log(s_1 + z) n_h(z|f_1) \\ \log(s_2 + z) n_h(z|f_2) \\ \vdots \\ \log(s_l + z) n_h(z|f_l) \end{pmatrix} \begin{pmatrix} y(s_1, D_{M1}; a) \\ y(s_2, D_{M2}; a) \\ \vdots \\ y(s_l, D_{Ml}; a) \end{pmatrix} \begin{pmatrix} 3 \\ 3 \\ 5 \\ 5 \end{pmatrix}$$

Having made this adjustment to the definition of $X(a)$, the stage II moment condition defined by equation (15) is unchanged.

Although these modifications are simple in principal, they greatly increase the computational burden. In our benchmark model, the estimator described in Appendix B.1 can be implemented

Table 20. Point Estimates of a and $s_{\#}$ for choices of s_h^2/s_f^2

Parameter	Benchmark					
	0	0.2	0.4	0.6	0.8	1.0
s_h^2/s_f^2						
s_f , Std. Dev. of Matching Shock + Econometric Error	0.0275	0.0247	0.0239	0.0247	0.0253	0.0265
a_p , College Quality	0.1349	0.1169	0.0912	0.0565	1.27510 ⁶	1.503 10 ⁴
a_s , Human Capital	1.181 10 ⁶	1.116 10 ⁶	1.731 10 ⁶	0.467	0.2003	0.1914
a_q , Learning Cost	0.0574	0.0444	0.0676	0.0601	0.0372	0.0407
a_0 , Intercept	\$79,108	\$78,206	\$74,726	\$75,072	\$82,376	\$82,376

in two sub-stages. First, the single index and graduation probability parameters are estimated, and then $s_{\#}$ is estimated in a second sub-stage. When the estimator is modified as laid out in this Appendix, s_f (which is equal to $s_{\#}$ by definition in our benchmark model) must be jointly estimated with the other stage I parameters. The computational burden of the joint estimation and the added numerical integrals is greatly lessened by the assumption of normally distributed h and $\#$ and the resulting closed-form solution of Bayes' rule.

We can now describe the results of our analysis. We focus on two sets of parameters: the total standard deviation of the matching shock and the error in the HC index, s_f ; and the HHI parameters, a . Neither of these can be assumed to be identical to our benchmark results because the f variable influences both stages of our estimator. We vary the magnitude of s_h^2/s_f^2 , which varies the fraction of f that represents unobserved HC. Recall that this ratio is an assumption chosen to test the robustness of our model, not a product of variation in our data.

Table 20 provides the results of our analysis, and two aspects of the results warrant comment. First, the estimates of s_f , a_q and a_0 remain stable across all choices of s_h^2/s_f^2 . Although estimated from a different model than the benchmark results in the main text, s_f only ranges slightly outside of the very tight confidence bounds provided in Table 3, and a_q and a_0 remain within the 95% confidence intervals provided in Table 5. This is not surprising since the choice of s_h^2/s_f^2 does not affect the identification arguments for these aspects of the model.

Second, the magnitudes of a_p and a_s remain stable until s_h^2/s_f^2 exceeds 0.5. This is expected since the identification of a_p and a_s hinges on the existence of a matching shock that is exogenous to HC. As the exogenous component of the matching shock vanishes (i.e., as s_h^2/s_f^2 approaches 1), then identification of a_p and a_s begins to fail as a practical matter. However, the net effect on HHI of the unobservable traits of the students (q) and the net effect of the observable college and student traits (p and s) remains stable across s_h^2/s_f^2 .

H.4.3. Computational Issues, Unobserved Human Capital Model. Computational issues cause computational difficulties in the unobserved HC model. First, since the single index function only provides the observed component of HC, we must compute an integral over a Bayesian posterior that

Table 21. Estimates of the Wealth Equation with Minority HHI Penalty

Parameter	Point Est.
a_p , College Quality	0.1313
a_s , Human Capital	$2.832 \cdot 10^{-7}$
a_q , Learning Cost	0.0465
a_0 , Intercept	\$80,334

describes the distribution of the unobserved component. This occurs both in stage I (the graduation probability moment condition) and stage II (the HHI moment condition). Second, we must simultaneously estimate all of the stage I parameters, whereas in the benchmark model we can estimate the matching shock after the other stage I parameters. The nal algorithm had the following structure, where the inner loop is step (3) and the remaining steps describe the outer loop bisection routine used to pin down s_f .

- (1) Initialize upper and lower ranges of \hat{s}_f .
- (2) Set candidate \hat{s}_f equal to average of upper and lower bounds.
- (3) Estimate the remaining stage I parameters, \hat{p} , by solving the constrained optimization problem in equation (17).
- (4) Compute the model generated Kendall's t statistic given \hat{s}_f and \hat{p} .
- (5) Compute the Kendall's t statistic in the data given the estimate of θ^S .
- (6) If the model generated Kendall t statistics are larger (smaller) than the Kendall t statistics generated by the data, set the lower (upper) bound of the range of \hat{s}_f equal to the candidate value of \hat{s}_f .
- (7) Unless the upper and lower bound on \hat{s}_f are sufficiently close, return to step 2.

We used a tolerance on the upper and lower bounds of 10^{-5} for the sake of accuracy. In order to ensure that multiple restarts could detect multiple equilibria, we used a very liberal initial bound on the possible values of \hat{s}_f and allowed it to assume values between 0 and 0.5, which required the stage I estimation to be repeated 16 times for each random start.

The only remaining computational issue in stage II is to integrate over the unobserved component of HC during wage equation estimation. This also adds to computational burden, but not nearly to the same degree as in stage I.

H.5. Minority HHI Penalty. In this section we present both the estimation results and the counterfactuals after imposing an 8.3% HHI penalty for minority students as found in Fryer et al. [10] (see footnote 31). Our log utility structure means that such a penalty will primarily have the effect of shifting the levels of the utility of minority students. To the extent that this has an effect on the learning-cost estimates, it can only be through the change in the HHI equation parameters, which are close to our benchmark estimates (Table 21). Therefore, the learning costs attributed to each member of our dataset are essentially identical to those in the benchmark analysis.

Table 22. Counterfactual Minority Graduation Probability with Minority HHI Penalty

...BY LEARNING-COST QUINTILE			
Learning-Cost Type Tier	Status Quo	Proportional Quota	Color-Blind
Top Achiever Quintile	0.562	0.556	0.584
Second Achiever Quintile	0.389	0.412	0.399
Third Achiever Quintile	0.299	0.329	0.300
Fourth Achiever Quintile	0.236	0.263	0.230
Bottom Achiever Quintile	0.159	0.178	0.155
...BY COLLEGE QUALITY QUINTILE			
College Quality Tier	Status Quo	Proportional Quota	Color-Blind
Top College Quintile	0.612	0.555	0.667
Second College Quintile	0.439	0.411	0.485
Third College Quintile	0.353	0.328	0.391
Fourth College Quintile	0.270	0.265	0.305
Bottom College Quintile	0.178	0.179	0.197

Table 23. Counterfactual Minority HHI by Achievement Quintile with Minority HHI Penalty

Learning Cost Type Tier	Status Quo	Proportional Quota	Color-Blind
Top Achiever Quintile	\$72,298	\$72,861	\$71,837
Second Achiever Quintile	\$68,064	\$69,401	\$67,187
Third Achiever Quintile	\$64,365	\$66,210	\$63,075
Fourth Achiever Quintile	\$61,023	\$62,581	\$59,172
Bottom Achiever Quintile	\$56,138	\$57,067	\$54,725

The enrollment statistics are identical to the third decimal point, so we omit this table. We include the graduation probability statistics due to small differences that we observe for the proportional system. The HHI counterfactuals for the minority students differ from the baseline model primarily due to the mechanical imposition of the 8.3% penalty. These small differences suggest our model is robust to the inclusion of a uniform minority HHI penalty.

H.6. Estimation With Increased Matching Shock Variance. One of the three components of our stage I estimator's objective function is to maximize the correlation between (p, s) in our data.⁶⁵ We view the correlation-maximization exercise as an attempt to maximize the fit of our contest model to our data, which is analogous to the criteria that OLS minimize the mean squared error in order to maximize the fit of a linear model to the data. However, by trying to maximize the

⁶⁵The other two objectives are to accurately fit the graduation rates in the IPEDS data and provide a good fit between the copula parameters and the empirical CDF of the data.

Table 24. Estimates of the Wealth Equation After Increased $s_{\#}$

Parameter	Assumed Value of $s_{\#}$			
	Benchmark = $b_{\#} = 0.0275$	$b_{\#} = 4 = 0.11$	$b_{\#} = 8 = 0.22$	$b_{\#} = 16 = 0.44$
a_p , College Quality	0.1349	0.1349	0.1342	0.1341
a_s , Human Capital	$1.181 \cdot 10^{-6}$	$4.530 \cdot 10^{-6}$	$1.557 \cdot 10^{-5}$	$6.039 \cdot 10^{-5}$
a_q , Learning Cost	0.0574	0.0688	0.0921	0.1172
a_0 , Intercept	\$79,108	\$77,906	\$75,116	\$71,266
$a_q \text{ std}(\log(q))$	0.0242	0.0234	0.0237	0.0237

t of our contest model, one might worry that our estimates overstate the assortativity of the market, and that this inflated assortativity may be driving our stage II results.

In this section we artificially increase the magnitude of $s_{\#}$, the standard deviation of the matching shock, to assess what would happen to our stage II estimates if we had overstated the assortativity of the match. Recall that the assignment becomes perfectly assortative between p and s when $s_{\#} = 0$ and completely random as $s_{\#} \rightarrow \infty$. Therefore, increasing $s_{\#}$ decreases the assortativity of the college assignment. We discuss the effect of changing $s_{\#}$ explicitly in the preamble to Section 6. As $s_{\#}$ changes, the assignment mapping from NHC to college quality changes, meaning the first-order condition underlying our stage II estimator is altered. For example, if the $s_{\#}$ is so large that assignment is essentially random, then the competitive channel of incentives is effectively turned off, equilibrium HC choices will fall and our empirical model would generate much lower learning costs to explain the observed levels of HC in the data.^{66,67}

The values of the HHI parameters that result from our robustness check are reported in Table (24). The values of $s_{\#}$ studied correspond to 2, 4, 8, and 16 times the benchmark point estimate $b_{\#}$. The only point estimate that changes with $s_{\#}$ is a_q , but one must recall that the distribution of $\log(q)$ also changes as a_q changes. The values of the standard deviation of $\log(q)$ multiplied by a_q , which reflects the effect of q in the HHI equation, are reported in the final row, and these values are quite stable across choices of $s_{\#}$.

H.7. Estimating Without Low Quality Colleges. The structure underlying stage II of our estimator uses the model's first-order condition to infer the types of each of the students in our B&B sample. This exercise assumes that the students understand the mapping from their own effort into the accumulation of pre-college HC; the mapping from HC into a distribution of college assignments; and then the mapping from the quality of the college they graduate from, their own

⁶⁶This line of reasoning would seem to suggest that our control functions should also lose their identifying power, but it is important to note that while this would be true in the limit, for finite values of $s_{\#}$ we can still separately identify the influences of s and q in the HHI equation. To summarize, our identification strategy requires that $s_{\#}$ is both positive and finite.

⁶⁷Note also that under essentially random assignment, the productive channel of HC incentives is also affected. For example, academically talented students would base their investment decisions on the idea of being matched with the average quality college, rather than with the best colleges where they would otherwise expect to enroll.

Table 25. Estimates of the Wealth Equation After Dropping Worst College Seats

Parameter	Percentage of Lowest Quality College Seats Dropped					
	Benchmark	10% Dropped	20% Dropped	30% Dropped	40% Dropped	50% Dropped
a_p , College Quality	0.1349	0.1398	0.1469	0.1532	0.1397	0.1747
a_s , Human Capital	$1.181 \cdot 10^6$	$1.313 \cdot 10^6$	$3.173 \cdot 10^7$	$7.570 \cdot 10^6$	$1.563 \cdot 10^6$	$2.361 \cdot 10^6$
a_q , Learning Cost	0.0574	0.0543	0.0551	0.0668	0.0601	0.0496
a_0 , Intercept	\$79,108	\$79,591	\$79,901	\$79,363	\$79,278	\$81,375

pre-college HC, and their learning cost into post college HHI. Our estimates are less credible if some of the students either do not have rational expectations or cannot effectively optimize.

To this end, in this section we re-estimate stage II of our model after dropping some percentage of the worst performing college seats, as one might think these students/families are either less informed or less able to optimize than the students/families that attend the best schools.⁶⁸ By retaining only the students that graduate from the better colleges, we hope to make our stage II estimates more credible by deriving the estimates solely from the students that one might presume are most knowledgeable about the college admissions process and the returns to a college education.

The resulting changes to the HHI equation parameters, the only parameters estimated in stage II, are reported in Table 25. The point estimates are very stable, being all well within the 95% confidence bounds of the baseline point estimates, and the qualitative assessment, that HHI conditional on graduation is driven by p and q , remains unchanged.

H.8. Partial Adaptation to AA Rule Changes. One might wonder about the possible effects if, in the short-run, at least some of the students do not have the sophistication to immediately understand the implications of an AA rule change. In this section, we model this by considering a hypothetical scenario where only some of the students adapt their HC choice to the rule change. Students that do not adapt leave their HC unchanged from the status quo level. Those that do adapt best-respond to the resulting assignment mapping from NHC to college quality, given the actions of all competitors, including unsophisticated students whose behavior remains unchanged.

It is difficult to make general statements about the effect of partial adaptation to arbitrary changes in the AA policy due to the many moving parts of the model. To establish some intuition we consider a change to a color-blind policy, which we choose for its real-world relevance. As it happens this example is also convenient because of its straightforward connection to the numbers already reported in the last two columns of Table 6. Suppose that only 50% of students adapt to changes in the AA policy. It will be helpful to recall that AA essentially shields each demographic group from competition (either fully or in part) with the other group. Now consider the case of nonminority students, half of whom will adjust their behavior under a shift from the status

⁶⁸Due to the relatively small value of $s_{\#}$, this is almost equivalent to dropping the highest learning cost students.

quo AA to a counterfactual color-blind admissions rule. Their incentives change little since they make up 85% of the student population. Therefore, the status quo assignment mapping for nonminority students is close to the color-blind assignment mapping faced by all students because each individual nonminority student is already fully competing with most of the other students in the status quo. Since the nonminority assignment mapping changes little across these policies (see Table 6), their HC choices and outcomes do not vary much (see, for example, Table 10).

On the other hand, changing from the status quo AA to a color-blind rule has a significant effect on minority students regardless of whether they adapt or not. As per the argument above, minority students that adapt their HC choice face similar incentives to what they would face in a color-blind equilibrium, and their HC choices and outcomes thus resemble those under a color-blind equilibrium. Those minority students that don't adapt obtain outcomes similar to what they would receive in a χ -HC analysis of a color-blind admissions rule. Therefore, the average outcome for the minority students given the AA rule change lies roughly halfway between the outcome received in a full equilibrium under the color-blind rule and the outcome received under a χ -HC analysis of the color-blind rule. Finally, it is clear that the average minority outcome moves smoothly between that of the χ -HC color-blind analysis to the full equilibrium color-blind analysis as the percentage of adapting minority students moves from 0% to 100%.

H.9. Other Robustness Checks. In this section we explore various robustness checks to probe for whether our main results are stable with respect to deviations from our functional form assumptions.

H.9.1. Cost Function. We experimented with adding an additional parameter to the cost model, $c(s) = \exp(d(s - \underline{s}))$, for added exhibility. Identification would then require additional conditions, for example, an incentive compatibility assumption for marginal market participants who are indifferent to college attendance versus entering the workforce. In practice the extra parameter adds little explanatory power. The estimate of d is most influenced by the behavior of very low achieving college students where data are sparse. Elsewhere, changes in the curvature parameter d are compensated for by corresponding changes to the scale of q with little change to the economic implications of the point estimates. We also estimated the model with an alternative power law cost function of the form $C(s, q) = q(s + 1)^2$ and found that model estimates were qualitatively very similar, suggesting that our functional form assumption for costs is not unduly driving the results.

H.9.2. Non-Constant Matching Shock Variance. We also estimated a model wherein the matching shock was linear in the HC choice of the agents. This would allow, for example, the matching shock to be larger for students with lower HC choices, which might reflect higher effort costs that impede searches for potential colleges to apply to. If we include a linear trend in s , then we estimate that $s_{\#}(s) = 0.0279 - .0006s$, which is essentially the same as the result in Table 3. We also experimented with allowing a distinct matching shock parameter for minority and nonminority

students and requiring that each group's parameter best rationalize the correlation of p and s amongst students in that group. We did not find that the parameters differed significantly.

H.9.3. Alternate forms of the Income Equation We considered alternative forms for the HHI equation. For example, we explored including the estimated matching shock, $\#$, as an additional driver of HHI by estimating a HHI equation with the form $u(P, S, q) = a_0 p^{a_p} s^{a_s} q^{a_q} \#^{a_{\#}}$.⁶⁹ Our estimates imply that the marginal effect of a one standard deviation change of p and q are almost equal, the effect of s is negligible, and the effect of a one standard deviation change of $\#$ is 10% as important as that of either p or q . In short, the estimates are qualitatively the same as in our baseline model.

We estimated a model where the HHI equation is $u(p, s, q) = a_0 (p + g_p p^2)^{a_p} (s + g_s s^2)^{a_s} q^{a_q}$.⁷⁰ Our estimates with this more flexible form of the model also indicated that P and q are the primary drivers of HHI, whereas s is an insignificant driver of HHI production except for the best students placing at the best colleges. Even for these students, s is only marginally important.

Appendix I. Interpreting the utility function

The benefits that motivate agents to accumulate HC are captured by two elements of the utility function: graduation probability and future HHI conditional on graduating. Our goal in this section is to propose different interpretations of the utility function including one that admits statistical discrimination.

I.1. An Interpretation Without Statistical Discrimination. Let $f(p, s, q)$ be a primitive production function that describes how students and colleges jointly produce college-level human capital (CHC) from their “inputs,” being the college's endowment of resources p , the student's (pre-college) HC s , and her exogenous learning cost type q . One possible interpretation of $u(p, s, q)$ that does not admit statistical discrimination is that the resulting CHC is observable to firms, completely determining the marginal productivity of the graduate, and the student's HHI is determined in a competitive labor market according to her observable marginal productivity—in other words, $u(p, s, q)$ is completely determined by $f(p, s, q)$. In particular, under this first interpretation, q has no direct effect on marginal productivity after college, so firms need not infer this to determine the competitive wage of the worker.⁷¹

I.2. An Interpretation With Statistical Discrimination. We now provide a second interpretation that admits a role for signaling (and therefore statistical discrimination on the basis of observed race) between college graduates and employers. Under this interpretation, $u(p, s, q)$ is a value

⁶⁹The effect of $\#$ is identified off of the functional form assumption, so we do not mean to place great weight on this analysis.

⁷⁰Note that the form of the nonlinear control function $y(s, D_M; a, g_p, g_s)$, which stands in for q , is already included by the additional terms in p and s . Since the form of y is determined endogenously within the model, we do not add a quadratic term in q .

⁷¹Alternatively, if part of the value of attending a high quality college is the existence of a referral network allowing access to high HHI jobs, then this could be formalized as p entering $u(p, s, q)$ separately from, or in addition to, the CHC production function $f(p, s, q)$.

function for some interaction between college students, colleges, and firms. Suppose that the marginal productivity of a college graduate is a function of the CHC accrued during college (i.e., $f(p, s, q)$) and her learning cost q , which could (for example) influence the student's ability to acquire job- and firm-specific human capital later in life. We assume the student's underlying type q is not directly observable to the firms, but the firms can observe the student's CHC and make inferences about q given the information at hand.

Statistical discrimination can occur only if demography carries information about q even in the presence of the other data observable to the firm.⁷² For example, events that could introduce noise into CHC during a student's 4-year college career may include both negative life shocks (e.g., sickness or turbulence in a romantic relationship) and positive life shocks (e.g., finding an excellent personal mentor) that affect her learning process. Moreover, these shocks could imply either that the stock of CHC accrued during college is not fully under the student's control (e.g., sickness interferes with mid-semester attendance) or that the firm's observations of the underlying CHC are imperfect (e.g., after performing well, a family emergency prevents the student from adequately concentrating during a final exam). In either case, the student's demography can carry useful information about q even when p and the (potentially stochastic) choice of CHC and/or (potentially imperfect) observation of CHC are known.⁷³

I.3. Structural Primitives vs Reduced-Form Utility. If $u(p, s, q)$ is a reduced form for a post-enrollment game, then one must take care to consider the interpretation of it if the AA system changes. If $f(p, s, q)$ is a structural primitive, then $u(p, s, q)$ is by definition stable to the extent that it is driven by observable CHC. However, $u(p, s, q)$ could change if the firm's beliefs about the student's type change significantly. Consider a particular minority student. If the AA system changes, then the college quality assigned to that student changes significantly. If firms treated school quality as a sufficient statistic of q (based on the average learning costs of all the students enrolled at the school), then the firm's beliefs about this student would change significantly since the quality of the school she enrolls in (and the learning-costs of her nonminority peers) change significantly.⁷⁴ However, if firms also condition on a student's race, the equilibrium beliefs of the firms need not change significantly because the firms' inferences take into account the disproportionate effect that changes to the AA system have on the college at which minorities with each realization of q enroll. Under our interpretation, this would mean that $u(p, s, q)$ would continue to be approximately correct under the equilibrium in the new AA system.

⁷²Recall that type distributions may differ by demographic status so that q , which derives from both innate and early-life environmental factors, is correlated with race.

⁷³If the students play a signaling game to convey information about their learning cost through the CHC observation of the firm, then the firm's equilibrium beliefs about their learning costs are non-degenerate and the firm does not perfectly know the student's realization of q . Thus, firms might have an incentive to use observable race as an additional control for predicting q .

⁷⁴Since all minority students are affected simultaneously, the learning-costs of her minority peers would change much less as the AA system changes.

