LAW SCHOOL ADMISSIONS UNDER THE UC AFFIRMATIVE ACTION BAN

Danny Yagan*
University of California, Berkeley

December 2012

ABSTRACT

The consequences of banning affirmative action depend on schools' ability and willingness to avoid it. This paper uses a seventeen-year sample of law school applications to measure how completely UC law schools avoided the 1996 UC affirmative action ban. The pre-ban black admission rate was 61%. Controlling for selective attrition from applicant pools, I find that the ban reduced the black admission rate to 31%—half of the pre-ban rate but still four times higher than the 8% rate that would prevail under observed white admission standards. Observed black admission advantages at intermediate credential levels were as large as 99 percentage points before the ban and 63 percentage points after the ban. The results have implications for modeling affirmative action bans, sustaining racial diversity under a ban, affirmative action constitutionality, the effectiveness of mandating nondiscrimination, and identifying discrimination in cross-sectional data. *JEL* Codes: I23, J7, K42.

*Email: yagan@econ.berkeley.edu. Mail: 530 Evans Hall #3880; Berkeley, CA. Phone: (510) 642-0822. Fax: (510) 642-6615. I thank Joseph Altonji, Joshua Angrist, Elizabeth Bartholet, David Card, Raj Chetty, David Cutler, John N. Friedman, Roland Fryer, Claudia Goldin, Joshua Gottlieb, Nathaniel Hilger, Caroline Hoxby, Lisa Kahn, Louis Kaplow, Lawrence Katz, Ilyana Kuziemko, Jessica Laird, N. Gregory Mankiw, Justin McCrary, Sendhil Mullainathan, Jesse Rothstein, Emmanuel Saez, Andrei Shleifer, Joseph Singer, Matthew Weinzierl, and Martin West for their comments. Sarah Abraham, Michel Kim, Amol Pai, and Michael Stepner provided excellent research assistance.

I Introduction

Selective U.S. universities practice affirmative action—using race in favor of black applicants and other underrepresented minorities—in admission decisions. Seven states have banned the practice at their public universities, and the U.S. Supreme Court may extend these bans nationwide when it decides *Fisher v. University of Texas* this year. Key to economic, policy, and legal debates is the ability and willingness of schools to blunt a ban's impact on black admission rates by increasing admission weight on low family income and other characteristics that correlate with black status.

This paper analyzes admission outcomes before and after Proposition 209, the ballot initiative that banned affirmative action at the University of California beginning in 1996. I focus on law school admissions at Berkeley and UCLA; in this paper, "Berkeley" and "UCLA" always refer to these campuses' law schools. Law school admissions hold center stage in affirmative action debates: the most recent Supreme Court decision concerned a rejected white law school applicant suing for damages, and recent empirical work on the effect of banning affirmative action on student outcomes has utilized cross-sectional data on law school admissions and bar passage rates (Sander 2004; Rothstein and Yoon 2008).

Berkeley and UCLA had stated policies of affirmative action before the ban, denied using race after the ban, and remained opposed to the ban throughout. Empirically, Berkeley and UCLA's unconditional black-white admission rate differences remained approximately unchanged—consistent with near-complete avoidance, or alternatively with the complete elimination of black admission advantages combined with substantial selective attrition from the applicant pool by relatively weak black applicants.

This paper controls for selective attrition using new application-level data on all 25,499 law school applications submitted to law schools nationwide between 1990 and 2006 by 5,353 undergraduates from an elite college. All results are local to this dataset, which I call the Elite Applications to Law School ("EALS"). Cross-sectional admission differences in the EALS are similar to those documented in the universe of law school applications to elite schools and in other university contexts (Kane 1998; Bowen and Bok 2000; Espenshade, Chung, and Walling 2004; Rothstein and Yoon 2008).

As a baseline, the pre-ban black admission rate at Berkeley and UCLA (averaged across the two schools) was 61% in the EALS. Had all pre-ban applicants been subject to observed white admission standards, I estimate that the black admission rate would have been 8%. EALS covariates are sufficiently rich to identify not only this 53-percentage-point overall admission rate disparity but also the particular margins where race was decisive: at certain intermediate credential levels where applicants were neither all accepted nor all rejected ("the accept/reject margin"), pre-ban black applicants to Berkeley and UCLA were 99 percentage points more likely to be accepted than whites.

I find that when controlling for academic credentials and nationwide trends, the UC affirmative action ban permanently reduced black admission rates by 30 percentage points. Thus the ban caused a large reduction in black admission advantages over similarly credentialed whites, unless the black-white difference in non-academic admission determinants (e.g. recommendation letter strength) changed substantially after the ban. Controlling additionally for a measure of unobserved applicant strength—inferred from admission decisions at non-UC schools (similar to Dale and Krueger 2002) and empirically powerful in the EALS—leaves the main estimates nearly unchanged, as does controlling for the UC-specific admissions factor of California residency. I conclude that the ban caused a large reduction in black admission advantages.

This large reduction occurred in spite of Berkeley and UCLA sustaining large observed black admission advantages under the ban. In particular, at the accept/reject margin, I estimate that black applicants to Berkeley and UCLA were 63 percentage points more likely to gain admission than whites after the ban, with 95% confidence lower bounds of approximately 35 percentage points and upper bounds above 75 percentage points. These post-ban disparities are net of Dale-Krueger measures of unobserved applicant strength, so they reflect unique weight given to black-correlates (e.g. low family income) at UC schools and any continued use of race. These post-ban disparities were relatively systematic: admission is a binary outcome, so a 63-percentage-point admission rate difference implies that the *fraction* of blacks strong enough to be admitted was 63 percentage points larger than the corresponding *fraction* of whites.

Auxiliary data provide a sense of the information that UC schools could have weighted uniquely in order to generate relatively systematic post-ban advantages. In an auxiliary dataset of which the EALS is a part, the maximum black-white admission rate difference that can be generated with family income is 37 percentage points—much, but not all, of the way to 63 percentage points. Berkeley and UCLA responded to the ban by soliciting new legal black-correlates such as essays on applicants' anticipated contributions to cohort diversity. The schools also had access to signals of black status that would likely have been illegal to use: in data collected from the elite college's yearbooks, I find that 84% of black students and 0% of white students listed participation in a black-focused extracurricular group. The main empirical estimates show that the net effect of Berkeley and UCLA weighting any or all of these black-correlates was to generate 63-percentage-point admission rate advantages for marginal black candidates and to allow overall black admission rates to fall by half.

The results have five implications. At the highest level, this paper contributes to a large literature on whether nondiscrimination mandates in fact constrain the behavior of agents who wish to use race in regulated decisions (Becker 1968; Arrow 1998; Heckman 1998). Freeman (1973), Heckman and Payner (1989), and Chay (1998) find that the 1960s Civil Rights Acts and related policies improved black labor market outcomes, though possibly by giving employers an excuse to break social codes against hiring black workers rather than by constraining employers who wanted to discriminate (Heckman and Payner; Heckman and Verkerke 1990). More recent acts have not been found to have improved the outcomes of protected groups (Acemoglu and Angrist 2001; Oyer and Schaefer 2002), perhaps because cross-group differences in labor market outcomes are due to skills differences and not discrimination (Neal and Johnson 1996; Fryer 2011). This paper shows that in university admissions—a prominent part of the modern economy in which selectors use race—California's nondiscrimination mandate was effective at constraining selectors' decisions.

Yet the constraint far from eliminated racial disparities in admission outcomes, with important implications for simulating the consequences of affirmative action bans. Recent quantitative predictions of the effects of eliminating affirmative action have considered the hypothetical *elimination* of observed racial disparities in admission outcomes (Arcidiacono 2005; Krueger, Rothstein, and Turner 2006; Rothstein and Yoon 2008). In the EALS and holding applicant characteristics fixed, eliminating observed racial disparities in UC admissions would have reduced the black admission rate from 61% to 8%. I find that the affirmative action ban reduced it only to 31%. Hence, assuming black-white parity in admission outcomes can substantially overstate the effects of banning affirmative action, perhaps the primary policy-relevant counterfactual. This finding is consistent with earlier theoretical work that emphasized avoidance possibilities (Chan and Eyster 2003; Fryer, Loury, and Yuret 2007; Epple, Romano, and Sieg 2008).

The decline in black admission prospects was nevertheless substantial, and a specific policy implication is that maintaining high levels of racial diversity under a ban may require forcing schools to use black-correlates more aggressively than the schools themselves desire. Post-ban UC schools generated 63-percentage-point observed black admission advantages at the accept/reject margin by weighting some unknown index of characteristics that correlated with black status (e.g. low family income and diversity essays). If they had placed arbitrarily high weight on that index, they would have admitted black applicants at higher rates, potentially even restoring overall pre-ban black admission rates.[1] UC schools did not do so—a rational choice if they valued non-racial strength (e.g. high GPAs and

[1] As a group, pre-ban blacks were admitted at a 53-percentage-point-higher rate than whites with similar credential levels. If post-ban schools had sufficient information to generate 63 percentage point differences at all credential levels (and not just at the accept/reject margin), then they of course could have restored a 53-percentage-point overall admission rate difference.

test scores) sufficiently highly or if higher black admission rates would have attracted litigation for violating the ban. Perhaps anticipating such voluntarily incomplete avoidance, the State of Texas followed its court-ordered affirmative action ban with a "Top 10%" law that required each University of Texas campus to admit the majority of their undergraduates based on a single race-blind criterion (high-school class rank) that very disproportionately benefitted minorities.

The UC's substantially incomplete avoidance can also be seen as bearing on the constitutionality of affirmative action, which currently requires there to be no "workable race-neutral alternatives that will achieve the diversity the university seeks" (*Grutter v. Bollinger* 2003).[2] Arithmetically based on public aggregates, the ban reduced Berkeley and UCLA black enrollment by reducing black application rates.[3] However, under a nationwide ban in which blacks have fewer alternative schools to apply to, the effect on applicant pools could be much smaller (Arcidiacono 2005; Epple, Romano, and Sieg 2008; Hinrichs 2012). This paper's findings suggest that even when holding the applicant pool constant, UC schools would have suffered a large decline in black enrollment. Hence if one were to use universities' revealed preferences to determine "workability", the above constitutional requirement could be judged to have been met in the EALS.

Finally, this paper demonstrates that when the outcome is binary, real-world cross-sectional data can be rich enough to identify discrimination. I estimate that pre-ban black applicants were 99 percentage points more likely to gain admission at certain credential levels than whites. In order for this difference to have been caused without the use of race, the share of blacks that were strong enough on unobserved factors to be admitted would have to have been nearly 100%, and the corresponding share of whites would have to have been nearly 0%. It is unlikely that *any* combination of characteristics other than race could have been used to differentiate between races so well. This highlights a statistical moment that may be useful in discrimination litigation—where contested outcomes are frequently binary (e.g. discrimination in loan approval), plaintiffs can obtain very rich individual-level data (such as mortgage application data that are required to be compiled under the Home Mortgage Disclosure Act), and courts typically demand more than statistically significant overall differences in order to infer discrimination (Selmi 2000). It also constitutes an alternative to recent experimental and quasi-experimental strategies for identifying discrimination and other causal impacts of race and gender (Goldin and Rouse 2000; Bertrand and Mullainathan 2004; List 2004; Price and Wolfers 2011; Anwar, Bayer, and Hjalmarsson 2012).

[2] The other two requirements are a "compelling government interest" in diverse campuses (currently "the educational benefits that flow from a diverse student body") and an eventual sunset.

[3] The black share of enrollment equals the product of black share of applicants, the unconditional (not holding applicant characteristics fixed) admission rate, and the yield rate. The latter two terms remained largely unchanged.

The remainder of the paper is organized as follows. Section II introduces the UC affirmative action ban and the EALS dataset. Section III estimates the effect of the ban on black admission rates. Section IV documents racial disparities in post-ban admissions. Section V discusses mechanisms. Section VI concludes.

II Institutional Background and Data

II.A The UC Affirmative Action Ban

Affirmative action—awarding admission preference to underrepresented minorities on the basis of race—is currently legal at all private U.S. universities and is legal at all public ones except those in the seven states that have banned it.[4] On November 5, 1996, California voters approved Proposition 209 which amended the state constitution to read: "The state shall not discriminate against, or grant preferential treatment to, any individual or group on the basis of race, sex, color, ethnicity, or national origin in the operation of public employment, public education, or public contracting." In particular, no University of California applicant is to be preferred to another on the basis of race. The ban went into effect immediately at UC law schools: application forms were amended to declare that race was not a criterion for admission, and the page requesting applicant race was diverted to a UC statistical department and was not reported to admission offices.[5] Institutional features like the number of first-year enrollees (275 at Berkeley Law, 310 at UCLA Law) remained nearly unchanged; admitted cohorts are approximately three times as large as enrollment cohorts.[6] Judging from public aggregates, the ban reduced black enrollment mainly by reducing the black share of applicants, rather than by reducing black admission or yield rates; *a priori*, this reduction in black applicants may or may not have been driven by a reduction in black admission advantages (see Section III.A below).

Law school admission decisions are made by a small group of selectors applying subjective criteria with little transparency. The UC publishes annually the total number of applicants by race and the total number of admitted students by race but not academic characteristics by race (see Figure II below). Soon after the ban, the State of California guaranteed that high school seniors graduating in the top 4% of their high schools would gain admission to at least one UC campus but not necessarily the one of their choice. With eight UC campuses, this had little binding effect on undergraduate admissions at the elite campuses of Berkeley and UCLA and likely had little effect on Berkeley and

[4] Other than California, the states currently under affirmative action bans—all implemented after California's—are Arizona, Florida, Michigan, Nebraska, New Hampshire, and Washington. Georgia and Texas had temporary bans.

[5] The ban went into effect one year later at UC undergraduate campuses.

[6] Similar to other top law schools over the 2000s recession, Berkeley's admitted cohort size fell 10% in 2001 as its yield rose. UCLA's admitted cohort size declined temporarily in 1998 after a temporary spike in yield. The EALS does not offer sufficient statistical power to analyze yield.

UCLA law school admissions. Thus the main policy potentially affecting Berkeley and UCLA law school admission offices was the affirmative action ban itself, rather than any additional regulations.

This contrasts, for example, with the undergraduate admission office at the University of Texas at Austin, the state system's flagship campus that was temporarily bound by an affirmative action ban and analyzed by Long and Tienda (2008). Austin was forced by a state law to automatically admit most of its students based on a single race-blind criterion that very disproportionately aided minority applicants at the substantial expense of the school's preferred measures of academic qualifications (University of Texas 2007).[7] Austin admissions were also conducted on such a large scale (an order of magnitude larger than Berkeley and UCLA) that the remaining non-automatic admissions were themselves determined by an academic strength index and a subjective index assigned to each application by one of forty-four application readers. The University of Texas published audits of a subset of these application readers, testing for scoring consistency according to the official race-blind rubric (University of Texas 2005). These additional constraints on admission office discretion may explain why Long and Tienda find that minority applicants under Texas's ban did not enjoy cross-sectional admission advantages in non-automatic admissions.

II.B Data Source, Variables, and Sample Restrictions

This paper's primary dataset—which I call the Elite Applications to Law School (EALS)—contains confidential individual-level data on 67% of an elite college's seniors and graduates who applied to law schools nationwide between the fall of 1990 and the fall of 2006. Applications to almost every U.S. law school are submitted through the Law School Admissions Council, which records application information and the admission decision for every application filed. Academic credentials are verified through third-party reports, and race is reported by applicants where dishonest answers are grounds for revocation of an admission offer, expulsion from law school, or disbarment.[8] Applicants choose whether to release their data to their colleges' administrators, and I obtained and digitized seventeen years of a single college's data. The college is elite, is not on the west coast, and has never been subject to an affirmative action ban. Subsection C investigates possible selection over time into the EALS, and Section III addresses selection over time into the Berkeley and UCLA applicant pools.

[7] After Texas's temporary ban, the legislature required *every* University of Texas campus to admit every undergraduate applicant who graduated in roughly the top 10% of his or her Texas high school class. As of 2007, over two-thirds of Texas residents admitted to the flagship undergraduate campus at Austin were admitted via this automatic guarantee. Austin admitted 12,000 undergraduates per year.

[8] To the extent that any applicants misreported their race, the EALS race variable nevertheless represents the race that was reported to schools on application forms. LSAC's Credential Assembly Service details and its Misconduct and Irregularities policy can be found at http://www.lsac.org/jd/apply/cas.asp and http://www.lsac.org/jd/apply/misconduct-and-irregularities.asp, respectively.

The EALS contains six variables for each application: applicant race, LSAT test score (integers between 120 and 180), undergraduate grade point average (GPA) to two decimal places on a 4.00 scale, application year, law school submitted to, and admission decision.[9] I standardize LSAT and GPA to each have mean zero and standard deviation one across applicants. Motivated semi-parametrically in Subsection D and used in figures, I summarize applicants' LSAT and GPA scores with a scalar measure I call "academic strength" equal to the standardized sum of standardized LSAT and standardized GPA, similar to the rescaling that Kling, Liebman, and Katz (2007) employ in a different context. Application years 1990-1991 through 2001-2002 as well as 2005-2006 also contain applicant state of permanent residence; for these years, I digitized a California resident indicator for Berkeley and UCLA applications only.[10]

The raw data contain 38,200 applications of 6,072 applicants to 187 law schools. I restrict the analysis sample to the 94.3% of applicants listed as white, Asian, black, or Hispanic and the 78.9% of applications submitted to UC Berkeley, UCLA, or one of the fifteen most-applied-to schools that were never subject to an affirmative action ban. These fifteen schools correspond closely to the top-ranked law schools according to *U.S. News and World Report*, so I refer to them only somewhat imprecisely as the "top fifteen non-UC law schools."[11] The 170 other schools received relatively few applications in the EALS and are poor control schools for Berkeley and UCLA because these 170 other schools are less selective. The final seventeen-school EALS sample comprises 25,499 applications submitted by 5,353 applicants. See Online Appendix A for additional data-coding details.

II.C Summary Statistics

Table I lists summary statistics. The EALS sample is 61% white, 10% black, 19% Asian, and 10% Hispanic. Black applicants on average possess LSAT scores and GPA's 1.1 and 1.0 standard deviations lower, respectively, than white applicants. Online Appendix Figures Ia-c use non-parametric densities

[9] The ability of GPA to predict EALS admissions decisions (documented below) may derive in part from the fact that all EALS applicants attended the same undergraduate institution, rather than different institutions with potentially very different grading critera as is typically the case in undergraduate admissions data. Arcidiacono, Aucejo, Coate, and Hotz (2011) study the effect the of the ban on black graduation rates using confidential undergraduate data obtained directly from the UC through a California Public Records Act request. These other data would have limitations for a study of admissions: they are aggregated into three-year intervals, do not distinguish blacks from other non-Asian minorities, bin test score and GPA variables rather than reporting the precise numbers available to the admissions office, and do not contain identifiers for previous educational institution.

[10] The raw data comprise approximately one thousand pages of print-outs from a spreadsheet program. State of permanent residence was not printed in some years, apparently related to spreadsheet format changes. The spreadsheet format facilitated digitization using scanners and optical character recognition software. The software erroneously and inconsistently read two-letter state abbreviations as various two-letter words, so I manually coded a California resident indicator for Berkeley and UCLA applications only. This California resident indicator is significant with $p < 0.001$ in probit regressions of admission on the California resident indicator, LSAT, GPA, race indicators, and school-year fixed effects among Berkeley and UCLA applications.

[11] Deviations from *U.S. News* rankings are usually explained by a lower-ranked school being located in a large city.

of these academic characteristics to illustrate the first order stochastic dominance of the black and Hispanic distributions by the white and Asian distributions. Online Appendix Figure Id plots means of a summary measure of LSAT and GPA over time by race among EALS applicants; post-ban and pre-ban means are very similar within races, suggesting little differential selection over time into the EALS. Section III addresses differential selection over time into the Berkeley and UCLA applicant pools.

Berkeley received applications from 28% of all applicants (1,594, making it the seventh-most-applied-to school in this sample) and UCLA received applications from 14% of all applicants (777, the thirteenth most in this sample); see Online Appendix Table I for additional comparisons. These schools received relatively few applications from black students—60 before the ban and 67 after the ban at Berkeley, and 31 before the ban and 27 after the ban at UCLA—but as I demonstrate next, the EALS provides sufficient statistical power because within-race admission decisions are largely determined by academic credentials.

II.D Race and Admission in the Pre-Ban Cross Section

Figure Ia displays the semi-parametric relationship between LSAT, GPA, and admission within race-school-years in the EALS, using a 5% random sample of all 23,128 applications submitted to non-UC schools (Online Appendix Figure II displays the 100% sample, intelligible only in color). Each application's admission decision is plotted in (GPA, LSAT) space, where each application's LSAT score has been adjusted by the estimated race-school-year fixed effect in order to visually account for selectivity differences across races, schools, and years. Specifically I fit a probit regression of admission on standardized LSAT (mean zero and standard deviation one), standardized GPA, and school-year-race fixed effects; add each application's estimated school-year-race effect to its LSAT value; and plot individual application decisions in GPA vs. adjusted LSAT space. Applications above and to the right of the best-fit admission threshold line have high enough LSAT and GPA scores to have a predicted admission probability of more than 50%, while those below and to the left do not.[12]

The best-fit line correctly predicts 85.4% of all admission decisions, and incorrect predictions are concentrated near the line. The ratio of the coefficients on LSAT and GPA in the underlying probit is 0.95, indicating that a one standard deviation higher LSAT is about as valuable in the admissions cross section as a one standard deviation higher GPA. When useful for subsequent illustrations, I therefore summarize an applicant's "academic strength" as the standardized (mean zero, standard deviation one) unweighted sum of standardized LSAT and standardized GPA. Figure Ib shows that

[12] The probit model is $\Pr(ADMITTED_{istr}) = \Phi(\beta_1 LSAT_i + \beta_2 GPA_i + \gamma_{str})$ where γ_{str} denotes the fixed effects. "Adjusted LSAT" equals $LSAT_i + \hat{\gamma}_{str}/\hat{\beta}_1$. This slope of the best-fit admission threshold line is 0.95, equal to $-\hat{\beta}_1/\hat{\beta}_2$.

the semi-parametric relationship between academic strength and admission within race-school-years is well-approximated by a univariate probit regression of admission on academic strength alone. I refer to such a curve relating admission to academic strength as an "admission rule" in academic strength.

Figure Ic plots fitted admission rules for blacks and whites in pre-ban Berkeley and UCLA admissions.[13] For ease of comparison, each school's fitted rules have been shifted horizontally by an additive constant so that the admission probability for whites equals 0.5 at academic strength 0 to facilitate comparison. The graph shows that there are levels of academic strength at each school where blacks were nearly assured admission and whites were nearly assured rejection. Berkeley's black and white admission rules are separated by 1.90 standard deviations of academic strength, implying black status is observed to be worth more than the difference between an A- GPA and a B- GPA for a given LSAT in the pre-ban cross section.[14] At UCLA, the difference is 1.39 standard deviations. Had pre-ban black applicants to each school been subjected to the observed pre-ban white admission standards, Berkeley's black admission rate is predicted to have been 6% rather than the actual 57%, and UCLA's to have been 10% rather than 65%. I formally document these differences in Section IV. These black-white differences in the EALS are similar in magnitude to those found in the universe of law school applicants to elite schools like Berkeley and UCLA (Rothstein and Yoon 2008).[15]

Online Appendix Figure III shows that black and white admission rules have similar steepness when admission is allowed to respond differently to academic strength for each race in the underlying regressions. This appendix figure also illustrates pre-ban admission rules for Hispanics and Asians, as well as non-parametric densities of applicant academic strength by race at Berkeley, UCLA, and the average non-UC school. Roughly speaking, Hispanic applicants enjoyed smaller cross-sectional admission advantages than blacks. For simplicity and statistical power, this paper focuses on black admission outcomes. In unreported results, the effects of the ban on Hispanic admission decisions are similar to those reported for blacks in Section III (they are large and statistically significant), while the analysis of cross-sectional admission advantages conducted in Section IV produces relatively uninformative confidence intervals when done for Hispanics.[16]

[13] For each school I estimate the probit model $\Pr(ADMITTED_{it}) = \Phi(\beta_1 ACADEMICSTRENGTH_i + \beta_2 BLACK_i + \gamma_t)$ using pre-ban black and white applications, where $BLACK_i$ is a black indicator and γ_t denotes year fixed effects.

[14] That is, $\hat{\beta}_2/\hat{\beta}_1 = 1.90$ in the underlying Berkeley regression.

[15] Rothstein and Yoon report that in population data from the 1990-1991 admissions cycle, black enrollment at elite law schools would have been 90% lower under white admission standards.

[16] Visually apparent in the relatively flat Hispanic admission rules shown in Appendix III, Hispanic admissions decisions in the pre-ban baseline, especially at UCLA, were noisier than black admissions decisions. Hispanic outcomes may be subject to confounds that black outcomes are not, for example due to California's large and growing Hispanic population over this time period; the California-residency indicator is unavailable for approximately half of the post-ban era.

III Effect of the Ban on Black Admission Rates

This paper's core empirical analysis proceeds in two steps. First, this section investigates whether the ban *changed* black-white admission rate differences, using difference-in-differences regressions to estimate the effect of the ban on black admission rates controlling for post-minus-pre changes in white admission rates and for national trends and holding applicant characteristics constant. Second, Section IV investigates whether the ban *eliminated* black-white admission rate differences at UC schools, using cross-sectional regressions on post-ban admissions only. These are complementary analyses but use distinct specifications, so I present their results sequentially.

III.A Ambiguous Indications from Unconditional Aggregates

Figure II uses public aggregates on the universe of law school applicants (not just the EALS) to plot the time series of overall admission rates by race as well as the black share of the applicant pool at Berkeley and UCLA. After a temporary large decline, the black-white admission rate gap narrowed at each school and at Berkeley even exceeded pre-ban levels.

This may appear to imply that admission offices fully avoided or evaded the ban in the long run, but the permanent 45% decline in the black share of the applicant pool at each school leaves the door open to potentially strong compositional effects. For example if the ban dissuaded applications from black students of all credential levels equally—perhaps because post-ban black students perceived Berkeley to be a less black-friendly place to study—it may have failed to affect black admission advantages. But if the ban dissuaded applications from weakly-credentialed black students because they suddenly expected rejection as suggested by Card and Krueger (2005) in the Texas context, it may have completely eliminated black admission advantages. Indeed I find some evidence of the latter selection pattern in the EALS: at marginal significance levels, less-academically-credentialed black applicants in my sample were substantially less likely to apply after the ban.[17] I use rich applicant characteristics to control for such compositional changes in the EALS.

III.B Visual Evidence

Figure IIIa displays the time series of black and white admission rates at Berkeley, UCLA, and non-UC schools, where applicant characteristics have been held constant at pre-ban levels using simple semi-parametric reweighting as in DiNardo, Fortin, and Lemieux (1996). In particular, to construct the

[17] EALS covariates explain admissions decisions very well but explain *application* decisions with r^2-statistics an order of magnitude smaller, apparently because applicants choose among similarly-ranked schools based on geographical preferences and other factors omitted from the EALS. I consequently have little power to analyze the effect of the ban on application decisions, such as whether the ban reduced application rates among high-credentialed black students—the focus of Card and Krueger (2005).

time series of black admission rates at Berkeley, I first compute terciles of academic strength (the scalar summary measure of LSAT and GPA defined and motivated in Section II.D) among pre-ban Black applications to Berkeley.[18] Then for each time period shown in the figure, I weight black applications to Berkeley so that each pre-ban-defined tercile receives equal weight when computing the displayed admission rate. I repeat this process for whites at Berkeley and for whites and blacks separately at UCLA and at each non-UC school, averaging across non-UC schools to construct the non-UC series. This semi-parametric reweighting is data-demanding, so I group the data into two pre-ban time periods (1990-1992 and 1993-1995) and two post-ban time periods (1996-2000 and 2001-2006).

The figure shows that at non-UC schools, there was little change over time in the difference between black and white admission rates. At Berkeley the black admission rate rose between 1990-1992 and 1993-1995 about as much as the white admission rate did. Between 1993-1995 and 1996-2000, the black admission rate fell from 64% to 33% and did not subsequently recover relative to the white admission rate. Figure IIIb shows a similar decline at UCLA.[19]

One can use such reweighted admission rates to compute a difference-in-differences-in-differences (DDD) estimate of the effect of the ban on the black admission rate at each UC school that controls for changes in academic strength—equal to the change in black admission rates at the UC school, net of the change in white admission rates at the UC school and changes in the black-white admission rate difference at non-UC schools.[20] Pooling pre-ban years and post-ban years, the DDD estimate of the effect of the affirmative action ban on Berkeley's black admission rate is -30 percentage points, relative to the actual pre-ban black admission rate of 57%. For UCLA, the estimate is -41 percentage points, relative to the actual pre-ban black admission rate of 65%. See Online Appendix Table II for the arithmetic. These declines were much larger than those observed at any non-UC school, so the empirical p value of each of these declines across non-UC schools is 0.

III.C Regression Estimates

Table II reports regression estimates of the effect of the ban on black admission outcomes at each UC school, computed by fitting probit and OLS models based on the following difference-in-differences (DD) specification:

$$\text{(1)} \qquad \Pr(ADMITTED_{it}) = \Phi(\mathbf{X}_i\boldsymbol{\alpha} + \beta_1 BLACK_i + \beta_2 BLACK_i \times POST_t + \gamma_t)$$

[18] Quartiles yield similar results; I use terciles because some bin counts are small.

[19] See Online Appendix Figure IV and its notes for the non-reweighted time series and a comparison to the nationwide data presented in Figure II. The non-reweighted change is negative and uneven due to differences in black academic strength over time; the reweighting in Figure III adjusts for these differences.

[20] Note that this statistic does not account for the fact that a decline in the black admission rate opens up space for some more applicants of all races; I account for this in the analogous parametric estimates presented below in Table 2.

using black and white applications to either UC school, where $ADMITTED_{it}$ is an indicator for whether applicant i's application in year t earned an admission offer; $BLACK_i$ is an indicator for applicant race; $POST_t$ is an indicator for the application being submitted after the ban; $\mathbf{X}_i$ is a vector containing LSAT score, GPA, and other covariates depending on the specification; and γ_t is a vector of year fixed effects. Online Appendix Tables III and IV replicate Table II using alternative specifications that include all races and control for more interactions. When producing DDD estimates that account for national trends, I include all black and white applications to the top-fifteen non-UC schools and interact the second and third terms with an indicator for the application being submitted to a non-UC school.[21] Standard errors are clustered at the applicant level.

Columns (3)-(4) display the basic probit results. Column (3) of panel A shows that when confining attention to applications to Berkeley, the ban is estimated to have caused a 40 percentage point reduction in the probability of admission, averaged over the characteristics of pre-ban black applicants and relative to the actual pre-ban black admission rate of 57%. Controlling for trends at non-UC schools, Column (4) displays a DDD estimate of -36 percentage points. Panel B reports a DDD estimate for UCLA of -33 percentage points, relative to the actual pre-ban black admission rate of 65%. These effects are statistically significant with t statistics between 3 and 7.

Holding all else equal, a decline in black admission rates relative to whites opens up space in the admitted cohort for both black and white applicants, suggesting that these estimates may somewhat overstate the effect of the ban on black admission rates. I therefore compute an adjusted estimate of the effect of the ban on the black admission rate at each UC school by using the UC-specific coefficients of each regression to compute a probit latent variable value for each black and white pre-ban applicant according to post-ban criteria, and then adding a constant to every applicant's value until the mean predicted admission probability across applicants equals the actual admission rate observed among these applicants.[22] These estimates are reported in the bottom row of each panel of Table II. The resulting estimates are only 3 to 5 percentage points lower than the DDD estimates reported above.

The identifying assumption of the DDD regression in column (4) is that any post-minus-pre changes in the unobserved strength of black applicants relative to white applicants was not local to applicants to UC schools. Applicants choose which schools to apply to, so one may be concerned that the

[21] The DDD specification is $\Pr(ADMITTED_{ist}) = \Phi(\mathbf{X}_i\boldsymbol{\alpha} + \beta_1 BLACK_i + \beta_2 BLACK_i \times POST_t + \beta_3 BLACK_i \times UC_s + \beta_4 BLACK_i \times POST_t \times UC_s + \gamma_{st})$, where UC_s is an indicator for whether the application was submitted to the UC school being analyzed and γ_{st} is a vector of school-year fixed effects. I weight applications so that each school carries equal weight in each time period (pre-ban and post-ban).

[22] Adding a constant is a way to vary selectivity uniformly across applicants. I obtain similar results using an alternative method: using the UC-specific coefficients to rank pre-ban applicants and then "admitting" the N highest-ranked applicants, where N equals the total number of black and white pre-ban EALS applicants that the UC school admitted.

ban induced differential selection across races into UC applicant pools such that post-ban blacks were relatively much weaker on unobserved admission determinants like recommendation letter strength. I address this first by augmenting Equation (1) with an additional "inferred strength" control, which is based on independent admission decisions akin to Dale and Krueger (2002). Preferences of admission offices are highly correlated across law schools; Figure I and Online Appendix Figure III showed this to be the case for directly observed applicant characteristics (LSAT, GPA, and race).[23] All top law schools solicit and are believed to value applicant characteristics like recommendation letters, leadership experience, and a background of no criminal behavior or academic dishonesty. I proxy for such commonly-valued unobserved admission determinants using the intuition that if an applicant predicted to be rejected based on LSAT, GPA, and race is in fact consistently admitted across schools in the EALS, this applicant is likely strong on unobserved characteristics like recommendation letters. Specifically, I construct an "inferred strength" variable for an application submitted by applicant i to school s equal to the mean across all applications submitted by applicant *i to schools other than s* of residuals from within-school regressions of admission on LSAT, GPA, race indicators, and time-period fixed effects.[24] Online Appendix Figure V illustrates the incremental predictive power of this inferred strength variable.

Column (6) of Table II reports the results of repeating the DDD specification of column (4) with the additional linear control of inferred strength. Both the Berkeley and UCLA results are nearly unchanged: a negative DDD effect of 34 percentage points, implying 30 percentage point (roughly 50%) declines in the black admission rate after accounting for space-opening effects. These are my preferred estimates because this specification uses all of the controls that are available for the full sample. These declines were much larger than those estimated in the EALS at any non-UC school, so the empirical p value of each of these declines across non-UC schools is 0. Dividing these 30 percentage point declines by the mean percentage point advantage enjoyed by pre-ban black applicants over whites with similar observed characteristics (detailed below in Section IV), these DDD estimates imply the ban reduced black-white admission rate differences by 59% at Berkeley and 56% at UCLA.

Finally, one may yet be concerned about differential selection on admission determinants that are specific to UC schools. A leading candidate for such a determinant is California residency, which

[23] Characteristics that are valued inconsistently across admissions offices include the applicant's geographic preference and intended legal specialty.

[24] Specifically, I fit: $\Pr(ADMITTED_{ist}) = \Phi(\beta_1 LSAT_i + \beta_2 GPA_i + \beta_3 BLACK_i + \beta_4 HISPANIC_i + \beta_5 ASIAN_i + \gamma_t)$ separately for each of seventeen schools in the EALS main sample and in each of two time periods (pre-ban and post-ban). I use all races here so that I can duplicate Table II in the online appendix for all races using the same underlying data. The resulting inferred strength variable ranges from -1 to 1. To flexibly handle the small share of applicants who applied to only one school, I assign their applications the same arbitrary inferred strength value and include an indicator for these applicants in all regressions where inferred strength is used.

positively predicts admission to UC schools in the EALS. Column (7) reports DD results including the California residency indicator as an additional control, using all applications for which the variable is available (see Section II.B). The estimates are similar to those of column (5), which uses the same specification but excludes California residency. I conclude that the affirmative action ban caused a large reduction in black admission advantages over similarly-credentialed whites at UC schools in the EALS.

Online Appendix Table III replicates Table II using applications from all races (white, black, Hispanic, and Asian); the results are very similar to those in Table II. Online Appendix Table IV replicates Online Appendix Table III while also fully interacting covariates $\mathbf{X}_i$ with race indicators, the post-ban indicator, and the non-UC indicator; the DD results are somewhat larger in magnitude (more negative) than those in Table II. Finally and for general reference, Online Appendix Table V displays OLS estimates of admission regressed on LSAT, GPA, race indicators, and school-year fixed effects for each school type and time period.

IV Continued Racial Disparities under the Ban

The previous section used difference-in-differences regressions to show that the ban substantially reduced black-white admission rate differences at Berkeley and UCLA, conditional on applicant characteristics. This section uses *cross-sectional* regressions to quantify *continued* black-white admission rate differences. It also sets the stage for Section V's discussion of mechanisms by discussing what kind of information may have been required in order to achieve the estimated post-ban differences.

IV.A Average Conditional Admission Rate Differences

Table III uses cross-sectional regressions among EALS applicants to report estimates of black-white admission rate differences at Berkeley and UCLA, conditional on observed covariates and averaged over the empirical distribution of black applicants' covariate levels. Panel B reports estimates using only post-ban applicants to each school and is the focus of this section; for reference, panel A repeats the analysis using pre-ban applicants. In order to obtain confidence intervals that lie inside the range of zero to one, bootstrapped 95% confidence intervals are reported below each estimate, based on one thousand bootstrapped samples for each school-time period.

Column (1) reports the actual observed black admission rate. Column (2) reports the hypothetical black admission rate estimated to prevail if all applicants were subjected to observed white admission standards. Column 3 reports the difference between columns (1) and (2), which I call the "average conditional admission rate difference" between blacks and whites because it is the estimated difference in admission rates between black applicants and observationally identical white applicants, averaged over the empirical distribution of black applicants.

I use the cross-sectional analogue to difference-in-differences Equation (1) to compute the estimates in column (2) of the hypothetical black admission rate under observed white admission standards. For each school and time period, I estimate the probit regression:

$$\Pr\left(ADMITTED_{it}\right) = \Phi(\mathbf{X}_i\boldsymbol{\alpha} + \beta_1 BLACK_i + \gamma_t) \tag{2}$$

where $\mathbf{X}_i$ is a vector of LSAT, GPA, and inferred strength based on independent admission decisions (see Section III.C for its definition and motivation) and γ_t are year fixed effects. I then use only the estimated coefficient vector $\hat{\boldsymbol{\alpha}}$ and the year fixed effects to compute a probit latent variable value for each applicant, and then to account for the fact that a decline in the black admission rate opens up space in the admitted cohort, I add a constant to every applicant's value until the mean predicted admission probability across applicants equals the actual admission rate among these applicants. Results are similar when omitting inferred strength.[25]

Panel B reports that whereas Berkeley actually admitted 31% of post-ban black applicants in the EALS, 13% are predicted to have been admitted under observed white admission standards. At UCLA, the estimates are 41% and 21%, respectively. These average conditional admission rate differences of 18 percentage points and 20 percentage points are very statistically significant; the actual black admission rate is higher than the hypothetical black admission rate in each of the one thousand bootstrapped samples of post-ban Berkeley applications (i.e. bootstrapped $p < 0.001$) and in all but one of the one thousand bootstrapped samples of post-ban UCLA applications (i.e. bootstrapped $p = 0.001$).

IV.B Maximum Conditional Admission Rate Differences

Applicants with very strong non-racial covariates (LSAT, GPA, and inferred strength) are accepted at near 100% rates regardless of race, and applicants with very weak non-racial covariates are accepted at near 0% rates regardless of race. Thus at either extreme, the black-white admission rate difference is nearly zero. But at intermediate covariate levels where applicants are neither all accepted nor all rejected ("the accept/reject margin"), black-white admission rate differences can be large. The *average* conditional admission rate differences reported above are an average among near-zero admission rate differences and large admission rate differences, reflecting the empirical distribution of black applicants'

[25] I do not include the California residency indicator as a control because it is unavailable for about half of the post-ban period. I obtain similar results when I include Hispanics and Asians in the regression along with Hispanic and Asian indicators.

covariate levels. This subsection estimates the largest black-white admission rate difference across covariate levels, which I call the "*maximum* conditional admission rate difference"—intuitively equal to the black-white admission rate difference among applicants who are on the margin of being accepted or rejected. I now formalize and discuss how these estimates have implications for the information used in admissions.

The maximum conditional admission rate difference can be easily visualized when considering only a scalar index of non-racial covariates. Figure IV displays probit-fitted post-ban black and white "admission rules" in academic strength, exactly as done in Figure Ic for pre-ban applicants.[26] The vertical distance between each school's black and white admission rules is an estimate of the black-white admission rate difference at each level of academic strength (equal to the probit marginal effect on the black indicator in Equation 2). The figure shows that at intermediate levels of academic strength, the estimated black-white admission rate difference is large. Intuitively, the ability of academic strength to predict within-race admission decisions (see Figure I) implies that admission rules are steep rather than flat, which is necessary for the econometrician to observe much larger black-white differences at intermediate levels of academic strength than at the extremes.

Column (4) of Table III presents estimates of the maximum conditional admission rate between blacks and whites, when Equation (2) is once again estimated using LSAT, GPA, and inferred strength as non-racial covariates. As a benchmark, Panel A reports that the maximum conditional admission rate difference in pre-ban admissions is estimated at 99 percentage points for both Berkeley and UCLA.[27] Panel B reports that in post-ban admissions, the estimated maximum differences are 57 percentage points and 69 percentage points, respectively, with lower bounds on the 95% confidence intervals of 37 and 33 percentage points and upper bounds of 76 percentage points and 99 percentage points, respectively.[28]

IV.C Discussion of Maximum Conditional Admission Rate Differences

Though it is a simple and potentially obvious point, I now explain how the large conditional admission rate differences estimated above can be informative of the racial information used in admissions, similar to an assessment made in the most recent Supreme Court affirmative action case (see Section V.C). Let $\mathbf{X}$ continue to denote the vector of observed characteristics that determine admission (LSAT, GPA, and inferred strength), $BLACK$ denote the black indicator, and $\mathbf{U}$ denote a vector of every

[26] That is, I estimate Equation (2) when $\mathbf{X}_i$ contains only the academic strength variable and then plot predicted admission probabilities for each race at each level of academic strength.

[27] Maximum conditional differences are 99 percentage points even when using academic strength as the only non-racial covariate, as depicted visually in Figure Ic.

[28] These maximum conditional differences are slightly smaller when using academic strength as the only non-racial covariate; see the notes to Figure IV for these estimates.

other characteristic that determines admission.[29] Consider the null hypothesis that race is not used in admissions. The null can be stated as: there exists a "total strength" function $s(\cdot,\cdot)$ such that the admission process can be represented as $ADMITTED_i = 1(s(\mathbf{X}_i, \mathbf{U}_i) > 0)$.[30] That is, admission is an arbitrary function of observed non-racial admission determinants $\mathbf{X}$ and unobserved non-racial admission determinants $\mathbf{U}$ but not of the black indicator $BLACK$.

Hence when race is not used in admissions, the black-white admission rate difference at any particular value of the observables $\mathbf{X} = \mathbf{x}$ equals the fraction of blacks with strong enough unobservables to be admitted, minus the fraction of whites with strong enough unobservables to be admitted:

$$\begin{aligned} (3) \qquad & E\left(ADMITTED_i|BLACK_i = 1, \mathbf{X}_i = \mathbf{x}\right) - E\left(ADMITTED_i|BLACK_i = 0, \mathbf{X}_i = \mathbf{x}\right) \\ = \quad & F_{WHITES,\mathbf{X}=\mathbf{x}}(0) - F_{BLACKS,\mathbf{X}=\mathbf{x}}(0) \end{aligned}$$

where $F_{WHITES,\mathbf{X}=\mathbf{x}}(\cdot)$ and $F_{BLACKS,\mathbf{X}=\mathbf{x}}(\cdot)$ denote the cumulative distribution functions of total strength $s(\mathbf{x}, \mathbf{U})$ for whites and blacks, respectively. This one-to-one correspondence under the null between black-white differences in the outcome and black-white differences *in the share of people* with strong unobservables is unique to binary outcomes like admission.[31] I refer to the right-hand side of Equation (3) as the "black-white stochastic dominance" along unobserved admission determinants $\mathbf{U}$ at a specific covariate value $\mathbf{x}$ that is implied under the null.

To connect the algebra here to the terminology used in the empirical analysis above, the left-hand side of Equation (3) is the "conditional black-white admission rate difference" at a specific covariate value $\mathbf{x}$. The previous subsection presented estimates of the maximum conditional admission rate difference across all values of $\mathbf{X}$ in the EALS. With sufficient data, one would ideally estimate the maximum conditional black-white admission rate difference non-parametrically by binning applicants into fine cells based on non-racial observables $\mathbf{X}$, computing black-white admission rate differences within each cell, and finding the maximum difference across cells. Because the EALS has limited observations when cutting the data by school, race, and covariate levels, the previous subsection made the probit functional form assumption with covariates entering linearly, which Section II.D suggests is reasonable in the EALS.

Empirically, conditional black-white differences can sometimes be too large to have been plausi-

[29] Empirical results are very similar when excluding inferred strength from $\mathbf{X}$. Excluding inferred strength from $\mathbf{X}$ implies that $\mathbf{U}$ includes every admission determinant outside of LSAT and GPA, rather than just those that are weighted uniquely weakly or strongly by UC admissions offices relative to non-UC admissions offices.

[30] The choice of zero as the admission threshold is without loss of generality.

[31] In particular, no such stochastic dominance condition holds for outcomes with an infinite range like earnings. As an extreme hypothetical example, a large black-white difference in mean earnings can be generated by an unobservable that causes an arbitrarily small fraction of one race to found blockbuster internet start-ups. Of course, the stochastic dominance condition holds for binary earnings outcomes such as having earnings above a certain threshold.

bly achieved using non-racial characteristics. For example, Table III reported that the maximum conditional black-white admission rate difference in pre-ban Berkeley and UCLA admissions was 99 percentage points, with relatively tight 95% confidence intervals. Under the null hypothesis of no use of race, the share of blacks with strong enough unobservables to be admitted must have been 99 percentage points larger than the corresponding share of whites. It seems highly unlikely that *any* combination of non-racial variables could be constructed to predominate so completely and exclusively among one race—strongly suggesting that pre-ban admission offices used race.

To clarify what drives large maximum conditional black-white differences empirically, it is worth noting that although the use of race was part of pre-ban admission offices' stated policies, there was nothing inevitable about the econometrician observing a near-100-percentage-point difference in the data. For example, if admission offices had weighted unobserved non-racial covariates (e.g. preferred legal specialty) heavily in the admission process, the non-racial covariates available in the EALS (LSAT, GPA, and inferred strength) would have done a worse job of predicting within-race admission decisions, causing the econometrician to observe more moderate admission rates at all covariate levels and thus smaller maximum black-white differences. Visually, the admission rules in Figure Ic would have been flatter, reducing the maximum vertical distance between black and white admission rules. The intuition is that powerful non-racial observables can allow the econometrician to observe the marginal cases in which race is decisive.

Reported in Table III, the maximum conditional black-white differences in post-ban admission rates are estimated at 57 percentage points at Berkeley and 63 percentage points at UCLA, with 95% confidence lower bounds close to 35 percentage points and upper bounds above 75 percentage points. These point estimates are large but at least more plausibly generated without the use of race or other prohibited factors. In the next subsection, I use auxiliary datasets to discuss the types of admission factors other than race that could have generated such large differences. Before doing so, it is worth noting that one could attempt to test the no-use-of-race null hypothesis in post-ban admissions with the Altonji, Elder, and Taber (2005) method for evaluating selection bias. Applied here, their exercise suggests that post-ban admission offices used race, but the admissions context would appear to violate the method's assumptions, so I do not pursue that strategy.[32]

[32] Roughly speaking in the current setting, Altonji et al. ask: would the econometrician still observe a black-white admission rate gap if unobserved determinants of admission were as correlated with being black as the observed determinants (LSAT, GPA, and inferred strength) are? The inferred strength variable is nearly orthogonal to race as constructed, but the strong *negative* correlation between black status and both LSAT and GPA (see Table I and Appendix Figures I and III) would suggest that, if anything, Table III underestimates the magnitude of black-white admission rate gaps. However, this exercise may be particularly unsuited to the present analysis because there are legal unobserved characteristics (e.g. low family income) that are known to correlate strongly *positively* with black status, unlike LSAT and GPA.

IV.D Black-Correlates Available to Admission offices

Useful for the next section's discussion of mechanisms, I now investigate the type of applicant information that post-ban admission offices may have been able to use to generate 63-percentage-point differences in black and white admission rates. The correlations presented here derive from data from outside the EALS and do not necessarily reflect the information listed on EALS applications, so this exploration must be viewed purely as providing basic context for interpreting the empirical results and not as testing hypotheses of actual admission office behavior.

The ban prohibited the use of race in admissions but allowed admission offices to increase admissions weight on characteristics correlated with black status ("black-correlates") that have plausible non-racial justification such as low family income.[33] A somewhat-coarse family income variable is available in a dataset closely related to the EALS: individual-level data on matriculants at the top fifty-two law schools in 1991.[34] Overall in this auxiliary dataset, the maximum black-white stochastic dominance in family income is 26 percentage points; within academic strength terciles, that maximum is 37 percentage points.[35]

Post-ban application forms contained other black-correlates that admission offices could have been used to legally increase black-white stochastic dominance along dimensions omitted from the EALS. Before the ban, Berkeley gave applicants ten short unconnected prompts for the personal statement, eight of which did not refer to diversity or disadvantages. Immediately after the ban and ever since, all ten were replaced by a single lengthy one that invited applicants to discuss their contributions to "the diversity of the entering class" and their backgrounds including "a personal or family history of cultural, educational, or socioeconomic disadvantage" (see Online Appendix Figure VI). In 1998, Berkeley added a full-page socioeconomic questionnaire to its application form requesting information

[33] This was not written in the text of the affirmative action ban, but this is almost certainly true based both on court rulings and on the fact that various UC schools took clear and immediate steps to increase admissions weight on black-correlates and were not litigated for noncompliance. Note that this follows legal distinctions made in discrimination lawsuits: "disparate treatment" (i.e. the direct use of race) in regulated decisions is always illegal, and policies that have racially "disparate impacts" (i.e. that substantially benefit one race over another) are typically considered illegal unless those policies have sufficient non-racial justification (Siskin and Trippi 2005; Selmi 2006).

[34] The dataset is called the Bar Passage Study (BPS) and was released by the Law School Admission Council—the same administrative entity that originally compiled the data in the EALS. LSAC required consent by schools and individuals for the BPS, so the full dataset contains approximately 27,000 of the 40,000 matriculants to U.S. law schools in 1991 (Wightman 1998). The data indicate law school tier; I confine attention to the most elite category. The data contain each matriculant's race, LSAT score, undergraduate GPA, and family income in five bins that capture substantial racial differences: from lowest income to highest income, the five bins contain 7%, 18%, 35%, 34%, and 6% of black students and 1%, 6%, 27%, 51%, and 15% of white students. The data do not contain information on law school applications or the undergraduate institution attended.

[35] I construct academic strength in the BPS the same way I construct it in the EALS: the standardized sum of standardized LSAT and standardized undergraduate GPA. Within academic strength quartiles, the maximum is 38 percentage points. Terciles are defined using black applicants only; when defining terciles using whites as well, the maximum is 33 percentage points.

such as college attendance rates of high-school friends and whether the applicant was raised by a single parent. Beginning in 2001, UCLA solicited declarations of interest in a Critical Race Studies program and instituted admission preference for interested applicants.

Post-ban application forms may also have contained strong black-correlates that would almost certainly be illegal to use in admissions because they lack plausible non-racial justification. At least Berkeley was stripped of the applicant race question from the application form (the application stated the page was diverted to a UC statistical office instead), and these schools rarely interviewed candidates, but extracurricular group participation listed on applicant résumés may have contained strong racial information. To investigate this, a research assistant used only pictures and names to subjectively identify as many black students as possible in the 1998 and 2004 yearbooks of the elite college, as well as a similar number of non-Hispanic white students.[36] The RA then used the student-provided information listed in the yearbooks to code whether the student was awarded a GPA-based honor and the number of extracurricular groups the student participated in that were explicitly dedicated to black issues or culture. Overall in this auxiliary dataset, I find that 72% of the elite college's black students participated in a black-focused extracurricular group, compared to 0% of white students; within the 60% of the sample that listed any GPA-based honor, the figures are 84% and 0%.[37] Applicant name (Bertrand and Mullainathan 2004; Fryer and Levitt 2004) may have provided additional racial information.

V Mechanisms

The previous two sections established this paper's core empirical results: the affirmative action ban substantially reduced but far from eliminated black-white admission rate differences at Berkeley and UCLA in the EALS. After considering whether changes in admission office preferences were likely responsible for these effects, this section uses a simple framework to briefly detail three informational and legal mechanisms that can explain the results. The data do not single out the particular mechanism (or mechanisms) that bound in practice. Instead, the focus of this section is on what one would have to believe about the admission technology, university objective function, and enforcement regime in order for each mechanism to explain the results and the consequent implications.

[36] The RA identified 193 black students and then identified a nearly equal number of white students by beginning at a randomly chosen student in each yearbook and searching for white-looking students in regular intervals of printed pictures.

[37] In addition to restricting the sample to applicants who were more likely to be competitive at top law schools, conditioning on the GPA-based honor excludes any "non-reporting" students who opted to report nothing in their yearbook entries.

V.A Changes in Constraints vs. Changes in Preferences

Both changes in constraints and changes in preferences can in principle explain changes in economic behavior. Though preference changes are often considered unrealistic, they are plausible here because UC administrators are formally employed by the State of California to fulfill constitutional mandates, including abiding by the affirmative action ban. In particular, UC administrators could have voluntarily abandoned preferences for racially diverse cohorts in the wake of the ban, or affirmative action proponents could simply have been replaced with affirmative action opponents.

The actions that Berkeley and UCLA admission offices took to collect and weight black-correlates in admissions after the ban (see Section IV.D) suggest that they continued to value racial diversity and searched for ways to increase it. Contemporary quotes corroborate this assessment. UC administrators strongly opposed the ban: as the California political climate turned against affirmative action in 1995, the UC president, UC vice-presidents, and the chancellor of each UC campus united to "unanimously urge, in the strongest possible terms," the continuation of affirmative action.[38] Berkeley's dean added "The need to diversify the legal profession is not a vague liberal ideal: it is an essential component to the administration of justice."[39] The day after voters approved the ban, the UC president announced that the question facing the university was "How do we establish new paths to diversity consistent with the law?"[40] One year after the ban, Berkeley's dean launched an audit of policies and procedures "to see whether we can achieve greater diversity" after "dire" admission results.[41] UC administrators were not systematically replaced in the subsequent years; for example Berkeley's dean and the UC president continued in their posts through 2000 and 2003, respectively. Christopher Edley, a vocal proponent of affirmative action and formal adviser to President Bill Clinton on the topic, has served as Berkeley's dean since 2004.

V.B Constrained Behavior

Assuming stable preferences, the empirical results reject the possibility that admission offices were unconstrained by the ban. Three mechanisms in a very simple model can explain the findings. The model adopts the basic setup of recent models of legal avoidance (Chan and Eyster 2003; Fryer, Loury, and Yuret 2007; Epple, Romano, and Sieg 2008) and considers alternative enforcement mechanisms

[38] 1995 "Statement Supporting Affirmative Action by UC President, Chancellors, and Vice Presidents", http://www.development.umd.edu/Diversity/Response/Action/policy.

[39] 1995 press release, http://www.berkeley.edu/news/berkeleyan/1995/0524/regents.html. (Recall that "Berkeley" and "UCLA" always refer in this paper to these campuses' law schools.)

[40] 1996 "Letter from President Richard C. Atkinson to the University Community Re: Passage of Proposition 209", http://www.universityofcalifornia.edu/news/
article/20607.

[41] 1997 Berkeley press release, http://berkeley.edu/news/berkeleyan/1997/0820/kay.html.

within this single framework. See Appendix B for a formal exposition and Online Appendix Figure 7 for an illustration. The key assumptions are that admission offices have preferences only over the racial diversity and the aggregate non-racial applicant strength (some combination of LSAT, GPA, recommendation letters, etc.) of admitted cohorts; whites stochastically dominate blacks along non-racial strength (consistent with schools practicing affirmative action and with the empirical distributions of LSAT and GPA across races in the EALS); admission offices can admit applicants on a combination of non-racial applicant strength and a binary signal of black status; admission offices can admit a fixed number of applicants; and applicant pools are held constant for simplicity.

(i) Diluted Racial Information. First, the ban may have placed no binding legal constraint on admission offices's use of race or other illegal admission factors like black extracurricular group participation, for example because courts may not be able to observe how decisions are arrived at behind closed doors. Instead, the ban may have substantially diluted the signal of black status available to admission offices, namely by stripping them of the applicant race question (as Berkeley's application form claims) and leaving them with access only to imperfect racial information gleaned from other parts of the application. Relative to a pure signal of black status, a diluted black signal increases the opportunity cost of admitting black applicants: admission offices must forego more non-racial strength in order to admit each additional black student, because by valuing the diluted black signal, admission offices will sometimes admit weak white applicants or reject strong black applicants. If preferences for racial diversity are not Giffen, a ban will cause admission offices to substitute away from racial diversity and toward non-racial strength, admitting blacks at a lower rate than before the ban but (as long as the diluted signal is not too weak) still at a higher rate than comparable whites.

Section IV.B estimated that post-ban admission offices used information in EALS admission decisions that predominated among black applicants approximately 63 percentage points more than among whites, and Section IV.D reported correlations from the elite college's yearbooks that raise the possibility that résumés contained additionally powerful black-correlates. The diluted racial information mechanism implies that the purer was the black signal available to admission offices, the smaller was the increase in the opportunity cost of admitting black applicants and thus the stronger the substitutability between racial diversity and non-racial strength had to be in order for diluted information to explain the large decline in black admission rates.

(ii) Whistleblower Threat. Second, post-ban admissions may have been forced through threat of litigation to abstain from using strong black-correlates. The use of black-correlates like black-group participation, black-sounding names, and revelations of black status through diversity essays would almost certainly be judged illegal under the ban because they lack credible non-racial justification, and

courts may have been able to observe the use of such information through the testimony of insiders with knowledge of how admission decisions were made ("whistle blowers"). Such a constraint would generate a larger increase in the opportunity cost to admitting black applicants than the diluted racial information mechanism, thus requiring weaker substitutability between racial diversity and aggregate non-racial strength in order to explain the empirical results. This mechanism would imply that legal black-correlates like low family income and diversity essays were strong enough to generate the large estimated racial disparities under the ban.

(iii) Cap on Measurable Disparities. Third and similar to the enforcement assumption in Coate and Loury (1993), the ban may have left the admission office unconstrained in its access to and use of racial information, so long as the black-white admission rate differences that can be measured by courts (e.g. conditional on LSAT, GPA, and inferred strength) were small enough to have plausibly been generated legally. The rationale would be that courts may find it difficult to observe the information used in admissions and thus may instead infer noncompliance from the extremeness of outcomes.[42] For example in the admissions context, a court may judge that the 99-percentage-point conditional black-white admission rate differences estimated in pre-ban admissions are unlikely to have been generated without the use of race, but that the approximately 63-percentage-point differences estimated in post-ban admissions could plausibly have been generated without the use of race. With continued knowledge of applicant race, optimal admission office behavior in the model is to reduce such measurable black-white admission differences down to the maximum level that could plausibly be generated without race—but to nevertheless use race to achieve that level, in violation of the ban. Thus a binding cap on measurable disparities would imply that post-ban admission offices continued to use race or strong illegal black-correlates, but modestly enough to escape litigation.

V.C Judicial Record

The three possible mechanisms differ in their legal mechanisms: diluted racial information implies that courts are powerless in enforcing affirmative action bans, whistleblower threat implies that the possibility of insider defection is strong, and a cap on measurable disparities implies that whistleblower threat is weak but that the threat of enforcement based on measurable outcome disparities is strong enough to limit noncompliance. Courts use a host of information in adjudicating discrimination lawsuits in general, yet in the case of affirmative action, the limited judicial record appears to have placed special emphasis so far on the magnitude of measurable black-white admission rate differences

[42] The evidentiary standard of proof in most discrimination and other civil cases is the "preponderance of the evidence", which is weaker than the "beyond reasonable doubt" standard used in criminal cases and which is often interpreted as merely probably guilt.

and implicitly on diluted racial information, rather than whistleblower threat.

First, the Supreme Court used conditional black-white differences to gauge the use of race in its most recent affirmative action case: "In an attempt to quantify the extent to which the [University of Michigan] Law School actually considers race in admission decisions," expert witness testimony provided "cell-by-cell comparisons between applicants of different races"—where the cells were defined by LSAT scores and undergraduate GPA, a non-parametric analogue to this paper's Figure IV and Table III—documenting that minority applicants "are given an extremely large allowance for admission" (Opinion of the Court in *Grutter v. Bollinger* 2003, internal quotations removed). Second, dissenting Supreme Court comments suggest that admission offices under a ban would "resort to camouflage" by using participation in "minority group associations" and related information to sustain high black admission rates "through winks, nods, and disguises."[43] This suggests that regulated admission offices would have to construct (possibly imperfect) racial proxies but could not be stopped from using illegal admission factors. Third, a kind of whistleblower report on alleged Proposition 209 noncompliance entitled "Report on Suspected Malfeasance in UCLA Admissions and the Accompanying Cover-Up" written by a former UCLA admissions faculty overseer (not himself an admission officer) has not led to litigation in the four years since its release, to the best of my knowledge (Groseclose 2008).

V.D Broader Predictions and Implications

This paper's empirical analysis deals exclusively with admission decisions in the EALS. However, the framework utilized in this section can in principle apply to other binary decisions governed by nondiscrimination mandates. Future empirical evidence could suggest whether each of the three enforcement regimes detailed above is quantitatively important outside of admissions.

Under a diluted information regime in which selectors have free reign to use their racial information and holding all else equal, one may expect smaller black-white differences when courts can deprive selectors of racial information (e.g. disallowing an applicant race question on a credit card application) than when it cannot (e.g. hiring). Under the whistleblower threat regime in which a knowledgeable insider can blow the whistle on the use of race, one may expect smaller black-white differences when a group makes a collective decision (e.g. a corporate board hiring a CEO) than when an individual makes one (e.g. a CEO hiring a manager), as well as when a firm operates on a large scale rather than

[43] Without recourse to affirmative action, "institutions of higher education may resort to camouflage. For example, schools may encourage applicants to write of their cultural traditions in the essays they submit, or to indicate whether English is their second language. Seeking to improve their chances for admission, applicants may highlight the minority group associations to which they belong, or the Hispanic surnames of their mothers or grandparents...If honesty is the best policy, surely Michigan's accurately described, fully disclosed College affirmative action program [sic] is preferable to achieving similar numbers through winks, nods, and disguises" (Ginsberg dissent, joined by Souter, *Gratz v. Bollinger* 539 U.S. 244 2003).

a small scale. Under the capped measurable disparity regime in which courts' ability to enforce a nondiscrimination mandate is limited by its ability to observe large black-white outcome differences, one may expect smaller black-white differences when courts have access to powerful observables (e.g. in mortgage lending, where courts have access to Home Mortgage Disclosure Act data on the universe of individual mortgage applications) than when it does not (e.g. hiring).

Answers could suggest policy changes. Under the capped measurable disparity regime, reporting requirements such as the Home Mortgage Disclosure Act can improve a court's ability to observe large black-white differences in outcomes and thus find a defendant at fault for noncompliance. Under the whistleblower threat regime, enhanced protections and incentives for whistleblowers could be valuable. But under the diluted information regime, policymakers may be limited in their ability to strengthen nondiscrimination law enforcement (e.g. because few employers could be deprived of applicant race information) and may instead choose to pursue other policies to narrow black-white outcome differences such as promoting skill formation (Heckman 1998, Loury 1998).

VI Conclusion

Debates over affirmative action hinge in part on schools' ability and willingness to avoid an affirmative action ban. This paper used application-level data on a large sample of graduates from an elite college in order to estimate the effect of the UC affirmative action ban on admission outcomes at UC Berkeley and UCLA law schools. The novelty of the analysis derived from having data on applications from before and after the ban and with rich enough covariates and independent screens (decisions at non-UC schools) to control for selective attrition and non-racial applicant strength. I found that the affirmative action ban reduced the black admission rate from 61% to 31%, well above the 8% rate that would prevail under observed white admission standards. Observed black admission advantages at intermediate credential levels were as large as 99 percentage points before the ban and 63 percentage points after the ban.

The results have several implications. Despite the potential enforcement frictions, mandating nondiscrimination in university admissions—a prominent part of the modern economy in which selectors use race—can substantially constrain the use of race behind closed doors. Simulations of affirmative action bans should assume large continued black-white admission disparities under a ban. Maintaining racial diversity under a ban may require forcing schools to sacrifice substantially more non-racial applicant strength than the schools themselves desire, potentially bearing on how "workable" are the alternatives to affirmative action and hence the practice's constitutionality. Methodologically, real-world cross-sectional data can in certain circumstances be rich enough to reveal discrimination.

References

1. Acemoglu, Daron and Joshua D. Angrist. 2001. "Consequences of Employment Protection? The Case of the Americans with Disabilities Act." Journal of Political Economy, 109(5): 915-957.

2. Altonji, Joseph G. and Rebecca M. Blank. 1999. "Race and Gender in the Labor Market." In Handbook of Labor Economics, Volume 3, ed. Orley Ashenfelter and David Card, 3143-3259. Amsterdam: North Holland.

3. Anwar, Shamena, Patrick Bayer, and Randi Hjalmarsson. 2012. "The Impact of Jury Race in Criminal Trials." Quarterly Journal of Economics, 127(2): 1017-1055.

4. Arcidiacono, Peter. 2005. "Affirmative Action in Higher Education: How Do Admission and Financial Aid Rules Affect Future Earnings?" Econometrica, 73:5: 1477-1524.

5. Arcidiacono, Peter, Esteban Aucejo, Patrick Coate, and V. Joseph Hotz. 2011. "The Effects of Proposition 209 on College Enrollment and Graduation Rates in California." Duke University *mimeo*, http://public.econ.duke.edu/~psarcidi/prop209.pdf.

6. Arrow, Kenneth. 1973. "The Theory of Discrimination." In Discrimination in Labor Markets, ed. Orley Ashenfelter and Albert Rees, 3-33. Princeton: Princeton University Press.

7. Arrow, Kenneth. 1988. "What Has Economics to Say about Racial Discrimination?" Journal of Economic Perspectives, 12(2): 91-100.

8. Becker, Gary. 1968. "Crime and Punishment: An Economic Approach." Journal of Political Economy, 76: 169-217.

9. Becker, Gary. 1957. The Economics of Discrimination. Chicago: University of Chicago Press.

10. Bertrand, Marianne, and Sendhil Mullainathan. 2004. "Are Emily and Greg More Employable than Lakisha and Jamal? A Field Experiment on Labor Market Discrimination." American Economic Review, 94(4): 991-1013.

11. Bowen, William G., and Derek Bok. 2000. "Chapter 2: The Admissions Process and Race Neutrality." In The Shape of the River, 15-52. Princeton: Princeton University Press, (Orig. pub. 1998).

12. Card, David, and Alan B. Krueger. 1994. "Minimum Wages and Employment: A Case Study of the Fast-Food Industry in New Jersey and Pennsylvania." American Economic Review, 84(4): 772-793.

13. Card, David, and Alan B. Krueger. 2005. "Would the Elimination of Affirmative Action Affect Highly Qualified Minority Applicants? Evidence from California and Texas." Industrial and Labor Relations Review, 58(3): 416-434.

14. Chan, Jimmy, and Erik Eyster. 2003. "Does Banning Affirmative Action Lower College Student Quality?" American Economic Review, 93(3): 858-872.

15. Chay, Kenneth Y. 1998. "The Impact of Federal Civil Rights Policy on Black Economic Progress: Evidence from the Equal Employment Opportunity Act of 1972." Industrial and Labor Relations Review, 51(4): 608-632.

16. Coate, Stephen, and Glenn Loury. 1993. "Antidiscrimination Enforcement and the Problem of Patronization." American Economic Review, 83(2): 92-98.

17. Dale, Stacy Berg, and Alan B. Krueger. 2002. "Estimating the Payoff to Attending a More Selective College: An Application of Selection on Observables and Unobservables." Quarterly Journal of Economics, 117(4):1491-1527.

18. Darity, William A. and Patrick L. Mason. 1998. "Evidence on Discrimination in Employment: Codes of Color, Codes of Gender." Journal of Economic Perspectives, 12(2): 63-90.

19. DiNardo, John, Nicole Fortin, and Thomas Lemieux. 1996. "Labor Market Institutions and the Distribution of Wages, 1973-1992: A Semiparametric Approach." Econometrica, 64(5): 1001-1044.

20. Dobkin, Carlos, and Nancy Nicosia. 2009. "The War on Drugs: Methamphetamine, Public Health, and Crime." American Economic Review, 99(1): 324-349.

21. Epple, Dennis, Richard Romano, and Holger Sieg. 2008. "Diversity and Affirmative Action in Higher Education." Journal of Public Economic Theory, 10(4): 475-501.

22. Espenshade, Thomas J., Chang Y. Chung, and Joan L. Walling. 2004. "Admission Preferences for Minority Students, Athletes, and Legacies at Elite Universities." Social Science Quarterly, 85(5): 1423-1446.

23. Feldstein, Martin. 1995. "Behavioral Responses to Tax Rates: Evidence from the Tax Reform Act of 1986." American Economic Review, 85(2): 170-174.

24. Freeman, Richard. 1973. "Changes in the Labor Market for Black Americans, 1948-72." Brookings Papers on Economic Activity, 1973(1): 67-131.

25. Fryer, Roland. 2011. "Racial Inequality in the 21st Century: The Declining Significance of Discrimination." In Handbook of Labor Economics, Volume 4, ed. Orley Ashenfelter and David Card, 855-971. Amsterdam: North Holland.

26. Fryer, Roland, and Steven Levitt. 2004. "The Causes and Consequences of Distinctively Black Names." Quarterly Journal of Economics, 119(3): 767-805.

27. Fryer, Roland, Glenn Loury, and Tolga Yuret. 2007. "An Economic Analysis of Color-Blind Affirmative Action." Journal of Law, Economics, and Organization, 24(2): 319-355.

28. Goldin, Claudia and Cecilia Rouse. 2000. "Orchestrating Impartiality: The Impact of 'Blind' Auditions on Female Musicians." American Economic Review, 90(4): 715-741.

29. Groseclose, Tim. 2008 "Report on Suspected Malfeasance in UCLA Admissions and the Accompanying Cover-Up." UCLA *mimeo*, http://www.sscnet.ucla.edu/polisci/faculty/groseclose/CUARS.Resignation.Report.pdf.

30. Heckman, James J. 1998. "Detecting Discrimination." Journal of Economic Perspectives, 12(2): 101-116.

31. Heckman, James J., and Brook S. Payner. 1989. "Determining the Impact of Federal Antidiscrimination Policy on the Economic Status of Blacks: A Study of South Carolina." American Economic Review. 79(1): 138-177.

32. Heckman, James J. and J. Hoult Verkerke. 1990. "Racial Disparity and Employment Discrimination Law: An Economic Perspective." Yale Law and Policy Review, 8(2): 276-298.

33. Hinrichs, Peter. 2012 “The Effects of Affirmative Action Bans on College Enrollment, Educational Attainment, and the Demographic Composition of Universities,” Review of Economics and Statistics, 94(3): 712-722.

34. Kane, Thomas. 1998. “Racial and Ethnic Preference in College Admissions.” In The Black-White Test Score Gap, ed. Christopher Jencks and Meredith Phillips, 431-456. Washington: Brookings Institution.

35. Kling, Jeffrey, Jeffrey Liebman, and Lawrence Katz. 2007. “Experimental Analysis of Neighborhood Effects.” Econometrica, 75(1): 83-119.

36. Krueger, Alan B, Jesse Rothstein, and Sarah Turner. 2006. “Race, Income, and College in 25 Years: Evaluating Justice O’Connor’s Conjecture.” American Law ans Economics Review, 8(2): 282-311.

37. Leonard, Jonathan S. 1984. “The Impact of Affirmative Action on Employment.” Journal of Labor Economics, 2(4): 439-463.

38. List, John A. 2004. “The Nature and Extent of Discrimination in the Marketplace: Evidence from the Field.” Quarterly Journal of Economics, 119(1): 49-89.

39. Long, Mark, and Marta Tienda. 2008. “Winners and losers: Changes in Texas university admissions post-Hopwood.” Educational Evaluation and Policy Analysis, 30(3): 255-280.

40. Loury, Glenn. 1998. “Discrimination in the Post-Civil Rights Era: Beyond Market Interactions.” Journal of Economic Perspectives, 12(2): 117-126.

41. McCrary, Justin. 2007. “The Effect of Court-Ordered Hiring Quotas on the Composition and Quality of Police.” American Economic Review, 97(1): 318-353.

42. Miller, Amalia R. and Carmit Segal. Forthcoming. “Does Temporary Affirmative Action Produce Persistent Effects? A Study of Black and Female Employment in Law Enforcement.” Review of Economics and Statistics.

43. National Resource Council. 2004. “Statistical Analysis of Observational Data.” In Measuring Racial Discrimination, ed. Rebecca M. Blank, Marilyn Daddy, and Constance F. Citro, 118-161. Washington: National Academies Press.

44. Oyer, Paul, and Scott Schaefer. 2002a. “Sorting, Quotas, and the Civil Rights Act of 1991: Who Hires When It’s Hard to Fire?” Journal of Law and Economics, 45(1): 41-68

45. Oyer, Paul, and Scott Schaefer. 2002b. “Litigation Costs and Returns to Experience.” American Economic Review, 92(3): 683-705.

46. Phelps, Edmund. 1972. “Statistical Theory of Racism and Sexism.” American Economic Review, 62(4): 659-661.

47. Price, Joseph and Justin Wolfers. 2010. “Racial Discrimination among NBA Referees.” Quarterly Journal of Economics, 125(4): 1859-1887.

48. Ramsey, Frank. 1927. “A Contribution to the Theory of Taxation.” Economic Journal, 37(March): 47-61.

49. Rothstein, Jesse and Albert H. Yoon. 2008. “Affirmative Action in Law School Admissions: What Do Racial Preferences Do?” *University of Chicago Law Review*, 75: 649-714.

50. Sander, Richard H. “A Systematic Analysis of Affirmative Action in American Law Schools.” *Stanford Law Review*, 57: 368-483.

51. Selmi, Michael. 2000. “Why Are Employment Discrimination Cases So Hard to Win?” Louisiana Law Review, 61: 555-575.

52. Selmi, Michael. 2006. “Was the Disparate Impact Theory a Mistake?” UCLA Law Review, 53(3): 701-782.

53. Siskin, Bernard, and Joseph Trippi. 2005. “Statistical Issues in Litigation.” In Employment Discrimination Litigation, ed. Frank Landy, 132-166. San Francisco: Jossey-Bass.

54. Stigler, George J. 1946. “The Economics of Minimum Wage Legislation.” American Economic Review, 36(3): 358-365.

55. University of Texas at Austin Office of Admissions. 2005. “Inter-rater Reliability of Holistic Measures Used in the Freshman Admissions Process of The University of Texas at Austin.” http://www.utexas.edu /student/admissions/research/Inter-raterReliability2005.pdf (last accessed October 9, 2012).

56. University of Texas at Austin Office of Admissions. 2007. “Report 10.” http://www.utexas.edu /student/admissions/research/HB588-Report10.pdf (last accessed October 9, 2012).

57. Wightman, Linda F. 1998. “LSAC National Longitudinal Bar Passage Study.” *LSAC Research Report Series*, originally accessed online and now available by email via Linda Reustle (lreustle@lsac.org).

58. Yinger, John. 1998. “Evidence on Discrimination in Consumer Markets.” Journal of Economic Perspectives, 12(2): 23-40.

Online Appendix A: Details of EALS Data Coding

The first application year's LSAT scores are in a more compact scale than all other years', and I convert them to the modern scale using percentile rank. I de-mean GPA by year to account for modest grade inflation over time. I code "Chicano/Mexican-American", "Hispanic", and "Puerto Rican" as Hispanic. Undergraduate major is available in some years' raw data; it has low statistical power in subsamples and its use would limit the years available for analysis so I omit it. The admission decision for a small percentage of accepted students is classified as rejected when the applicant in fact accepted and deferred an admission offer. The relatively minor importance of this measurement error is suggested visually in Figure Ib, where actual admission rates are close to 100% at high levels of academic strength, rather than plateauing at a smaller number; admission rates are closer to 100% when inferred strength is also used to rank applicants. Year of college graduation is available in all years; I omit it for simplicity but every result holds when also controlling for a quartic in graduation year. The only other information in the raw data are indicators for whether the applicant took the LSAT more than once, whether the applicant withdrew an application before an admission decision was made, and whether the applicant accepted an admission offer. I exclude withdrawn applications from the analysis, and I do not have sufficient power to analyze matriculation decisions.

The raw data do not contain applicant identifiers, so for each year I create applicant identifiers by treating as coming from the same applicant those applications that match on all of the application-invariant variables; this is a powerful method in large part because GPA is coded to two decimal places. I exclude the fewer than one percent of observations for which this implies that a single applicant submitted multiple applications to the same school.

I do not include the University of Michigan in the group of fifteen most-applied-to schools because it was subject to an affirmative action ban during the sample. I do not analyze Michigan as a treatment school because its bans were effective during the sample only in 2001 and 2006 and I do not have sufficient power to conduct year-by-year difference-in-differences. UC law schools at Davis and Hastings as well as public Texas law schools received few applications in the EALS and similarly do not permit robust inference.

Online Appendix B: A Simple Model of Behavior under an Affirmative Action Ban

I use a simple version of recent avoidance models (Chan and Eyster 2003; Fryer, Loury, and Yuret 2007; Epple, Romano, and Sieg 2008) to characterize optimal admission office behavior under an affirmative action ban, depending on the enforcement regime. The analysis uses terminology specific to admission decisions under an affirmative action ban but applies generally to acceptance decisions under nondiscrimination laws.

(i) The Admission office's Maximization Problem. The simplest way to model the admission office's maximization problem is to cast it as a simple two-good consumption problem: the admission office has concave preferences over the number of black applicants admitted $\overline{r}$ and the aggregate non-racial strength of the admitted cohort. Each applicant is either black or white, the applicant pool is the same pre-ban and post-ban, and all admitted students matriculate.[44] The admission office faces a binding capacity constraint: it can admit no more than a fixed number $\bar{N}$ of applicants and must reject some applicants.

[44] In a broader model where not all admitted students matriculate, the yield (the percentage of admitted students who matriculate) may be a relevant margin for understanding admissions office behavior. Using aggregate data covering all applicants to UC law schools, the black yield at UCLA rose a few percentage points relative to the white yield after the ban; at Berkeley it declined in the first year of the ban but then recovered the following year and rose slightly over time.

The admission office's problem is:

$$\max_{\overline{r},\overline{s}} u(\overline{r},\overline{s}) \quad \text{s.t.} \quad N(\overline{r},\overline{s}) \leq \bar{N}$$

where $N(\overline{r},\overline{s})$ is the minimum number of applicants that must be admitted in order to deliver $\overline{r}$ black admits and $\overline{s}$ aggregate non-racial strength. $N(\overline{r},\overline{s})$ is an implicit function of the joint distribution of race and non-racial strength in the applicant pool. The admission office faces a tradeoff in that the admission rule that maximizes the number of black admits is not the one that maximizes aggregate non-racial strength.

The admission office can admit applicants on the basis of two pieces of information: non-racial strength s_i and a binary signal $BLACKSIGNAL_i \in \{0,1\}$ of black status. The optimal admission rule can always be characterized as a "rank-and-yank" rule that admits the $\bar{N}$ applicants that have highest rank according to:

$$rank_i = s_i + \lambda BLACKSIGNAL_i$$

where λ is chosen to maximize admission office utility. This is true because for any number of admitted blacks, the admission office maximizes aggregate non-racial strength by adopting a threshold rule within each black signal whereby the only admitted applicants are black-signalled applicants with non-racial strength above some $s^*_{BLACKSIGNAL=1}$ and white-signalled applicants with non-racial strength above some $s^*_{BLACKSIGNAL=0}$. Rank-and-yank implements any such pair of threshold rules by setting weight λ equal to $s^*_{BLACKSIGNAL=0} - s^*_{BLACKSIGNAL=1}$.

(ii) Affirmative Action. When affirmative action is not banned, the admission office is permitted arbitrary use of a pure signal of race in admission decisions. The black signal is pure in that $BLACKSIGNAL_i = 1$ if and only if applicant i is black. Online Appendix Figure VIIa illustrates a feasible pair of optimal admission thresholds and illustrates its consequences for black and white applicants. To define the no-affirmative-action benchmark, let s^* be the level of non-racial strength above which there are exactly $\bar{N}$ applicants. This is the race-neutral threshold that would maximize aggregate non-racial strength and corresponds to a rank-and-yank admission rule with $\lambda = 0$. An admission office practicing affirmative action chooses $\lambda > 0$ and thus adopts a threshold admission rule for blacks at $s^*_{BLACKSIGNAL=1}$ and a separate threshold for whites at $s^*_{BLACKSIGNAL=0}$. Relative to the no-affirmative-action benchmark, the admission office practicing affirmative action admits extra blacks (the grid fill pattern) and rejects extra whites (the solid fill pattern).

Online Appendix Figure VIIc illustrates the affirmative action budget set in $(\overline{r},\overline{s})$ space for the simple case of uniform distributions of non-racial strength within each race. The range of weights $\lambda \in [0,\infty)$ traces out the budget constraint (the solid curve). Point A is a potentially optimal bundle under affirmative action. The budget constraint is strictly convex because the first black applicant admitted through affirmative action is almost as strong as the white applicant that must be rejected in order to make room. After that, stronger and stronger white applicants must be rejected to make room for weaker and weaker black applicants.

(iii) Diluted Racial Information and Whistleblower Threat. The letter of an affirmative action ban prohibits the admission office from using a pure signal of race in admission decisions but allows it to use race-blind factors like low family income that correlate imperfectly with race and that have plausible non-racial justification. This could be implemented in two ways: by preventing the admission office from acquiring pure racial information on applicant race, or by the threat of litigation for using pure racial information such as through an admission officer insider "blowing the whistle" and testifying in court against the admission office. Either way, implementation of the letter of the law dilutes the usable racial information available to the admission office.

I model this dilution as fraction p_{black} of black applicants and fraction p_{white} of white applicants possessing the binary signal $BLACKSIGNAL_i$, with $p_{black} - p_{white} < 1$ and for simplicity $p_{black}, p_{white} \perp s_i$. Online Appendix Figure VIIb illustrates that an admission office placing weight on an impure black signal makes "mistakes": the admission office is forced to reject some applicants that have higher non-racial strength than accepted applicants of the same race.

By forcing the admission office to make mistakes, an affirmative action ban raises the opportunity cost of admitting black applicants. It can be easily shown that, in the analytically tractable case of uniform distributions of non-racial strength within race,[45] the dilution of black signal purity raises the marginal rate of transformation of admitted blacks for non-racial strength by a factor that is decreasing in the purity of the signal $BLACKSIGNAL_i$:

$$\frac{MRT^{DI/WT}_{\bar{r},\bar{s}}}{MRT^{AA}_{\bar{r},\bar{s}}} = \frac{1}{(p_{black} - p_{white})^2} > 1$$

The higher opportunity cost puts the diluted racial information and whistleblower threat ("DRI/WT") budget set in the interior of the affirmative action budget set, illustrated by the constant-length dashed line in Online Appendix Figure VIIc. As in any two-good consumption problem when the price of one good rises, changes in the consumption bundle hinge on income and substitution effects and are indeterminate when utility is unspecified. Unless preferences are Giffen, the optimal post-ban bundle under diluted racial information or whistleblower threat involves fewer admitted black applicants. Bundle B is one such possible bundle.[46]

(iv) Cap on Measurable Disparities. Online Appendix Figure VIIc also depicts an alternative enforcement regime: unable to enforce the letter of the law because the information used in admissions may be unobservable, courts may instead impose a *de facto* limit on the black-white admission rate difference that can be measured by courts (e.g. conditional on LSAT, GPA, and inferred strength) or, equivalently here, the total number of admitted blacks. This constraint—whether known with or without certainty by the admission office—creates a kink in the admission office's budget constraint, and the optimal response of a post-ban admission office with a pure signal of race is to continue using race, only more modestly than before the ban.[47] This lands the admission office at a bundle like C where aggregate non-racial strength rises and the number of admitted blacks falls.

[45] Without this or a similar assumption, the budget set can be non-convex over some intervals.

[46] Earlier avoidance models put structure on the admissions office's preferences and thereby generate specific directional predictions. In Chan and Eyster (2003), $\bar{r}$ and $\bar{s}$ enter separably and linearly. Under this and technology restrictions, the admissions office may respond to a ban by deliberately introducing idiosyncratic noise—an imperfect racial proxy when blacks are concentrated at lower levels of the non-racial strength distribution—into admissions decisions and generate non-Giffen outcomes. Fryer, Loury, and Yuret (2007) assume that the post-ban admissions office uses imperfect racial proxies to admit the same number of black applicants as it did pre-ban.

[47] Even if the *de facto* limit is not known with certainty, the admissions office chooses the number of extra black applicants that optimally trades off diversity goals and litigation risk and then achieves that number using race in admissions as if that number were the actual court-enforced maximum.

TABLE I
Mean Applicant Characteristics by Race

	Share of applicants	LSAT score (sd 6.7)	Undergraduate GPA (sd 0.33)	Academic strength (mean 0, sd 1)	Admission rate
A. All Applicants (N = 5,353, collectively submitting 25,499 applications to top-17 schools)					
White	60.8%	167.3	3.47	0.24	41%
Black	9.7%	159.9	3.15	-0.98	56%
Asian	19.4%	167.6	3.52	0.33	41%
Hispanic	10.1%	162.8	3.31	-0.48	39%
B. Applicants to Berkeley (N = 1,594)					
White	56.6%	167.5	3.47	0.23	31%
Black	8.0%	160.8	3.13	-0.92	43%
Asian	24.2%	167.0	3.49	0.21	36%
Hispanic	11.3%	162.3	3.31	-0.53	34%
C. Applicants to UCLA (N = 777)					
White	55.0%	165.4	3.38	-0.09	54%
Black	7.5%	159.6	3.03	-1.17	53%
Asian	24.5%	165.2	3.43	-0.06	60%
Hispanic	13.1%	159.8	3.23	-0.89	35%

Notes - Panel A lists mean applicant characteristics for the Elite Applications to Law School sample used in the paper. The sample comprises the 5,353 applicants who together submitted 25,499 applications over seventeen years to Berkeley, UCLA, and the top-fifteen law schools that were never subject to an affirmative action ban. LSAT is the standardized test score used in law school admissions and ranges from 120 to 180. Undergraduate grade point average is the cumulative undergraduate GPA on a 4.00 scale. "Academic strength" is a scalar index of the strength of an applicant's academic credentials, equal to the standardized (mean zero and standard deviation one) sum of standardized LSAT and standardized GPA (see Figure I for the semi-parametric motivation). Panels B and C list the same statistics for applicants to Berkeley only and UCLA only, respectively. Online Appendix Table I lists summary statistics on application behavior and comparisons to the nationwide population of law school applicants.

TABLE II
Effect of the Ban on Black Admission Rates

Dependent Variable:	Admission						
	OLS		Probit (average marginal effect)				
	(pp)	(pp)	(pp)	(pp)	(pp)	(pp)	(pp)
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
A. Berkeley Difference-in-differences							
Black × Post-ban	-39.7	-31.8	-40.5	-35.5	-40.0	-33.9	-43.6
	(6.8)	(7.3)	(5.4)	(6.3)	(5.3)	(6.5)	(5.5)
National trend controls		x		x		x	
Inferred strength control					x	x	x
CA residency control							x
N (applications)	1,029	17,329	1,029	17,329	1,029	17,329	779
Clusters (applicants)	1,029	3,754	1,029	3,754	1,029	3,754	779
Actual pre-ban black admission rate	56.7	56.7	56.7	56.7	56.7	56.7	56.7
Δ implied by Black × Post-ban effect	-34.2	-27.4	-35.7	-31.1	-35.6	-30.0	-39.7
B. UCLA Difference-in-differences							
Black × Post-ban	-48.1	-41.6	-35.3	-32.8	-35.0	-33.5	-31.1
	(10.5)	(10.4)	(11.1)	(11.0)	(11.2)	(11.1)	(10.5)
National trend controls		x		x		x	
Inferred strength control					x	x	x
CA residency control							x
N (applications)	485	16,785	485	16,785	485	16,785	371
Clusters (applicants)	485	3,736	485	3,736	485	3,736	371
Actual pre-ban black admission rate	64.5	64.5	64.5	64.5	64.5	64.5	64.5
Δ implied by Black × Post-ban effect	-41.6	-35.9	-32.1	-29.2	-32.0	-30.2	-29.2

Notes - Each column reports a coefficient from a difference-in-differences (DD) regression using black and white applications in the Elite Applications to Law School dataset at Berkeley (panel A) or UCLA (panel B). Standard errors clustered by applicant are in parentheses. Columns (1)-(2) use OLS regressions while the remaining columns use probit regressions and report marginal effects averaged over the UC school's pre-ban black applicants. The odd-numbered columns use the DD specification of Equation (1): admission regressed on a black indicator, a black indicator interacted with a post-ban indicator, year fixed effects, LSAT, and GPA; column (5) additionally controls for inferred strength, and column (7) additionally controls for a California residency indicator that is available only for applications to UC schools and in certain years. "Inferred strength" uses independent admission decisions to proxy for admission determinants like recommendation letters that are omitted from the EALS, similar to Dale and Krueger (2002); see Section III.B or Online Appendix Figure V for details. The even-numbered columns use a DDD version of Equation (1) that controls for national trends by including in the regression all applications to the top-fifteen non-UC schools and by interacting the black indicator and the black-times-post-ban interaction with a UC school indicator; the coefficient on this latter interaction is reported. These regressions include school-year fixed effects and are weighted so that each school receives equal weight in each time period (pre-ban and post-ban). The final row in each panel reports estimates of the change in the admission rate that pre-ban black applicants are predicted to have experienced had the ban been effect, accounting for the space-opening effect of a decline in black admission rates. Each estimate is computed by using the UC-specific coefficients of each regression to compute a probit latent variable value for each black and white pre-ban applicant according to post-ban criteria, and then adding a constant to every applicant's value until the mean predicted admission probability across applicants equals the actual admission rate observed among these applicants. Online Appendix Tables III and IV replicate this table using alternative specifications.

TABLE III
Black-White Admission Rate Differences in the Pre-ban and Post-ban Cross Sections

	Actual black admission rate	Hypothetical black admission rate under white coefficients	Average conditional black-white admission rate difference (col. 1 minus col. 2)	Maximum conditional black-white admission rate difference across covariate values
	(%)	(%)	(pp)	(pp)
	(1)	(2)	(3)	(4)
A. Pre-ban				
Berkeley	56.7	5.6	51.1	99.1
	[43.6, 69.5]	[1.2, 11.4]	[38.7, 62.5]	[97.1, 100.0]
UCLA	64.5	10.4	54.1	98.8
	[46.7, 80.6]	[2.2, 21.0]	[37.0, 70.5]	[92.5, 100.0]
B. Post-ban				
Berkeley	31.3	13.5	17.8	56.8
	[20.4, 43.4]	[7.1, 20.6]	[9.3, 27.0]	[36.8, 75.6]
UCLA	40.7	21.1	19.6	68.7
	[23.1, 60.0]	[7.9, 37.6]	[6.2, 34.1]	[33.6, 98.9]

Notes - Each cell reports an estimate of either a black admission rate or a black-white admission rate difference using the Elite Applications to Law School dataset. Ninety-five percent confidence intervals are computed using one thousand bootstrapped samples of each school-time period and are listed in brackets. Only black and white applications are used. Column (1) lists the actual black admission rate in the specified school-time period. Column (2) reports the black admission rate that is predicted to have prevailed if black applicants had been subjected to observed white admission standards, calculated by estimating Equation (2) which is a probit regression of admission on LSAT, GPA, inferred strength, a black indicator, and year fixed effects and then using the coefficients other than on the black indicator to predict admission probabilities for each applicant and accounting for the minor space-opening effect of a decline in black admission rates as described in Section IV.A (results are similar without the correction). Reported estimates are means of these predict admission probabilities. See Section III.C or Online Appendix Figure V for the definition of inferred strength; results are similar when omitting it. Column (3) equals the difference between columns (2) and (1) and is an estimate of the average black-white admission rate difference for this school-time period's black applicants, conditional on observed covariates. Empirically, applications with high levels of LSAT, GPA, and inferred strength are accepted at high rates regardless of race, and applications with low levels are accepted at low rates regardless of race. But at intermediate covariate levels, the black-white admission rate difference is large. Column (4) reports an estimate of the maximum black-white admission rate difference conditional on covariates, equal to largest probit marginal effect on the black indicator across covariate levels. See Figure IV for an illustration of these maximum conditional differences and Figure Ib for an illustration of the reasonableness of the probit functional form in EALS decisions.

ONLINE APPENDIX TABLE I
Application Behavior and Comparison of Applicant Characteristics

A. Application Behavior in the Full EALS Dataset, 1990-2006

Applications per applicant	5.7
Applications per applicant who applied to Berkeley or UCLA	7.8
Percent of applications sent to schools ranked 1-10	59%
Percent of applications sent to schools ranked 11-20	20%
Percent of applicants who applied to Berkeley	28%
Percent of applicants who applied to UCLA	14%

B. Applications and Applicants in the 17-School EALS Sample Used in the Paper

Applications	25,499
Applicants	5,353
Applications and applicants to Berkeley (7th-most in the 17-school sample)	1,594
Applications and applicants to UCLA (13th-most in the 17-school sample)	777

C. Mean Applicant Characteristics in the 17-School EALS Sample Used in the Paper and Nationwide

	EALS (sd)	Nationwide
LSAT	166.2	151.5
	(6.7)	
GPA	3.43	3.16
	(0.33)	
White	60.8%	70.9%
Asian	19.4%	7.7%
Black	9.7%	12.4%
Hispanic	10.1%	9.1%
Post-ban	54.8%	

Notes - Panel A lists statistics on the application behavior of Elite Application to Law School applicants, using all complete observations (32,627 applications from 5,692 applicants). The rankings refer to the rankings from the 1998 issue of U.S. News and World Report's "America's Best Graduate Schools", which ranked Berkeley seventh and UCLA seventeenth out of 174 law schools. Panel B lists statistics on applications submitted to the seventeen law schools used in the paper; see the notes to Table I for details. Panel C lists mean applicant characteristics. The Nationwide column lists statistics for all U.S. law school applicants in application year 2000-2001, the closest available year to the midpoint of the EALS sample. LSAT is the standardized test score used in law school admissions and ranges from 120 to 180. Undergraduate grade point average is the cumulative undergraduate GPA on a 4.00 scale. "Hispanic" includes applicants classified as Chicano/Mexican-American, Hispanic, and Puerto Rican. Post-ban is an indicator for the applicant applying to law school in application year 1996-1997 or later. The 5.4% of EALS applicants who do not report race or list their race as American Indian/Alaskan Native, Canadian Aboriginal, or Other are omitted from EALS statistics in this table and all analyses. The corresponding 7.9% of U.S. applicants are omitted from the U.S. applicant race percentages as well. The nationwide data were collected from various tables at http://www.lsac.org/LSACResources/default.asp, accessed on 6/6/2010.

ONLINE APPENDIX TABLE II
Effect of the Ban on Black Admission Rates
Semi-Parametric Estimates

Admission Rates at Non-UC Schools

	Pre-ban	Post-ban	Difference (pp)
White	40.6%	46.1%	5.5
Black	61.2%	63.0%	1.9
Difference (pp)	20.6	16.9	-3.6

Admission Rates at Berkeley

	Pre-ban	Post-ban	Difference (pp)
White	31.0%	33.9%	2.9
Black	56.7%	26.0%	-30.6
Difference (pp)	25.7	-7.9	-33.6

DDD estimate (percentage points): -29.9
DDD estimate, as % of pre-ban black admission rate: -52.8%

Admission Rates at UCLA

	Pre-ban	Post-ban	Difference (pp)
White	48.0%	60.1%	12.2
Black	64.5%	32.4%	-32.1
Difference (pp)	16.6	-27.7	-44.3

DDD estimate (percentage points): -40.7
DDD estimate, as % of pre-ban black admission rate: -63.0%

Notes - This table constructs the semi-parametric difference-in-differences-in-differences (DDD) estimates of the change in black admission rates at Berkeley and UCLA reported in Section III.B. Each pre-ban admission rate is an actual admission rate. Each post-ban admission rate is a reweighted estimate of the admission rate that pre-ban applicants of each race and school are predicted to have experienced after the ban. See the notes to Figure III for the reweighting procedure. The differences computed in the DDD are between pre-ban and post-ban periods, UC and non-UC schools, and black and white races. The non-UC schools are the top-fifteen schools in the EALS that were never subject to an affirmative action ban. See Table II for analogous parametric DDD estimates that account for the fact that a decline in black admission rates opens up space in the admitted cohort for members of both races.

ONLINE APPENDIX TABLE III
Effect of the Ban on Black Admission Rates
Using Applications from All Races

Dependent Variable:	Admission						
	OLS		Probit (average marginal effect)				
	(pp)	(pp)	(pp)	(pp)	(pp)	(pp)	(pp)
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
A. Berkeley Difference-in-differences							
Black × Post-ban	-38.1	-30.8	-39.8	-34.8	-38.6	-32.7	-42.5
	(6.9)	(7.3)	(5.6)	(6.4)	(5.5)	(6.7)	(5.8)
National trend controls		x		x		x	
Inferred strength control					x	x	x
CA residency control							x
N (applications)	1,594	24,722	1,594	24,722	1,594	24,722	1,197
Clusters (applicants)	1,594	5,324	1,594	5,324	1,594	5,324	1,197
Actual pre-ban black admission rate	56.7	56.7	56.7	56.7	56.7	56.7	56.7
Δ implied by Black × Post-ban effect	-34.6	-28.0	-36.5	-31.7	-36.0	-30.3	-40.1
B. UCLA Difference-in-differences							
Black × Post-ban	-45.4	-39.0	-35.9	-32.1	-35.1	-32.0	-33.2
	(10.4)	(10.5)	(10.3)	(10.8)	(10.3)	(10.8)	(10.2)
National trend controls		x		x		x	
Inferred strength control					x	x	x
CA residency control							x
N (applications)	777	23,905	777	23,905	777	23,905	586
Clusters (applicants)	777	5,300	777	5,300	777	5,300	586
Actual pre-ban black admission rate	64.5	64.5	64.5	64.5	64.5	64.5	64.5
Δ implied by Black × Post-ban effect	-41.4	-35.5	-34.6	-30.4	-34.6	-30.9	-33.0

Notes - This table replicates Table II using applications from all races (black, white, Asian, and Hispanic). The regressions underlying this table are the same as those underlying Table II except that the black indicator is replaced by a vector of black, Asian, and Hispanic indicators.

ONLINE APPENDIX TABLE IV
Effect of the Ban on Black Admission Rates
Using Applications from All Races and Controlling for Full Interactions

Dependent Variable:	Admission						
	OLS		Probit (average marginal effect)				
	(pp) (1)	(pp) (2)	(pp) (3)	(pp) (4)	(pp) (5)	(pp) (6)	(pp) (7)
A. Berkeley Difference-in-differences							
Black × Post-ban	-47.3	-40.6	-47.2	-46.0	-49.1	-47.7	-50.3
	(7.1)	(7.4)	(4.8)	(5.3)	(4.6)	(5.1)	(4.7)
National trend controls		x		x		x	
Inferred strength control					x	x	x
CA residency control							x
N (applications)	1,594	24,722	1,594	24,722	1,588	24,716	1,192
Clusters (applicants)	1,594	5,324	1,594	5,324	1,588	5,318	1,192
Actual pre-ban black admission rate	56.7	56.7	56.7	56.7	56.7	56.7	56.7
Δ implied by Black × Post-ban effect	-42.9	-36.9	-44.2	-42.9	-44.5	-43.0	-45.6
B. UCLA Difference-in-differences							
Black × Post-ban	-46.0	-38.0	-44.9	-41.8	-44.7	-41.3	-46.4
	(10.8)	(10.8)	(8.9)	(9.7)	(8.7)	(9.8)	(7.7)
National trend controls		x		x		x	
Inferred strength control					x	x	x
CA residency control							x
N (applications)	777	23,905	777	23,905	777	23,905	586
Clusters (applicants)	777	5,300	777	5,300	777	5,300	586
Actual pre-ban black admission rate	64.5	64.5	64.5	64.5	64.5	64.5	64.5
Δ implied by Black × Post-ban effect	-41.9	-34.6	-42.1	-38.9	-42.4	-38.8	-45.1

Notes - This table replicates Table II using applications from all races (black, white, Asian, and Hispanic) and more controls. The regressions underlying this table are the same as those underlying Table II except for two changes. First, the black indicator is replaced by a vector of black, Asian, and Hispanic indicators. Second, each non-racial covariate (LSAT, GPA, inferred strength, and California residency, depending on the specification) is interacted with each of the DD or DDD variables (the vector of race indicators, the post-ban indicator, the UC-school indicator, and any interactions of these variables). For example, column (1) regresses admission on LSAT, GPA, race indicators, year fixed effects, the race indicators interacted with the post-ban indicator, LSAT interacted with the post-ban indicator, GPA interacted with the post-ban indicator, LSAT interacted with the race indicators, and GPA interacted with the race indicators.

ONLINE APPENDIX TABLE V
Relationship between Admission and Race by School and Time Period

Dependent Variable:	Admission					
	Non-UC		UC Berkeley		UCLA	
	Pre-ban (pp)	Post-ban (pp)	Pre-ban (pp)	Post-ban (pp)	Pre-ban (pp)	Post-ban (pp)
	(1)	(2)	(3)	(4)	(5)	(6)
Black	64.2	56.4	77.4	31.9	64.7	19.0
	(2.0)	(2.0)	(5.5)	(4.8)	(7.6)	(7.9)
Hispanic	27.0	24.8	48.0	21.1	30.2	3.1
	(2.5)	(1.8)	(6.0)	(3.9)	(8.4)	(5.4)
Asian	4.1	-0.1	8.2	2.6	8.1	3.3
	(1.4)	(1.4)	(3.4)	(3.1)	(4.5)	(4.2)
LSAT (mean=0, sd=1)	22.8	25.2	24.2	17.5	28.1	28.9
	(0.7)	(0.6)	(1.7)	(1.4)	(2.2)	(2.0)
GPA (mean=0, sd=1)	23.4	19.9	22.3	21.6	20.2	19.3
	(0.8)	(0.9)	(1.9)	(1.6)	(2.4)	(2.1)
N (applications)	9,922	13,206	651	943	347	430
Clusters (applicants)	2,374	2,880	651	943	347	430
R-squared	0.444	0.450	0.441	0.363	0.497	0.525

Notes - This table reports coefficient estimates in percentage point units from OLS regressions of admission on race indicators, LSAT score, undergraduate GPA, and school-year fixed effects. The non-UC schools are the top-fifteen schools in the EALS that were never subject to an affirmative action ban. LSAT and GPA are each standardized across all EALS applicants to have mean zero and standard deviation one. In columns 1-2, I weight applications so that each school carries equal weight. Standard errors are clustered at the applicant level.

FIGURE I

Race, Academic Characteristics, and Admission under Affirmative Action

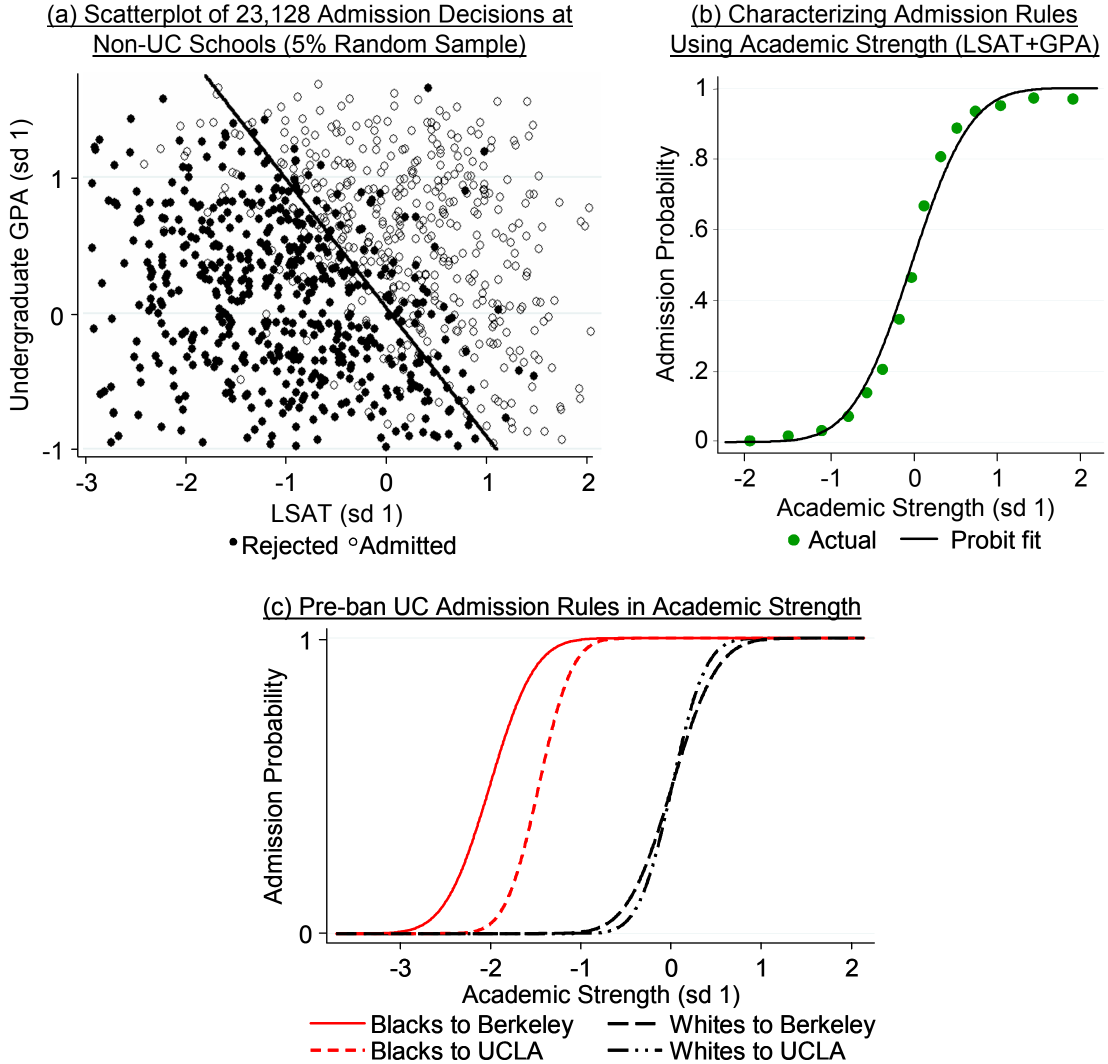

Notes – Figure Ia plots standardized LSAT score (mean zero and standard deviation one), standardized undergraduate GPA, and the actual admissions decision for a 5% random sample of the 23,128 Elite Applications to Law School (EALS) applications submitted to the top-fifteen non-UC schools that were never subject to an affirmative action ban. Online Appendix Figure II displays the full sample in color. To account for selectivity differences, each application's LSAT has been shifted by its school-year-race fixed effect from a probit regression of admission on LSAT, GPA, and these fixed effects (see Section II.D). The overlaid best-fit admission threshold line from the regression correctly predicts 85.4% of admissions decisions. The regression indicates that a one standard deviation higher LSAT is about as valuable in the admissions cross section as a one standard deviation higher GPA. Thus when useful, I summarize an application's LSAT and GPA with the scalar index "academic strength", equal to the standardized sum of standardized LSAT and standardized GPA. Figure Ib plots admission rates within fifteen academic strength bins using all 23,128 non-UC applications and overlays the univariate probit fit, where each application's academic strength has been shifted by its school-year-race fixed effect from a probit regression of admission on academic strength and these fixed effects. Figure Ic plots probit-fitted "admission rules" by race at UC schools before the 1996 affirmative action ban, derived from a regression of admission on academic strength, a black indicator, and year fixed effects using black and white pre-ban applicants to Berkeley, and separately for UCLA. For ease of comparison, each school's pair of admission rules has been shifted horizontally by an additive constant so that the predicted admission probability for whites equals 0.5 at academic strength 0.

FIGURE II
UC Admission Rates by Race Based on Public Aggregates
Not Holding Academic Strength Constant

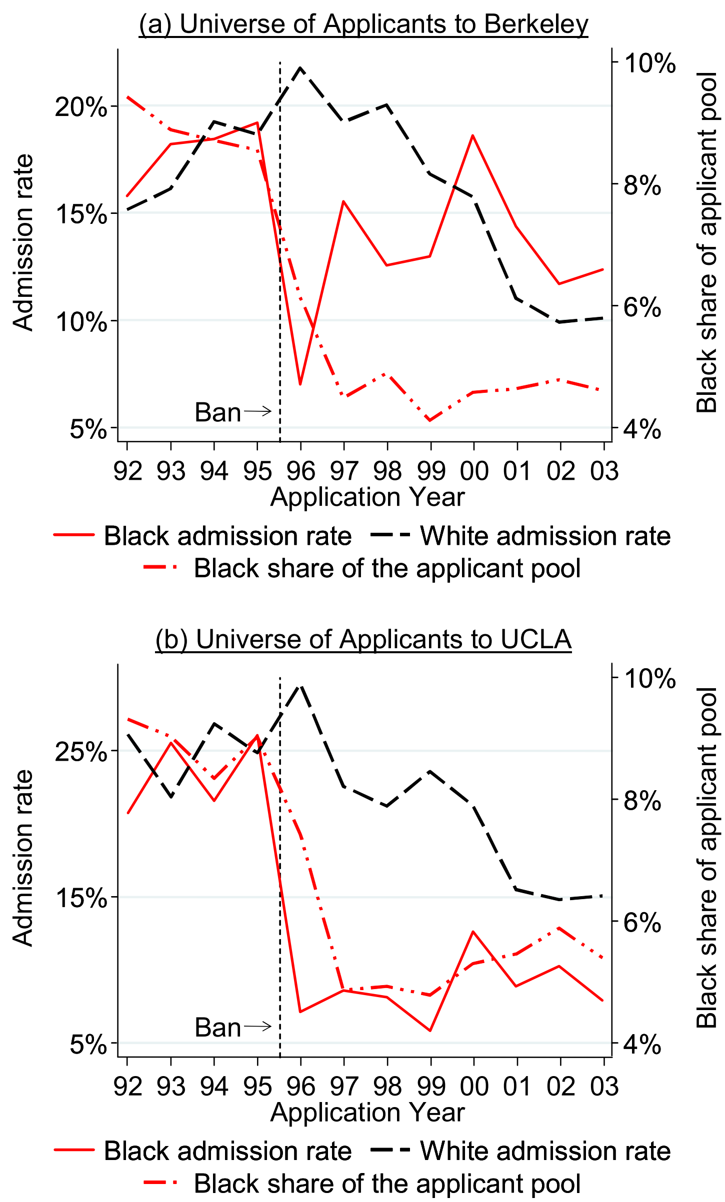

Notes – This graph uses public aggregates reported by the University of California on the universe of applicants (not just the EALS) to plot the time series of overall admission rates by race and the black share of the applicant pool at Berkeley and UCLA (which in this paper always refers to Berkeley and UCLA law schools). Application year refers to the autumn of the application year. These unconditional aggregates contain no information on applicant strength by race. The data were originally accessed on August 27, 2009, from the website of the UC Office of the President at http://www.ucop.edu/acadadv/datamgmt/graddata/lawnos.pdf; this file is no longer available online but it is available from the author.

FIGURE III
UC and Non-UC Admission Rates by Race in the EALS
Holding Academic Strength Constant

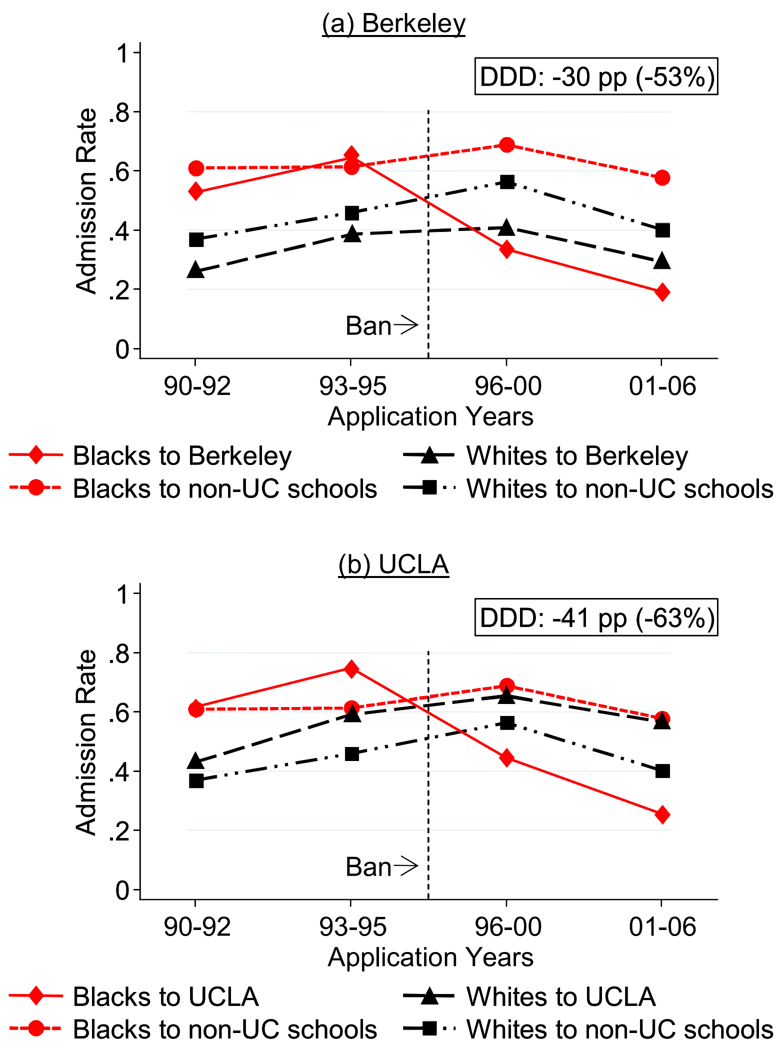

Notes – Figure IIIa displays the time series of black and white admission rates at Berkeley, UCLA, and non-UC schools, where applicant characteristics have been held constant at pre-ban levels using simple semi-parametric reweighting as in DiNardo, Fortin, and Lemieux (1996). To construct the time series of black admission rates at Berkeley, I first compute terciles of academic strength (the scalar summary measure of LSAT and GPA defined in Figure I) among pre-ban Black applications to Berkeley. Then for each time period shown in the figure, I weight black applications to Berkeley so that each pre-ban-defined tercile receives equal weight when computing the displayed admission rate. I repeat this process for whites at Berkeley and for whites and blacks separately at UCLA and at each non-UC school, averaging across non-UC schools to construct the non-UC series. This semi-parametric reweighting is data-demanding, so I group the data into two pre-ban time periods (1990-1992 and 1993-1995) and two post-ban time periods (1996-2000 and 2001-2006). Pooling all pre-ban years and all post-ban years, the difference-in-differences-in-differences estimate of the effect of the ban on the black admission rate at each UC school is overlaid, with the DDD estimate as a fraction of the pre-ban admission rate in parentheses. Online Appendix Table II lists the numbers underlying the DDD estimates. Table II reports parametric DDD estimates that account for the minor space-opening effect of a decline in black admission rates.

FIGURE IV
Black-White Differences in Post-ban Admissions

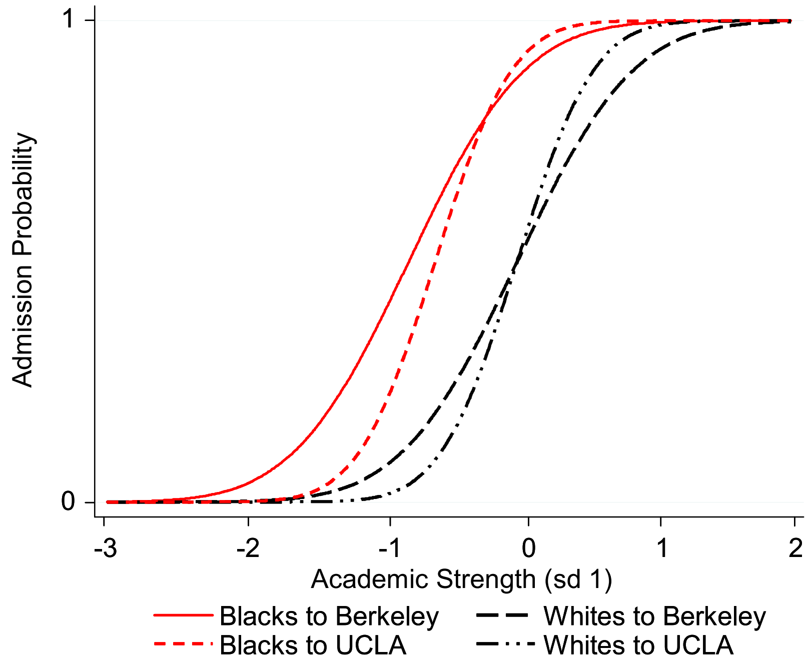

Notes – This figure replicates Figure Ic for post-ban applicants. It plots fitted admission rules by race at UC schools after the affirmative action ban, derived from a probit regression of admission on academic strength, a black indicator, and year fixed effects using black and white post-ban applicants to Berkeley, and separately for UCLA. See the notes to Figure Ic for the definition of academic strength. For ease of comparison, each school's pair of admission rules has been shifted horizontally by an additive constant so that the predicted admission probability for whites equals 0.5 at academic strength 0. The maximum vertical distance between the Berkeley curves is 56 percentage points and between the UCLA curves is 63 percentage points, slightly smaller than the estimates reported in Table III column (4) that condition more flexibly on covariates. (The horizontal distance between the Berkeley curves indicates that black status is observed to have been worth 0.86 standard deviations of academic strength in the post-ban cross section. For UCLA, the figure is 0.66 standard deviations.)

ONLINE APPENDIX FIGURE I
Distribution of Academic Characteristics By Race in the EALS

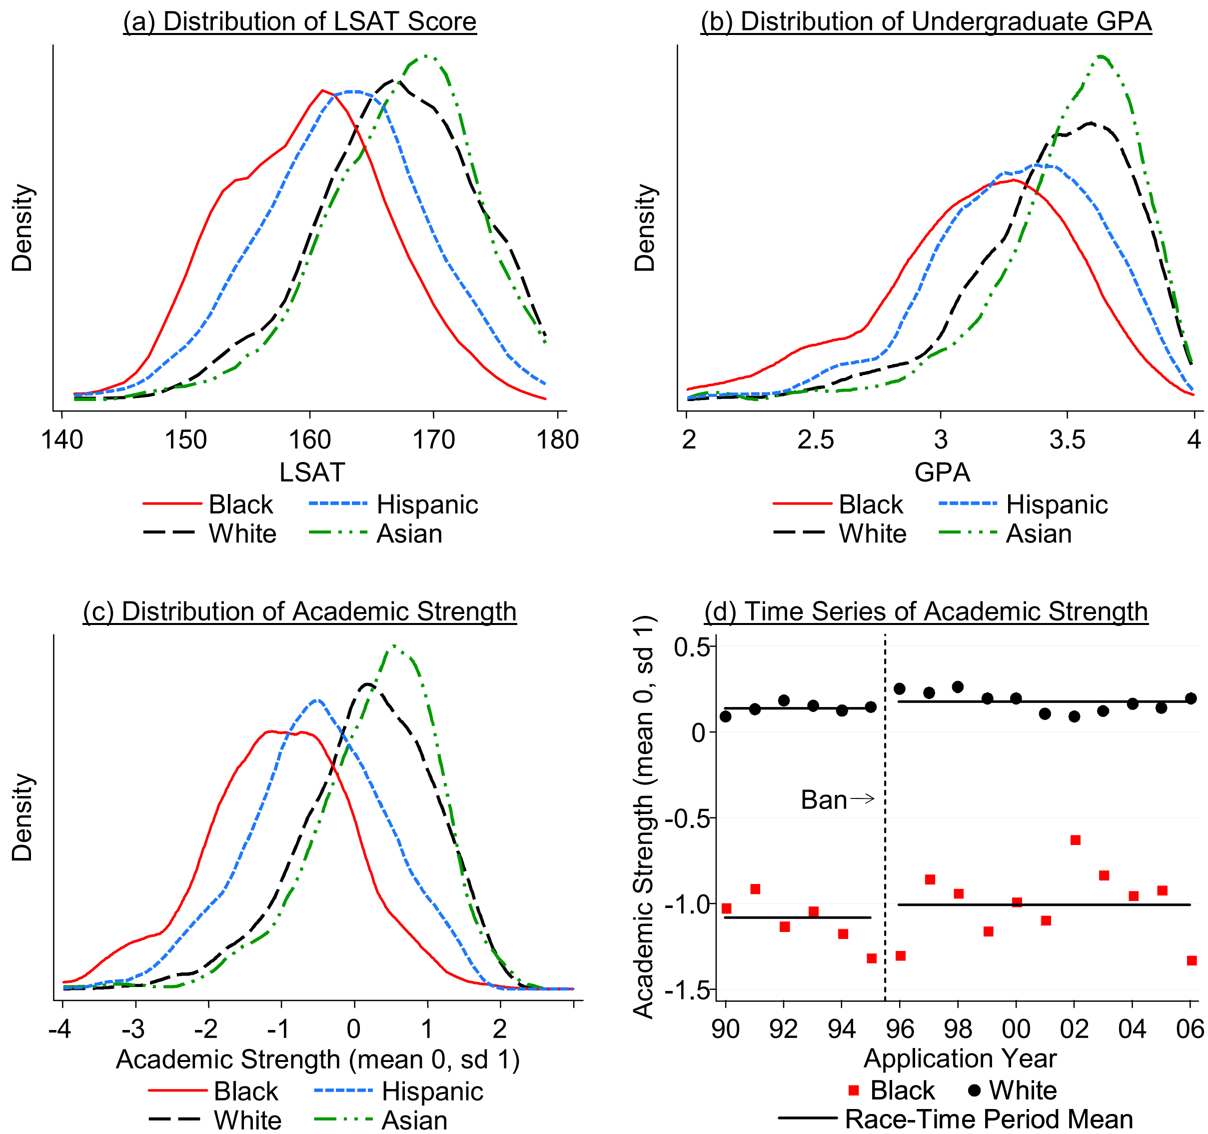

Notes – This figure displays the distribution of academic characteristics by race among Elite Application to Law School applicants in the paper's main sample: the 94% of EALS applicants who applied to Berkeley, UCLA, and/or one of the top-fifteen non-UC schools that were never subject to an affirmative action ban. LSAT is the standardized test score used in law school admissions and ranges from 120 to 180. Undergraduate grade point average is the cumulative undergraduate GPA on a 4.00 scale. "Academic strength" is a scalar index of the strength of an applicant's academic credentials, equal to the standardized (mean zero and standard deviation one) sum of standardized LSAT and standardized GPA (see Figure I for the semi-parametric motivation). Each displayed density is estimated non-parametrically using an Epanechnikov kernel with Silverman bandwidth. Application year refers to the autumn of the application year.

ONLINE APPENDIX FIGURE II
Scatterplot of 23,128 Admissions Decisions at Non-UC Schools

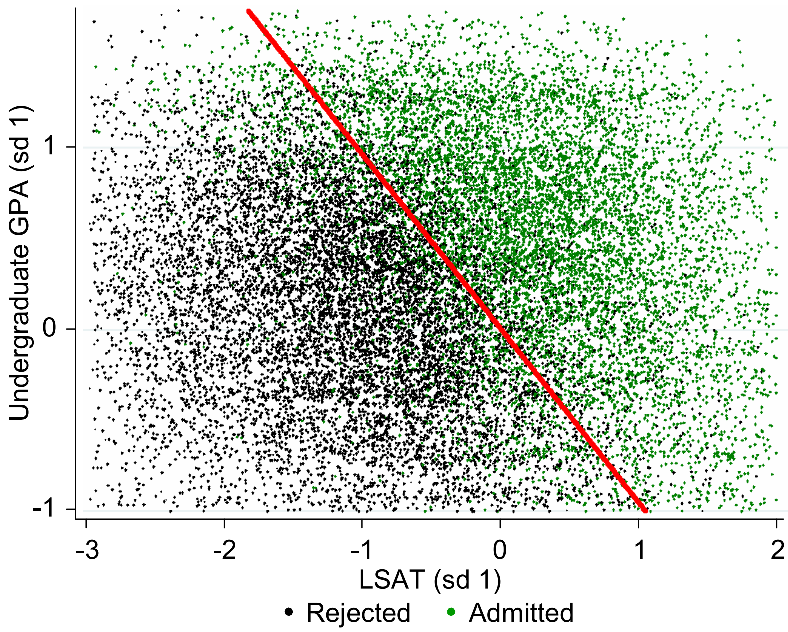

Notes – This figure is intelligible only in color. This figure replicates Figure Ia except that it plots all 23,128 applications to non-UC schools, rather than just a 5% random sample. See the notes to that figure for details.

ONLINE APPENDIX FIGURE III
Race and Admissions in the Cross Section under Affirmative Action

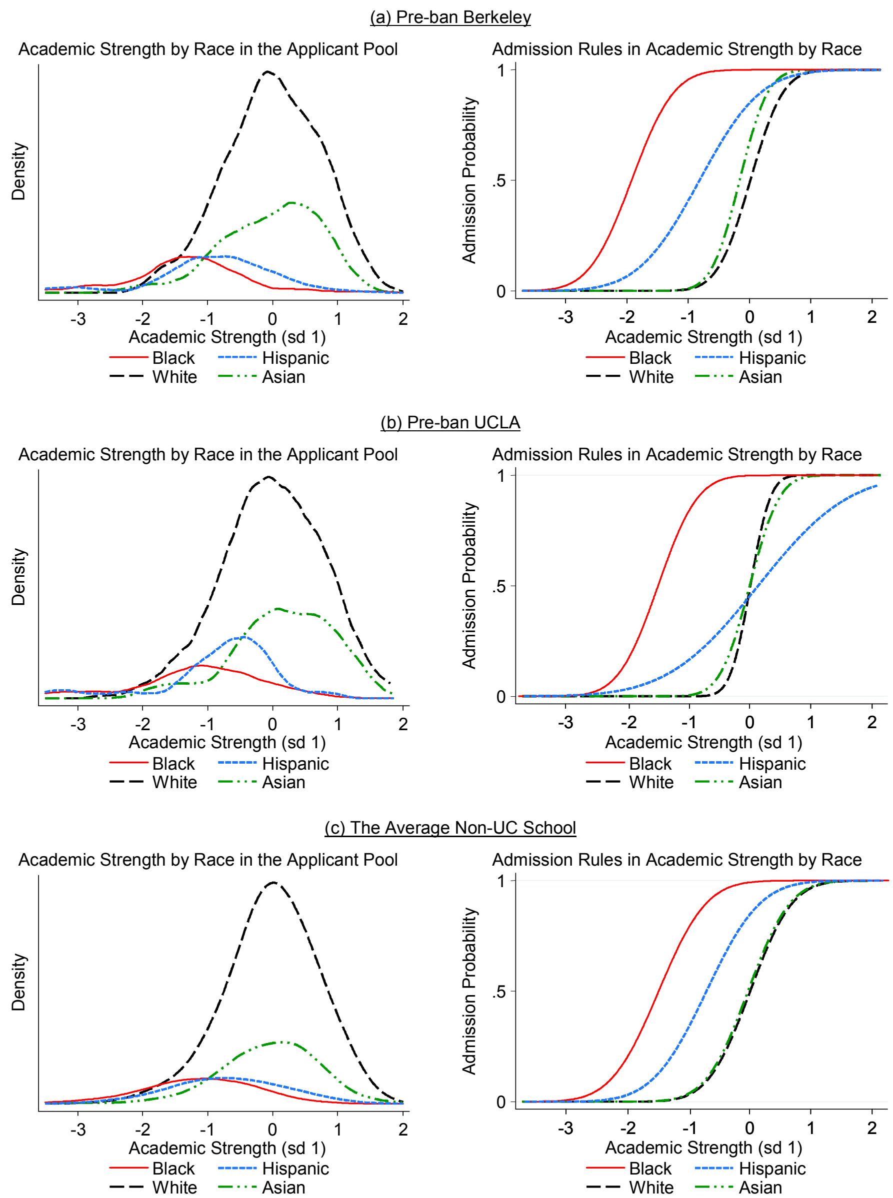

Notes – This figure plots academic strength by race in the applicant pool and the admission rules in academic strength (defined and motivated in Figure I) by race in the EALS at pre-ban Berkeley, pre-ban UCLA, and the average non-UC school in the average year. The left-hand-side panels display the density of applicants by academic strength and race. To construct these, I pool all years within each school-race, estimate each school-race's density non-parametrically using an Epanechnikov kernel with Silverman bandwidth, shift each school's distributions horizontally by an additive constant so that the white mode lies at academic strength 0, and then (for the set of non-UC schools only) average densities across schools. The right-hand-side panels display fitted admission rules by race constructed similarly to Figure Ic except that admission is allowed to respond to academic strength differently for each race. For each school type, the fits derive from a probit regression of admission on academic strength, race indicators, interactions among the race indicators and academic strength, and school-year fixed effects. The non-UC regression is weighted so that each school carries equal weight in each time period (pre-ban and post-ban), consistent with regressions elsewhere in the paper.

ONLINE APPENDIX FIGURE IV
UC and Non-UC Admission Rates by Race in the EALS
Not Holding Academic Strength Constant

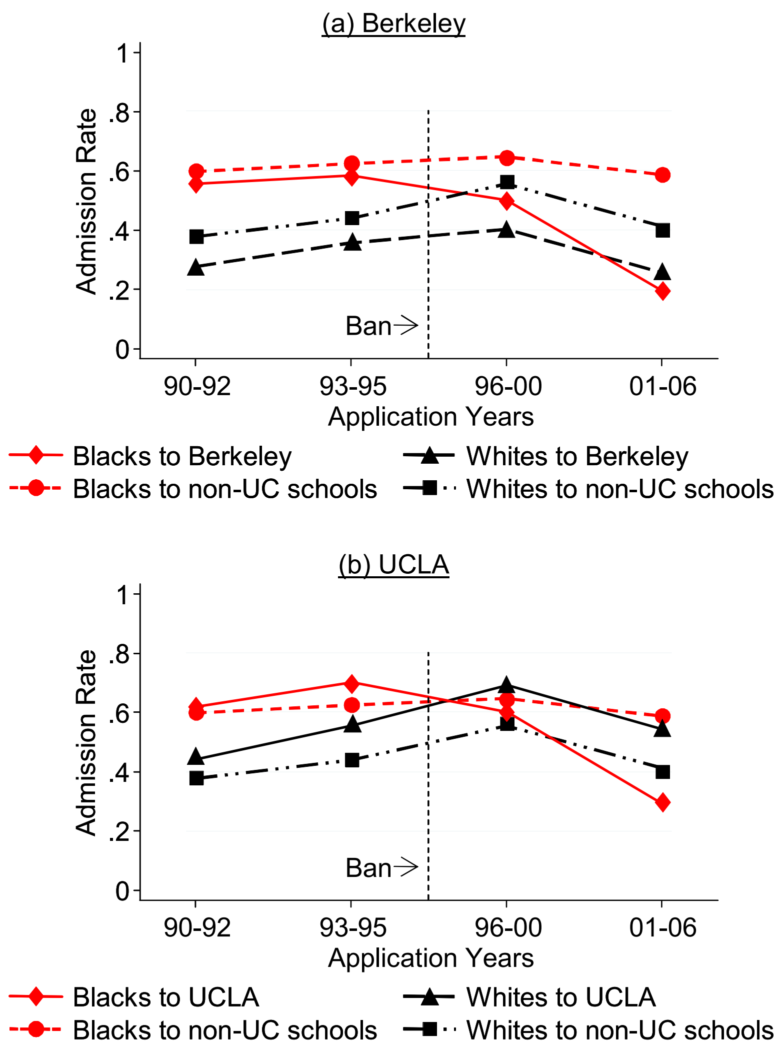

Notes – This figure replicates Figure III except that it does not reweight applications in any way. The change in admission rates is uneven due to differences in academic strength over time; the reweighting in Figure III adjusts for these differences. Although I find evidence that less-academically-credentialed black EALS applicants were less likely to apply to UC schools after the ban, this figure shows that unconditional black admission rates at UC schools still fell after the ban among EALS applicants more than they did among the universe of UC applicants (see Figure II). Graduates from the elite college that EALS students attended constitute a larger fraction of the Berkeley and UCLA applicant pools than graduates of most other individual colleges, so one possible explanation for the discrepancy between this figure and Figure II is that less-academically-credentialed black EALS applicants learned from the admissions outcomes of their predecessors that their admission prospects had in fact not fallen to zero and thus applied in greater numbers as time went on, causing unweighted black admission rates to fall. This explanation is consistent with the time series in this figure: black admission rates were higher in the first several years after the ban than in more recent years. As explained in the Section III.A, the EALS provides strong statistical power on admission decisions but not application decisions, so the EALS is not well-suited to formal testing of application behavior.

ONLINE APPENDIX FIGURE V
Relationship between Admission and the Inferred Strength Variable

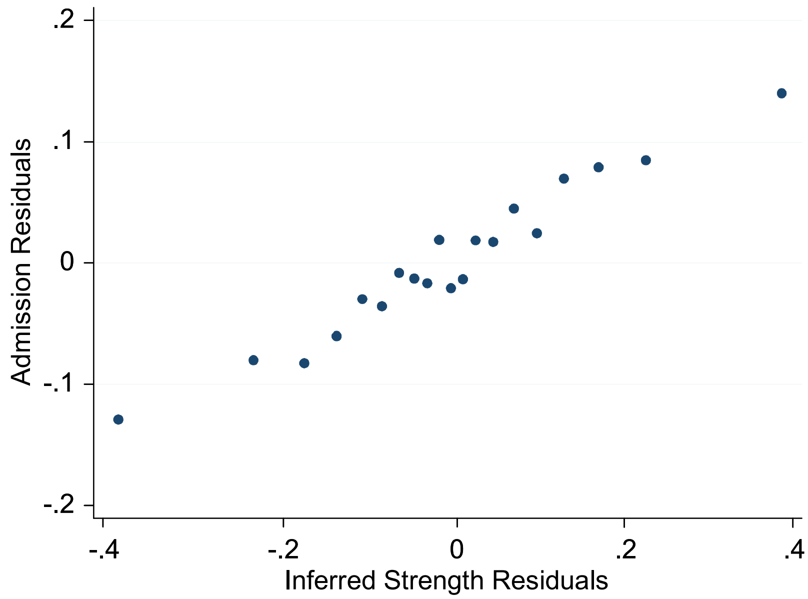

Notes – This graph displays the incremental predictive power of the inferred strength variable (motivated by Dale and Krueger 2002), conditional on other covariates. The difference-in-differences-in-differences regressions underlying Table II column (4) do not control for admission factors that are omitted from the EALS such as recommendation letter strength. I proxy for such commonly-valued unobserved admission determinants using the intuition that if an applicant predicted to be rejected based on LSAT, GPA, and race is in fact consistently admitted across schools in the EALS, this applicant is likely strong on unobserved characteristics like recommendation letters. Specifically, I construct an "inferred strength" variable for an application submitted by applicant *i* to school *s* equal to the mean across all applications submitted by applicant *í* to schools other than *s* of residuals from within-school regressions of admission on LSAT, GPA, race indicators, and time-period fixed effects. To construct this graph, I compute residuals from an OLS regression of inferred strength on all the covariates used in Table IIa column (4), group observations into twenty equal-sized (5 percentile-point) bins based on inferred strength residuals, and plot means within each bin. The few applicants who applied to only one school are omitted from this graph and are handled flexibly in the main regressions as specified in Section III.C.

ONLINE APPENDIX FIGURE VI
Personal Statement Prompts on Berkeley Application Forms

(a) 1995, the last pre-ban year

Personal Statement

The required personal statement is your opportunity to submit information you would like the reader of your file to have when evaluating your application. A résumé or chronological list of activities may accompany the personal statement.

The content of the statement is up to you; we would appreciate your limiting its length to about two typewritten pages. **Be sure to include your social security number and sign your statement.** You may wish to address the admissions criteria described in the paragraphs above as well as the following:

• Academic honors, awards or other recognition you have received not based solely on grade average;

• Influences on your life such as cultural, ethnic, or racial background or your circumstances while growing up;

• Physical or learning disabilities and any effect they may have had on your admissions credentials. Special resources and services are available to students with disabilities. Please contact the Disabled Students' Program, (510) 642-0518, for further information. (See page 45, Applicants with Disabilities.)

• Extracurricular, community, or other activities in the order of their importance to you;

• Work experience including nature and amount of outside employment while in college;

• For college graduates, the nature of your activities or employment since graduation;

• Description of your graduate studies, if any;

• Other test scores you would like to report (SAT, GRE, etc.);

• College grading and course selection, and college grade trends;

• Any additional information that would indicate that your admission would add diversity to the entering class.

(b) 1996, the first post-ban year

Personal Statement

Please submit a statement providing more information about yourself. The subject matter of the essay is up to you, but keep in mind that the reader will be seeking to get a sense of you as a person and as a potential student and graduate of Boalt Hall.

As stated in the Faculty Policy Governing Admission, Boalt Hall seeks to enroll a class with varied backgrounds and interests. Such diversity contributes to the learning environment of the law school and historically has produced graduates who have served all segments of society and who have become leaders in many fields of law. If you wish, you may separately discuss how your interests, background, life experiences and perspectives would contribute to the diversity of the entering class. If applicable, you may also describe any disadvantages that may have adversely affected your past performance, including such disadvantages as language barriers, or a personal or family history of educational or socioeconomic disadvantage.

Notes – This figure reprints the personal statement prompts from the 1995 (the last pre-ban year) and 1996 (the first post-ban year) Berkeley (formerly called "Boalt Hall") application forms. The 1995 personal statement prompt was nearly identical to the one from 1994. The personal statement prompts have remained almost exactly unchanged since 1996. Not depicted here, UCLA also changed its admissions process, giving explicit preference to students who indicated interest in critical race studies.

ONLINE APPENDIX FIGURE VII
Behavior under an Affirmative Action Ban

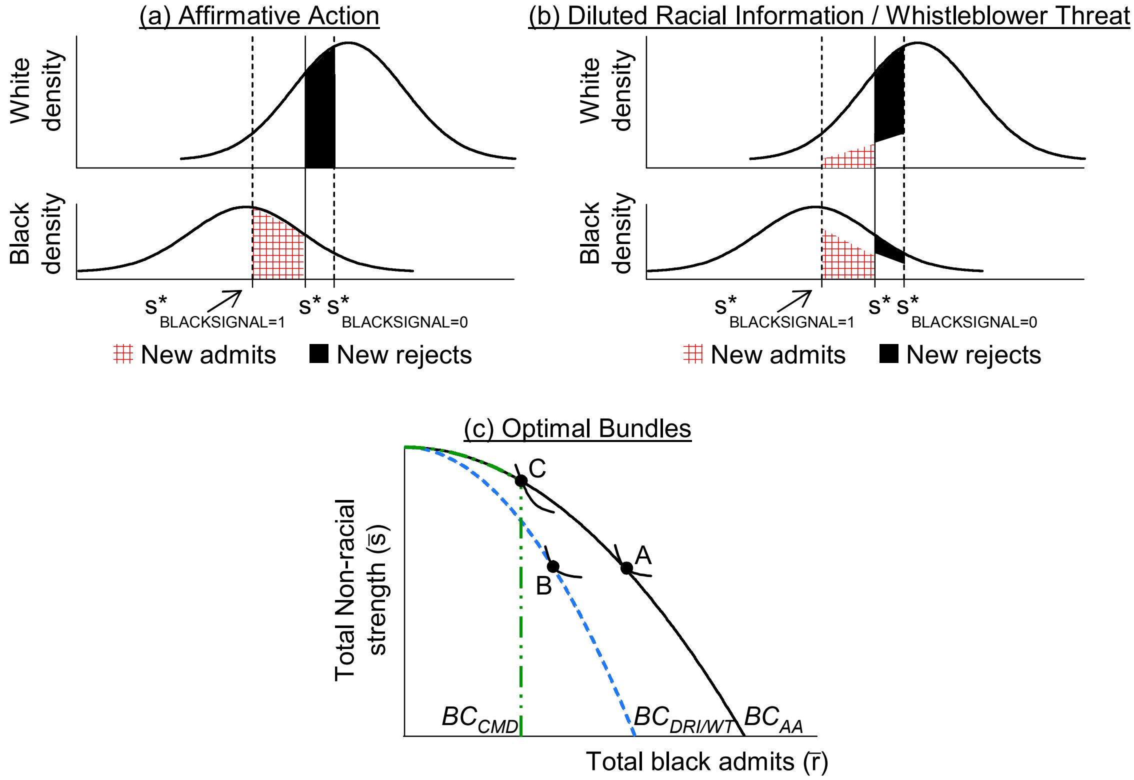

Notes – This figure illustrates the simple model detailed in Online Appendix B in which the applicant pool is held fixed and the admission office has concave preferences over the number of black applicants admitted and the aggregate non-racial strength of the admitted cohort. The admission office can admit applicants on two pieces of applicant information: non-racial strength and a signal of black status. Panels (a) and (b) depict applicant densities in non-racial strength; "new" refers to the effects of placing positive weight on the black signal. Panel (c) plots budget sets under the simplification of uniform distributions of non-racial strength; the graph omits feasible but always-dominated bundles by defining the x-intercept as the number of black applicants admitted if the admission office were to maximize only non-racial strength and the y-intercept as the aggregate non-racial strength achieved if the admission office were to maximize only the number of admitted blacks. Under affirmative action ("AA"), the black signal is pure. If a ban either dilutes the racial information available to the admission office ("DRI") or forces an admission office to abstain from using its pure racial information because an insider may expose the evasion (whistleblower threat "WT"), the admission office can use only an imperfect signal of black status. This increases the non-racial strength it must forego to admit each additional black applicant and pushes the DRI/WT budget side inside the AA budget set. If instead a ban places a cap on measurable black-white disparities ("CMD") without diluting the usable black signal, the post-ban admission office can use its pure black signal to achieve any bundle in its AA budget set, so long as the number of admitted blacks does not exceed the *de facto* limit.