

Affirmative Action and the Quality-Fit Tradeoff*

Peter Arcidiacono

Michael Lovenheim

Duke University & NBER Cornell University & NBER

January 15, 2015

1 Introduction

Race-based affirmative action policies, which seek to provide “affirmative” help to racial minorities, are among the most contentious policies that exist in US labor and education markets. Proponents argue that affirmative action helps level the playing field between majority and minority workers or students. Indeed, the original motivation for affirmative action is the legacy of discrimination in the U.S., particularly against African Americans. Lyndon Johnson summed it up well in his 1965 commencement address at Howard University, where he states

“You do not take a person who, for years, has been hobbled by chains and liberate him, bring him up to the starting line of a race and then say, ‘You are free to compete with all the others,’ and still justly believe that you have been completely fair.”

In this paper we focus on how affirmative action in higher education—and in particular how racial preferences in admissions—affect the outcomes of minority students.

While our focus is on higher education, affirmative action began in the United States with Lyndon Johnson’s 1965 executive order that affected the hiring practices of all federal contractors.

*We thank Kate Antonovics, Jivesh D’Sousa, Peter Hinrichs, and Richard Sander for helpful comments.

These regulations have since expanded to many states, and a large number of private employers have voluntarily enacted affirmative action hiring practices. Despite the fact that the focus of Johnson's original executive order was on hiring practices of federal contractors, colleges and universities throughout America quickly adopted these practices as well. The decision to engage in affirmative action in admissions is not federally-mandated; instead, both the decision to enact race-based affirmative action policies and the manner in which minority candidates are given admission preference are decisions made by each postsecondary institution. As a result of the rapid spread of affirmative action in to the higher education sector, race-based admissions preferences are one of the most prevalent and contentious fixtures of American postsecondary education.

As in the labor market, race-based preferences in college admissions evoke fierce and divided opinions on their merits, divisions that also hold among the Supreme Court justices. Of the three landmark court decisions on affirmative action that examined the use of applicant race in the admissions process at public universities, two were 5-4 decisions. While the use of explicit quota systems and point systems for minority applicants were ruled out in *Bakke v. California Board of Regents* and *Gratz v. Bollinger* respectively, *Grutter v. Bollinger* permitted the consideration of race as one of many factors in university admissions.¹ Writing for the majority that Michigan's law school admission process was constitutional, Justice O'Connor states that the constitution "does not prohibit the law school's narrowly tailored use of race in admissions decisions to further a compelling interest in obtaining the educational benefits that flow from a diverse student body."² More recently, the Supreme Court ruled in *Schuette v. Coalition to Defend Affirmative Action* that a state ban on universities considering race as part of their admissions process does not violate the Equal Protection Clause of the US Constitution. Thus, while universities can use race as a factor in admissions as long as it is "narrowly tailored," states can prohibit race-based admissions as well, at least among

¹In *Fisher v. Texas*, the Supreme Court ruled 7-1 to remand the case back to the lower court, stating that the lower court had not applied strict scrutiny.

²*Grutter v. Bollinger*, 539 U.S. 306 (2003).

public universities.

Whether many university admissions processes are narrowly tailored in their consideration of race – particularly at law schools – is debatable. The extent of race-based preferences varies by ethnic group, with affirmative action appearing to be largest for African Americans relative to Hispanics and other under-represented groups.³ For example, the average SAT score (Math+Verbal) for African Americans entering Duke University in 2001 or 2002 was 140 points lower than the average for white students, or about 1.4 standard deviations below the white mean. For Hispanics, the gap was half that (Arcidiacono et al. 2011).⁴

While *Grutter v. Bollinger* established that the legality of affirmative action policies is rooted in institutions' preferences for having a more diverse student body,⁵ President Johnson's quote and much of the rhetoric surrounding affirmative action focuses on the desire to support the educational attainment of under-represented minority students. Certainly, the initial impetus for these policies was the hope that affirmative action could play some role in reducing inequalities in life outcomes between minority and majority groups.

While racial preferences in admissions can be argued for based on the basis of remedying past discrimination or other equity-based reasons, there may be efficiency arguments in its support as well. Durlauf (2008) states that it is not obvious *a priori* that racial preferences are inefficient for two reasons. First, he highlights the fact that racial preferences will be inefficient when there

³There is a myth that the primary beneficiaries of affirmative action in admissions are women. This is true only to the extent that women are members of particular racial or ethnic groups. There is now a substantial gender gap in educational attainment between men and women. The gap is particularly large for African Americans, where there are two African American women enrolled in a four-year college for every African American man. The gender gap has become so substantial that men are now more likely to receive preferential treatment in admissions than women (Green et al. 2011). Preferences do, however, exist for athletes and legacies. For example, Espenshade, Chung, and Walling (2004) find preferences for athletes to be smaller than that for African Americans but similar to that for Hispanics, with lower (but still significant) preferences for legacies.

⁴Because the distribution of SAT scores for African Americans lies to the left of that of whites, even without race-based preferences African Americans would have lower SAT scores than whites within a particular school because they would more likely be on the margin of being admitted. But, it is certain that the gap would be much smaller absent race-based admission preferences. Arcidiacono, Khan, and Vigdor (2011) show that legacies have similar SAT scores to non-legacies. This is actually indicative of legacy preferences following similar arguments: legacies likely come from advantaged backgrounds so their test score distribution should be to the right of that of non-legacies.

⁵There is a small literature on the educational and social benefits to majority students of diversity that comes to mixed conclusions. This evidence is summarized in Arcidiacono, Lovenheim and Zhu (2014).

are complementarities between school quality and student preparation. Second, even if aggregate human capital would be higher without racial preferences, such admission preferences could increase the allocative efficiency of human capital across communities. That is, the social returns to raising human capital among individuals from heavily minority communities may be particularly high. In such a case, sacrificing productive efficiency in terms of the aggregate amount of human capital for allocative efficiency in terms of the distribution of human capital across different groups in the economy may be efficiency enhancing. While we are able to make progress on the first issue, the second is beyond the scope of this paper and provides an important caveat to our findings.

In this paper, we critically examine the evidence on how affirmative action in undergraduate and law school admissions affects both education and labor market outcomes.⁶ Although the effect of affirmative action on minority student outcomes is somewhat unrelated to the success of these policies in meeting institutions' preferences for ethnic diversity, these are important questions to consider because they highlight the potential costs or alternative benefits to these diversity goals. If, indeed, affirmative action harms long-run outcomes of minority students, it suggests that returns to diversity within an institution would need to be quite large in order to justify these programs. Conversely, if affirmative action acts to level the playing field, as originally intended, these policies may be particularly desirable because they reduce inequality and allow schools to act on their preferences for diversity.

Revealed preference suggests that more-selective and higher-resource colleges are preferred by students to their less-selective counterparts.⁷ One then would suspect that more-selective colleges also would lead to better student outcomes. In addition, resources such as expenditures per student are higher at selective colleges, and there is ample evidence that these expenditures translate into,

⁶We focus on undergraduate and law school admissions due to data availability and corresponding focus of the literature. There are some papers in the medical literature examining the importance of matching between the race of the doctor and the race of the patient, but these papers suffer from concerns about identification. See Durlauf (2008) for a discussion of this work.

⁷Long (2004) also shows evidence from conditional logit choice models that students place significant and increasing value on college quality when making admissions decisions.

for example, higher graduation rates. Thus, policies that sort minority students into more selective schools, which is what affirmative action aims to do, may lead to better educational outcomes for these students. But, is increasing college quality always beneficial? As a result of affirmative action, under-represented minorities, and in particular African Americans, are placed in schools where their academic preparation is significantly below that of their white counterparts. The *mismatch hypothesis* argues that many of the beneficiaries of preferences are so misplaced academically that they would actually be better off in the absence of affirmative action. Although its antecedents arose much earlier, this hypothesis was popularized in academia by Sander (2004) and in the mainstream by Sander and Taylor's 2012 book *Mismatch: How Affirmative Action Hurts Students It's Intended to Help, and Why Universities Won't Admit It*. While Sander's original article focuses on law school students, his work with Taylor examines the evidence for the mismatch hypothesis for undergraduates as well.

At first blush, economists should be very skeptical of the mismatch hypothesis. After all, affirmative action just expands the choice set. If an individual believes a particular school would be too difficult, then the individual could simply attend another school. By revealed preference, the individual must be better off. Yet, even in the context of rational expectations, where student beliefs are right on average given their information sets, there are cases where under-represented groups could be made better off in the absence of affirmative action. This comes about because schools may have private information about match quality but are letting in lower match quality students in order to satisfy diversity goals. Thus, the information sets of both schools and students are the critical components in driving mismatch.

Ultimately, the effect of affirmative action on student outcomes is an empirical question, and we evaluate the evidence to date on whether and how affirmative action in admissions helps its

purported beneficiaries.⁸ Our first goal in what follows is to address the following questions:

1. Under what conditions can affirmative action actually be bad for its beneficiaries? What are the relevant outcome measures in this case?
2. How can we make the empirical specification flexible enough to allow for the possibility that increasing school quality could be welfare-enhancing for some individuals but not for others?

While the second question may seem obvious given the first, it is surprising that many papers – even when explicitly testing for mismatch – estimate models that imply attending a more selective program is good for everyone or good for no one.

Given the answers to the first two sets of questions, we then turn to what the empirical evidence has to say about affirmative action more broadly and the quality-fit tradeoff in particular. Focusing on both undergraduates and law school students, we discuss the research findings on the following questions:

3. How extensive are race-based preferences?
4. What can we learn from the debate over mismatch in law schools? Does it matter for law school grades, bar passage, future earnings, or some combination thereof?
5. For undergraduates, on what margins does affirmative action matter? The extensive margin (whether individuals attend college at all) or the intensive margin (where individuals enroll)?
6. How does affirmative action affect graduation rates, choice of college major, and earnings?
7. Are “percent plans,” which provide automatic admission to public universities for students above a given percentile in class or state rank, a viable alternative to affirmative action? What effects do such admissions rules have on student outcomes and how do they compare to the estimated effects of affirmative action on those same outcomes?

⁸See Holzer and Neumark (2000) for a review of the effects of affirmative policies both in education and the labor market.

As we will see, there are many ways that the quality-fit tradeoff can manifest itself, but it will be difficult, if not impossible, to say whether the purported beneficiaries are actually made worse off under affirmative action. At the same time, there is consistent evidence that the fit between the student and the university matters, at least across some dimensions. We highlight areas for future research that can help shed some much needed light on these important and unresolved questions.

The rest of the paper proceeds as follows. In section 2 we discuss the mismatch hypothesis and the conditions for it to hold given that affirmative action expands the choice sets under-represented groups face. Important to this discussion will be the quality of information individuals have about their prospects for success. In section 3 we discuss ways to empirically test for heterogeneity in the fit between the school and the student. Section 4 examines mismatch in law school, considering the initial article by Sander as well as the multiple responses it has generated in light of the issues raised in sections 2 and 3. Section 5 focuses on mismatch at the undergraduate level, including the effects of statewide affirmative action bans. Section 6 discusses percent plans, and section 7 concludes.

2 The Quality-Fit Tradeoff

The scope for affirmative action policies to benefit minority students hinges in part on the extent to which more-selective colleges produce better outcomes for all students than less-selective colleges. This could be the case because the more selective colleges provide higher monetary inputs and have higher quality students who in turn generate positive peer effects. We refer to these as “quality effects.”

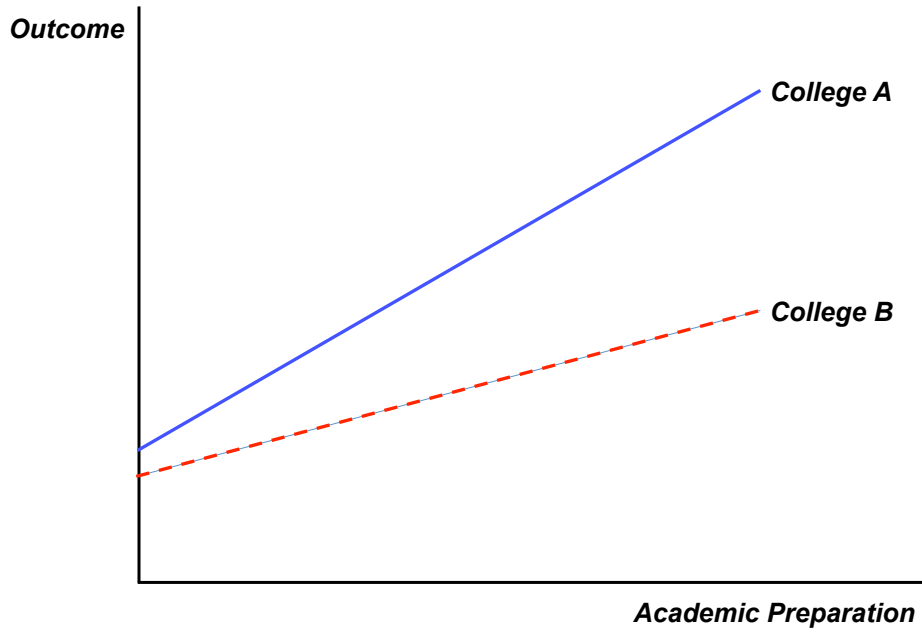
But the match between the college and the student also could be important. More selective colleges may pitch their material at a higher level, move at a faster pace, and assume that certain concepts and skills already are known. It then may be optimal for some students to attend a less-selective college even if a more-selective college is in their choice set. In other words, there may be

an optimal match between the pre-collegiate skills of students and schools of different types that varies across the student skill distribution. We refer to these as “match effects.”

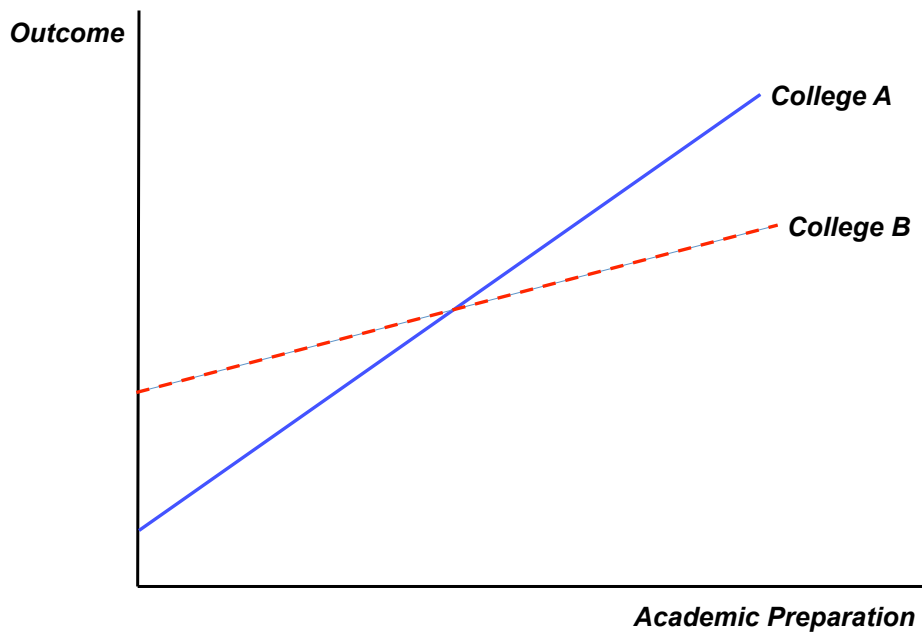
The extent to which match and quality effects are important may depend on how the benefits to the student are measured. For example, science courses may build on past knowledge in ways that other courses do not. Hence, match effects may be more relevant when considering graduating in the sciences rather than graduating in the humanities. Whether match or quality effects dominate may also depend on what part of the college quality distribution is being examined. For example, among selective colleges match effects may dominate. However, the significantly lower resource levels that prevail among non-selective colleges may lead quality effects to dominate in this part of the school quality distribution. Assessing the importance of the fit between the student and the school thus requires flexibility in how the outcome equation is specified. Namely, the returns to academic background need to vary in some way with the quality of the college.

Figure 1 shows one way in which heterogeneity in fit could be modeled. Panel (a) shows a case where College A produces better outcomes for all individuals. In contrast, Panel (b) shows a case where the match effects trump quality effects for a part of the skill distribution: higher levels of academic background are associated with better outcomes at both colleges, but now the college that maximizes the student’s expected outcome depends on the student’s academic background. College A is a particularly good fit for someone with a strong academic background. College B, however, will not be as stimulating for those who are more academically prepared for college, as College B is geared more towards those with relatively weaker backgrounds. In this example, College B will result in better outcomes for those who are less academically prepared for college. As we discuss below, the prior literature on school quality often estimates an average effect that assumes the relationship between quality and outcomes is monotonic for all students. This assumption usually is necessitated by data and sample size limitations, but with enough data we could see whether the

Figure 1: Heterogeneity in College Benefits



(a) College A dominates College B for all levels of academic preparation



(b) College A dominates College B for the most academically prepared, College B dominates College A for the least academically prepared

returns to academic background vary across different schools types.

3 The Role of Information

When students know their academic backgrounds and how their academic backgrounds translate into outcomes at different colleges, then students can optimally sort themselves among the colleges to which they have been admitted. Since racial preferences expand the choice set for under-represented minorities, when students have all the necessary information – or at least as much information as those of the colleges themselves – there would seem to be little scope for racial preferences to be harmful for minorities.⁹

As shown in Arcidiacono et al. (2011), when students are lacking information it is possible for minority students to be made worse off as a result of racial preferences, even when these students have rational expectations. The idea hinges on students not having full information regarding their abilities to succeed in school or about the quality of their matches with particular schools. If the university has private information about the quality of the match but does not reveal this information to the student, then it is possible that students admitted under affirmative action would have been better off attending a different school. This could be the case if schools admitted all students who were a good match as well as some minority students for whom the match quality is lower. Schools may have incentives to admit this latter group anyway in order to meet diversity goals. Suppose this latter group of students is the group that is admitted under affirmative action. Admitted minority students form expectations of their match quality based in part on observed average outcomes for all minority students admitted to the school and then make enrollment decisions based on the school that maximizes their expected utility. However, had the school revealed its private information, the student may have made a different decision.

Arcidiacono et al. (2011) distinguish between *local* mismatch and *global* mismatch. Under

⁹One place where negative impacts could arise would be through stereotyping which may negatively affect minority students who would have been admitted without racial preferences. But, those admitted under the preferences would still be advantaged as they would be pooled with stronger minority students.

local mismatch, the marginal admitted minority student is worse off under racial preferences but the average benefit for the group is still positive. Under global mismatch, the average benefit for the group turns negative. In their model, as long as schools face binding capacity constraints, some racial preferences will be advantageous for minority students. The model permits an interior optimal amount of racial preferences from the perspective of the minority group. Going past this point results in local mismatch, with further preferences potentially resulting in global mismatch.

That schools have private information about the quality of the match with the student is a necessary condition for mismatch when students have on average correct information. Information issues become even more salient when the rational expectations assumption is relaxed. Indeed, there is evidence suggesting that students – at least those from disadvantaged backgrounds – have surprisingly poor information about their educational prospects. Hoxby and Avery (2013) show that high-achieving, low-income high school students apply to schools that are very different from their high income counterparts, often not applying to a single selective institution despite the expected net costs of at least a set of these institutions being lower than the costs of the non-selective institutions they attend. Hoxby and Turner (2013) show that simply providing these high-achieving, low-income students information about their probabilities of admission to different tiers of schools as well as information about expected costs has significant effects on the types of colleges and universities these students attended. These students often have no desire to go to the local school but simply are unaware of the available options. This research demonstrates that considerable search frictions exist among students (particularly low-income students) in higher education markets. Although their focus is on students who match to lower-quality schools than they are able to attend, it demonstrates that students often do not have sufficient information to make informed decisions about where to attend college.

Racial preferences add another layer of uncertainty for students, because affected minority stu-

dents may not be aware they are being admitted with worse academic credentials than their peers as a result of these preferences. Universities generally state that race is one of many factors that are considered, so students may face significant uncertainty about the degree to which their race affected whether or not they were admitted. For example, Wightman (1998, page 70) shows that African American and white students had similar beliefs regarding their law school class rank.¹⁰ But, as we will see in Section 4, African American students have significantly lower grades in law school than their white counterparts.

Measuring whether affirmative action is harmful to its beneficiaries in a utility sense is a difficult task. Namely, it must be the case that by providing additional information about the prospects of success at a particular school, students would make different enrollment decisions. Interventions like the one used by Hoxby and Turner (2013) are geared towards the very top of the distribution. There are no interventions of which we are aware that give under-represented groups information about their prospects of success at different tiers of institutions.

In the absence of such interventions, researchers are forced to use correlates of match quality to estimate the extent of mismatch. If students have poor information about their match quality, once enrolled they may learn about the quality of the match and make adjustments if they determine that they enrolled in the wrong school. Indeed, using repeated surveys of students at Berea College that is designed to elicit beliefs about academic ability, Stinebrickner and Stinebrickner (2012) show that most low-income students¹¹ are over-optimistic about their grades when they enter college and quickly adjust their beliefs when they receive new information in the form of grades. Zafar (2011) finds similar learning dynamics among students at Northwestern University: students tend to be overly-optimistic about their major-specific outcomes at first and quickly update their beliefs once

¹⁰38% of African American and 34% of white law students expected to be in the top 10% of their class. The similar percentages for finishing in the top 25% of the class were 77% and 79% for African Americans and whites, respectively. The survey was conducted in the fall of 1991, and the sample was comprised of law students in their first semester of law school.

¹¹Berea College is a small liberal arts school in Kentucky that contains a student body comprised almost solely of students from low-income backgrounds.

they begin to receive grades.

The main adjustments students can make once they learn about a poor match is to switch to a less demanding major, transfer schools, or drop out. Stinebrickner and Stinebrickner (2012) show that learning about one's ability in the first year of college is strongly related to the decision to drop out, while Zafar (2011) finds that learning is associated with switching majors. A central reason for the differences in their findings is that dropout rates are extremely low at Northwestern (7%), while they are much larger at Berea (17%). Thus, these students adjust along different margins, which highlights the fact that any mismatch from affirmative action likely generates different responses depending on the characteristics of the school. This prediction has received little attention in the literature. In Section 5.4, we examine the evidence on college majors, while we discuss the relationship between affirmative action and college completion in Section 5.3. The effect of affirmative action on student transferring has not been examined in prior work, and we view this as an important area for future research.¹²

Although examining the relationship between affirmative action and these outcomes is of independent interest, it is important to highlight that these measures will be insufficient for establishing where minorities are actually worse off from affirmative action. For example, suppose we consider college graduation and could show that students admitted to a very selective school because of affirmative action would be more likely to graduate from a lower-tier school. The student may still have preferred to attend the more selective school even with full knowledge of the lower probability of graduating. This can occur because the individual may receive additional utility from graduating from the more selective school. In discussing the evidence on these outcomes, in addition to the varied identification concerns we highlight, it is important to keep in mind that that the outcomes

¹²Andrews, Li and Lovenheim (2014) provide a descriptive analysis of transferring behavior in Texas and show transferring prevalence is very high overall, and especially so at low levels of the college quality distribution. These results suggest that transferring is a margin on which students might adjust to learning about match quality, but no research currently exists that assesses whether this is the case.

being measured may not be capture all of the relevant aspects of how schooling choices affect utility, particularly in the long run.

4 The Law School Debate

The discussion over whether affirmative action benefits minority students has been particularly contentious in the realm of law school admissions. Estimating the effect of affirmative action in law schools differs from examining this question in the context of undergraduate admissions for several reasons. First, in contrast to what is seen for undergraduate institutions, affirmative action affects how many minority students attend law school as virtually all laws schools are selective. Hence, for law schools, it is not a matter of just where students attend but whether they attend at all. Second, the data on law schools are not as rich as the data for undergraduates in terms of containing specific information on schools attended and family background characteristics.

Weighed against these disadvantages is that to practice law one must take a bar exam. The bar exam gives a measure of learning that can be compared across law schools, which typically is lacking in the undergraduate literature.¹³ While bar passage is not a perfect measure of learning, especially if lower-ranked law schools teach more to the test than do higher-quality schools, the importance of passing the bar for one's legal career makes it a highly relevant measure of legal knowledge that also relates to the return one can expect on investment in law school. Furthermore, affirmative action is very aggressive in law school admissions, which makes this field of particular interest in understanding the effects of racial admissions preferences on student outcomes.

To illustrate the last point, we present results analyzing data on admissions decisions for Michigan's law school in 2002, the same year when the Supreme Court took up the *Grutter v. Bollinger* case.¹⁴ The data contain the applicant's race/ethnicity, LSAT scores, and undergraduate grades,

¹³Earnings could be another measure, but affirmative action in the labor market distorts comparisons across racial/ethnic groups. One caveat with using bar passage is that the difficulty of the exam varies across location.

¹⁴These data, along with data in 2003, were obtained by Richard Sander via a Freedom of Information Act request.

as well as whether the applicant was admitted. We create an ‘academic index’ to aggregate LSAT scores and undergraduate grades into one measure¹⁵ and then standardize the measure so that the mean is zero and has standard deviation one for the applicant pool.

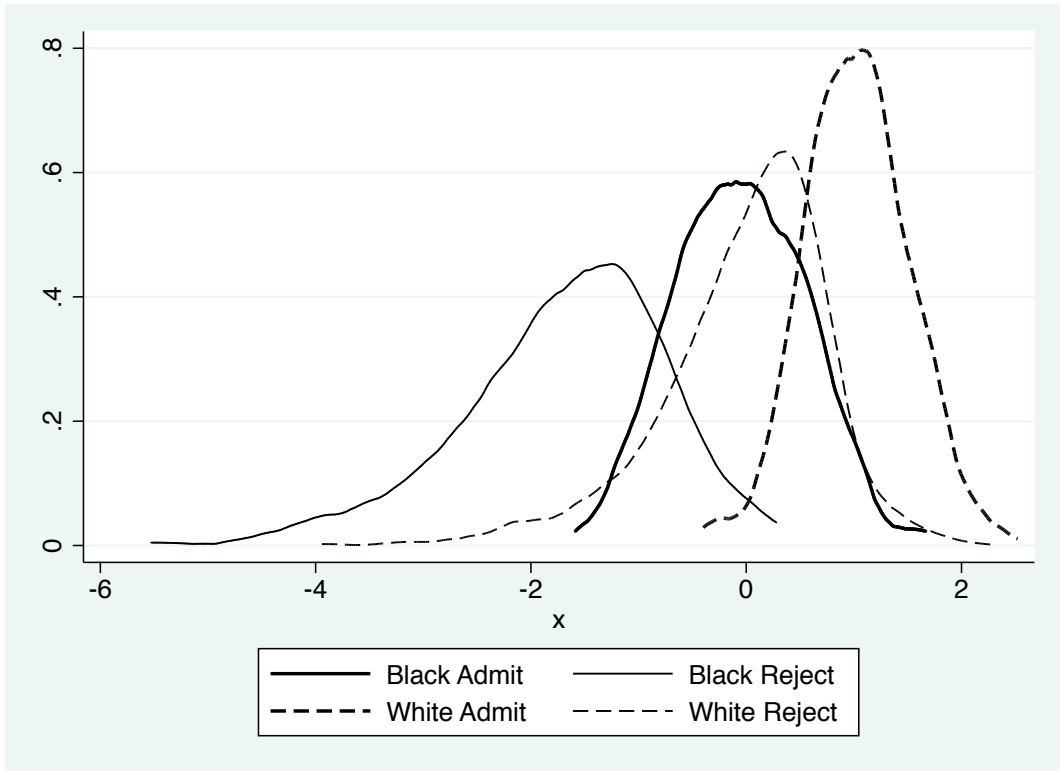
Figure 2 shows the distribution of the re-normed academic index separately for blacks and whites based on whether or not they were admitted. For both racial groups, those admitted had significantly higher index values. However, white students who were rejected generally look academically stronger than the black students who were admitted. Indeed, black applicants in 2002 were admitted to Michigan’s law school at a slightly higher rate than white students, 26.9% to 23.3%, despite the average academic index of black applicants being around 1.5 standard deviations lower than the mean white applicant.

More striking is where the black admit distribution lies relative to the white admit distribution. The median black admit had an academic index at the second percentile of the white distribution, and the seventy-fifth percentile of the black admit distribution was at the eighth percentile of the white distribution.¹⁶ The difference between the black and white admit distributions is not all due to affirmative action: if the African American academic index distribution is below the white distribution, this would produce a difference in the incoming qualifications of black versus white students even in the absence of affirmative action. However, the fact that these distributions are almost non-overlapping is suggestive of a large amount of race-based preferences in admissions being given to African American students. The University of Michigan is by no means an outlier either. The data show that the extent of preferential admissions for black students is even more pronounced at other elite public law schools, such as the University of Virginia and the University of Wisconsin. In both of these cases, the median black admit had an academic index that would place him below

¹⁵The formula follows Sander and Bambaauer (2012 page 896) and is given by $((LSAT - 120) \times 10) + (100 \times UGPA)$. One concern may be that our results are driven by lower black undergraduate grades due to affirmative action in undergraduate admissions. Using just the LSAT produces the same patterns.

¹⁶If we instead focus on enrollees, the median black enrollee has an academic index less than the first percentile of the white enrollee distribution. Indeed, no white enrollees had academic indexes that were lower than the 75th percentile of the black distribution.

Figure 2: Distribution of Academic Indexes for Applicants to Michigan Law by Race and Accept/Reject Status



Data from 2002. Academic index is a weighted average of the applicant’s LSAT score and undergraduate GPA and is normalized to be mean zero and standard deviation one for the population of applicants.

the first percentile of the white admit at the same school.¹⁷

The data also reveal that affirmative action works differently for blacks and Hispanics. While affirmative action is very much present for Hispanics (the median Hispanic admit at Michigan is at the 9th percentile of the white admit distribution), the median Hispanic admit is at the 78th percentile of the black admit distribution. Hispanic admission rates were also lower than those for blacks, despite having on average better test scores and undergraduate grades.¹⁸

These tabulations indicate that African American students receive a significant amount of pref-

¹⁷The median black admit at the University of Texas and the University of North Carolina was below the 2nd and 7th percentile respectively of the white admit distribution. These results refer to two admissions cycles for each school: 2006 and 2007 for Michigan and Wisconsin and 2005 and 2006 for North Carolina, Texas, and Virginia. These data were obtained by Richard Sander via a Freedom of Information Act request.

¹⁸Estimates from a logit model of admission with the right-hand-side variables being LSAT score, undergraduate grades, and indicators for race/ethnicity categories reveal that the admission advantage for blacks is roughly twice the advantage for Hispanics.

erential treatment in law school admissions. The extent of affirmative action in law schools is sufficiently large and the sorting on prior academic credentials is so prevalent that if mismatch is not present, law schools should be almost perfectly differentiated by average quality differences (as in Panel (a) of Figure 1). In other words, either the large differences in white and black pre-law school academic characteristics driven by affirmative action policies lead to mismatch or the effect of law school quality is ubiquitous. In the former case, affirmative action could lead to negative effects on African American law students, whereas in the latter case black law students would be unambiguously better off due to affirmative action because of the positive effects of law school quality on outcomes. We now turn to a discussion of the often contentious debate about the empirical support for both of these hypotheses.

4.1 Graduation and Bar Passage

4.1.1 A Framework for Analyzing the Quality-Fit Tradeoff in Law Schools

As we discuss in detail below, the empirical literature examining the returns to law school quality and the existence or extent of mismatch is extremely contentious. A core source of the disagreement among studies examining affirmative action in law schools is the difficulty in producing causal estimates of the effect of attending different-quality law schools on measured outcomes for black and white students. This difficulty is due to the fact that students select schools and schools select students based on factors that are both observable and unobservable to researchers. As a result, there is much concern that the estimates of law school quality are biased by the existence of unobserved differences in underlying student ability that exist across the school quality distribution.

Studies of the effect of affirmative action-driven mismatch on student outcomes seek to understand whether any overall positive effects of school quality are outweighed by the effect of lower academic performance relative to one's peers. Put differently, the empirical challenge is to under-

stand whether being lower in the academic ability distribution in a school has independent negative effects on outcomes, and if so how large these effects are in relationship to the gains from attending a higher quality law school. To fix ideas, denote Y_i , G_i , L_i , and X_i as individual i 's outcome (such as bar passage or earnings), normalized law school grades,¹⁹ law school quality tier, and other observable characteristics such as pre-law school academic credentials. Assume that Y_i and G_i can be decomposed as follows:

$$Y_i = \alpha_1 + G_i\alpha_2 + L_i\alpha_3 + X_i\alpha_4 + \epsilon_i \quad (1)$$

$$G_i = \gamma_1 + L_i\gamma_2 + X_i\gamma_3 + \eta_i \quad (2)$$

which can also be expressed as the following reduced form equation:

$$Y_i = \delta_1 + L_i\delta_2 + X_i\delta_3 + \mu_i. \quad (3)$$

If mismatch negatively influences post-graduation outcomes, then $0 > \delta_2$, where $\delta_2 = \alpha_3 + \alpha_2\gamma_2$. Taken together, equations (1)-(3) show that we can break out the effect of law school tier into two components: the direct effect and the indirect effect through law school grades,²⁰ where the first effect is expected to be positive and the second negative. Thus, if we can identify δ_2 and test whether it is greater than zero, it amounts to a test that matching effects are larger than college quality effects.

Virtually all models that researchers have used to test for mismatch in law schools are variants of equations (1)-(2) or (3), with different assumptions regarding the distribution of the errors. There

¹⁹By normalized grades we mean grades have been normalized to have zero mean within each school. This can alternatively be interpreted as a measure of student relative rank within each school.

²⁰Some researchers, (e.g., Ho 2005) have raised concerns about controlling for law school grades in the outcome equations because they are themselves affected by law school quality. Rosenbaum (1984) shows that controlling for a variable that is affected by the treatment can produced biased treatment effect estimates. However, he also shows that in a selection-on-observables setting, both the treatment and the mediating variable coefficients are identified as long as the observable characteristics are sufficient to account for selection into the treatment.

are two issues with this approach. First, it is restrictive in such a way that precludes the bottom panel of Figure 1. Namely, *everyone* is either hurt or helped by attending an elite school: the test for mismatch does not depend on the match between the school and the student. Hence, the tests in the literature are generally attempts to measure average effects of law school quality, with the average effects possibly varying by race.

This important caveat aside, the second issue is overcoming the selection problem. We now consider the barriers to recovering unbiased estimates of δ_2 , or alternatively α_3 and $\alpha_2\gamma_2$. For simplicity, we focus the discussion on identifying δ_2 , but the mechanical relationship between δ_2 , α_2 and γ_3 means that this discussion generalizes to these other parameters as well. The underlying difficulty with identifying δ_2 is that law schools have access to information about students that are unobservable to researchers and that also are likely to affect student outcomes. Thus students who are similar in terms of observables but who go to law schools of differing quality likely also differ in terms of these unobservable factors: the students attending more-elite law schools have higher unobserved ability, conditional on observed ability measures.

In this setting with selection-on-unobservables across law school quality tiers, researchers have used several methods to identify net effect of quality and matching. One way to proceed is to use what we term a “within-race estimator,” comparing outcomes among URM students at different law school quality tiers after adjusting for observed ability measures. To the extent there is overlap in observed ability across quality tiers, however, it is likely the case that unobserved ability is higher among students in the higher quality tier, biasing δ_2 upward. Hence, the within-race estimator likely provides an over-estimate of the gains to attending a more prestigious law school.

Another way to proceed is to combine within-race estimators in a difference-in-difference setting. This method ostensibly uses whites as a control group for selection amongst URM students. Whether such a comparison is valid and the direction of any bias rests heavily on assumptions about how

the joint distribution of observed and unobserved ability differ across groups. This difference-in-difference method can be used to see whether matching effects are important when the selection on unobservable ability is the same across races for each quality tier. Then, the whites can be used as a control group to purge the bias of the within-URM estimates because the bias in δ_2 among whites will be identical to the bias for non-whites. This is, to be sure, a very strong assumption, and if it is violated estimates of mismatch from equations (1)-(2) or (3) will be biased as well. Here, the bias can go in either direction. First, imagine that URM students have higher unobserved ability than whites (conditional on observed ability). Then, the tier effects will be biased upward more for minority students than for white students, which will lead one to overstate mismatch when comparing white versus minority students. Conversely, if unobservables among URM students are worse than for whites conditional on observed ability, then the difference-in-difference estimator will understate the extent of mismatch. Note that this is also only a test of whether matching effects matter as the coefficient on tier itself is biased upward for the reasons described above.

A third method of identifying the net effect of quality and matching is what we term the “cross-race estimator.” This method compares outcomes of black students to white students unconditional on quality tier. More formally, the cross-race estimator can be written as:

$$Y_i = \beta_1 + Black_i\beta_2 + X_i\beta_3 + \omega_i, \tag{4}$$

where *Black* is an indicator for whether the student is African American. Conditional on observed ability, the mismatch hypothesis predicts that black students will have worse outcomes than whites (i.e., $\beta_2 < 0$). This model identifies mismatch as long as the distribution of unobservables among African American students is the same as the distribution of unobservables among white students. As Rothstein and Yoon (2009) argue, it is likely that black students have lower unobserved ability than white students based on the fact that they tend to perform worse than white students at the

same school and with the same admissions credentials. In this case, β_2 will be biased downward, which will overstate the importance of matching.

As this discussion highlights, the validity of the different approaches to measuring quality and matching effects in law school rely strongly on the underlying assumptions about how student unobserved ability is distributed across college quality tiers, across racial groups within each tier, and across racial groups. Due to these identification concerns, no estimate of the extent of mismatch is perfect. However, we can use this framework to guide our interpretation of the differences in results across the methods used, which will help us synthesize this contentious literature. We now turn to a discussion of this literature, beginning with the seminal work of Sander (2004).

4.1.2 Sander (2004)

The law school debate began with Sander (2004), who argued that large admission preferences resulted in blacks learning less in law school and consequently failing the bar at a higher rate.²¹ Sander speculated that these learning effects were sufficiently large to actually result in fewer black lawyers than if racial preferences were not in place. He also argued that the loss of black lawyers occurs despite racial preferences increasing the number of black students who enroll in law schools as i) those students who are at the margin of being admitted have a relatively low probability of passing the bar exam and ii) being severely mismatched results in low bar passage rates.

Using data from the Law School Admission Council Bar Passage Study (BPS),²² Sander (2004) shows that, aside from those attending historically black universities, black students perform very poorly in law school in terms of grade point average. In particular, the bottom decile of the grade distribution contains between 44 and 52 percent of black students, depending on law school tier.²³

²¹Most of this literature focuses on black-white differences due to affirmative action. To our knowledge, there has been little attention paid to mismatch effects among Hispanic students in law school.

²²This data set follows 1991 law school matriculants. Participation in the study was high for both law schools and their students: 95% of law schools participated and 80% of the students signed consent forms. The primary disadvantage of the data is that school quality was aggregated into six broad tiers.

²³Recall that black students had similar expectations to white students with regard to their expected law school class

Since grades are a relative measure of performance and since there are large differences in entering credentials between blacks and whites, it is perhaps not surprising that black students are more likely to be at the bottom of the GPA distribution. Indeed, using 1995 National Survey data that includes school identifiers, Sander (2004) shows that the relative rank of one's undergraduate GPA and LSAT score within a school are strongly predictive of law school GPA. This finding suggests that policies designed to bring lower-ranked students into more elite schools could lead to them to perform substantially worse relative to their peers. While Sander (2004) does not explicitly estimate equation (2), this evidence forms the basis of his argument that γ_2 is both negative and sizable in magnitude.

Sander estimates equation (1) using the BPS data. While he estimates versions of equation (1) for both law school graduation and first-time bar passage, he focuses on the bar passage results. The resulting literature also focuses on this outcome, so we focus our attention on bar passage. The estimated coefficients on (normalized) law school grades, α_2 , and tier, α_3 , are positive. Furthermore, conditioning on law school grades and tier leads to no effects of race on bar passage. This drives the remainder of Sander's analysis that race is not relevant for bar passage. Sander's approach here is a mixture between the within-race and cross-race estimators. The estimates of α_2 and α_3 are mostly identified off of white students, because they comprise the majority of his sample. For the reasons described in the previous section, these coefficients are likely biased upward. Including an indicator for *Black* is in the spirit of the cross-race estimator, but how the biases on α_2 and α_3 combine with bias on the *Black* indicator due to differences in observables across racial groups is unclear. As a result, signing the bias in this regression on the *Black* coefficient is very difficult.

Sander then breaks out bar passage rates by LSAT scores bins for both blacks and whites. Within each bin, white students are more likely to pass the bar. Since Sander finds no race effects rank. To the extent that these subjective expectations reflect actual beliefs, the results suggest black students were either unaware of the extent of affirmative action or unaware of how credentials translate into law school performance.

on bar passage, he uses this to justify the assumption that black students would pass the bar within each of these bins at the same rate as whites had affirmative action not been present. He then conducts a decomposition analysis that estimates what first-time bar passage rates for blacks would have looked like absent racial preferences. This is done by assuming that all within-LSAT bin differences between blacks and whites are due to mismatch effects. The results of this exercise suggest that over 45% of the black bar failure rate (or about 57% of the difference in the black-white failure rate) is due to mismatch.²⁴ This finding implies that attending a higher-tiered law school lowers the probability of passing the bar on the first try, particularly for less-academically-qualified students. The decomposition estimates thus point to a substantial amount of mismatch from affirmative action that worsens the long-run outcomes of many black law students.

Sander argues that the reason for this finding is that affirmative action results in black students attending schools where they lack the necessary academic preparation, which in turn lowers the quality of their education. For several reasons, however, his estimates likely overstate the amount of mismatch. First, because black students have lower academic index values than whites overall, this is likely the case within each bin as well. As a result, even within LSAT bins, the black and white students are unlikely to have the same academic ability levels. Second, the assumptions that underlie this approach are identical to those that support the cross-race estimator, as Sander is assuming that there are no other factors (such as unobserved ability) that would drive differences in bar passage rates within each LSAT bin. Similar to the cross-race estimator, this leads to an upwardly-biased estimate of mismatch.

Note that Sander does not estimate (2), but he does do so in his response to Ho (2005a). In Sander (2005a), the estimated coefficient on tier in equation (2) for first-time bar passage, γ_2 , is sufficiently negative that $0 > \alpha_2 + \alpha_3\gamma_2$. This is surprising for two reasons. First, because the

²⁴The other reason African American students fail the bar less is because they are more likely to have lower academic index levels. As Sander (2004) highlights, even taking his calculations at face value, eliminating affirmative action could not close the black-white bar passage gap.

estimates are identified predominantly off of whites because they comprise the majority of the sample, they are within-race estimates. As discussed above, such estimates typically are biased away from finding mismatch effects. Second, because whites are the majority of the sample, the results imply that the negative net tier effects are also present for whites. Taken at face value, the estimates in Sander (2004) suggest that attending a more elite law school lowers one's chances of passing the bar regardless of one's entering credentials or race. The specification thus only allows for two possibilities: everyone benefits from attending a more elite school or no one benefits. A problematic conclusion one could draw from Sander's results is that *everyone* is harmed by going to a more elite law school, as the negative effect on GPA swamps the positive direct effect of school quality. Furthermore, if there are cross-race differences in mismatch effects, generalizing these estimates to a sample of African American students could yield misleading conclusions about the extent of mismatch.

All subsequent law school mismatch studies are subject to this same criticism that they imply everyone is either helped or harmed by attending a more elite school. The exception is one part of the analysis in Sander (2005b). Here, Sander allows curvature in the relationship between grades and outcomes. He bins percentile class rank and, within each bin and law school tier combination, estimates the joint probability of graduating and passing the bar on the first attempt. The results reveal a substantial penalty for finishing at the bottom of the class across law school tiers. This suggests the possibility of increasing eliteness being beneficial up to a point. As long as a student is reasonably competitive, the student can reap the benefits of the elite education. Hence, for the objectives of maximizing black bar passage, there may be an optimal affirmative action policy that falls in between the current policy and race-blind admissions.²⁵

²⁵This point is not explored further in Sander (2005b), as this piece was primarily a reply to criticisms levied against Sander (2004).

4.1.3 The Ensuing Debate

The difficulties posed by selection-on-unobservables in identifying mismatch effects in law school has led to a vigorous and often contentious debate about the validity of Sander's conclusions. Much of the ensuing work on law school mismatch focuses on identifying the effect of law school tier net of the grade effect (i.e., δ_2 from equation (2)). A negative estimate of the effect of quality tier in such a model provides evidence of mismatch. While examining the net effect of tier sidesteps the difficulties associated with the endogeneity of relative law school grades, these analyses face similar identification concerns related to selection-on-unobservables that form the basis for many of the objections to the findings in Sander (2004).

Rothstein and Yoon (2009) is the only other paper to pursue black-white comparisons in order to assess mismatch effects. They estimate probit models of bar passage on observed students credentials as well as an indicator for whether the student is black, akin to equation (4). Estimates of this model show that blacks are much less likely to graduate and to pass the bar even conditional on the observables, which is consistent with the mismatch hypothesis. However, at the bottom of the credential distribution there is very little overlap in credentials between black and white students due to racial preferences affecting the extensive margin. Those white students with poor credentials are likely to have stronger unobservables than similarly-credentialed black students, else they would not have been accepted to any law school. To address this concern, Rothstein and Yoon (2009) report estimates using both white and minority students that focus only on the top four quintiles of the academic index, which is where the white and black academic index distributions have the most overlap. Examining the top four quintiles only removes the negative effect of tier on law school graduation and cuts the effect on bar passage by more than half, rendering it insignificant.²⁶ The problem with focusing on the top four quintiles, however, is that three quarters of the black students

²⁶However, their estimate still is negative and sizable in magnitude, on the order of 3 percentage points. Furthermore, the 95% confidence interval cannot rule out large mismatch effects.

fall in this bottom quintile (Rothstein and Yoon 2009, page 18), and it is these students who are most at risk of being mismatched.

Rothstein and Yoon (2009) also estimate equation (3) separately by race, which is the within-race model. This is also the strategy pursued by Ho (2005) and one of the strategies in Ayres and Brooks (2005) and Williams (2013). Rothstein and Yoon (2009) argue that this model provides a lower bound of the extent of mismatch, similar the arguments described above. Ho (2005) and Rothstein and Yoon (2009) both find estimates that suggest law school tier has no effect on eventual bar passage using the within-race analysis. Because these findings are likely to be heavily influenced by selection-on-unobservables, they are not necessarily inconsistent with those in Sander (2004). Indeed, he finds evidence of mismatch in models that are biased towards finding mismatch, and the within-race estimates find no effect of law school quality tier in models that are biased against finding mismatch effects. A sensible reading of this evidence suggests the correct answer is somewhere in between, which would indicate modest mismatch effects exist.

Note that Rothstein and Yoon (2009) do not find quality tier effects for black students, while Sander (2005a) does. Part of this is due to the differences in their dependent variables. Rothstein and Yoon (2009) examine *eventual* bar passage, while Sander (2004) examines bar passage on the first try. There is evidence that the black-white gap in first time bar passage is larger than the gap in eventual passage (Williams 2013). This is an important distinction, however, as failing the bar on the first try has substantial costs in terms of increased study time, reducing earnings, and putting lawyers at a competitive disadvantage in the labor market. It is thus not surprising that Rothstein and Yoon find smaller mismatch effects than Sander, given that they are examining an outcome that has a smaller racial gap associated with it.

Another strategy pursued by Rothstein and Yoon (2009) is the across-race difference-in-difference approach discussed above. They estimate equation (3) separately for white and black students, and

then they test for equality of the δ_2 coefficients across groups. What is striking is that positive effects of law school quality emerge for whites but not for blacks. This suggests either that (i) selection on unobservables is present for whites but not for blacks or (ii) matching effects are effectively canceling out quality effects for blacks. Under the assumption that selection on unobservables is similar for blacks and whites, however, these estimates are consistent with mismatch: African American students at selective law schools are about 2.5 percentage points less likely to pass the bar than white students at selective law schools.

Williams (2013) argues that their results further understate mismatch due to the noisiness of the law school tier measure as well as their focus on eventual rather than on earlier bar passage. Williams (2013) finds negative effects of selectivity on both first-time and eventual bar passage, although only the former are statistically different from zero. Additionally, the estimated effects increase when the top two tiers are compared to the bottom two tiers as opposed to comparing the top two tiers to the bottom four, as in Rothstein and Yoon (2009). One caveat is that historically black colleges and universities are in this bottom tier and may operate differently from traditional law school environments.²⁷

As a way of mitigating differences in unobservables among black students attending different law school tiers, Ayres and Brooks (2005) and Williams (2013) pursue an identification strategy in the spirit of Dale and Krueger (2002).²⁸ Namely, comparisons are made between black students who attended their first-choice law school and those who reported being admitted to their first-choice law school but attending their second-choice (or lower-quality) law school,²⁹ either due to cost or distance considerations. The descriptive statistics reveal that black students who attend their first-choice school have similar credentials and background characteristics to those who attend their

²⁷Ayres and Brooks (2005) remove these schools in their analysis of eventual bar passage and positive effects of selectivity for blacks when tier is included linearly. Sander (2004) also includes tier in this way. In both cases, it is more appropriate to have only two tiers or dummy out the effects of each tier.

²⁸See Section 5.5 for a more complete description of this methodology.

²⁹For ease of exposition, we refer to these students as attending their second choice school even though the school they attend could be even further down their list.

second choice. However, first-choice students placed more importance on academic reputation and less importance on cost than second-choice students. Second-choice students then attended schools that were relatively less selective. A core assumption underlying this method is that students who were admitted to the same types of schools but who attend schools of differing quality have the same unobservable characteristics.

Results using this identification strategy point towards small positive effects of attending a more selective school on graduation that often are not statistically significant. But, both Ayres and Brooks (2005) and Williams (2013) find that black students who attend their second choice (and thus lower-quality) school are more likely to pass the bar the first time they take it, which is evidence consistent with mismatch. Ayres and Brooks (2005) find insignificant effects for passing the bar at a later date, however. Consistent with earlier bar passage results, estimated mismatch effects on the bar exam are strongest earlier on in the career. One interpretation of this finding is that the negative effects of mismatch will be undone over time, but the time pattern of mismatch is important because of the costs associated with delaying bar passage.³⁰

The law school debate over Sander (2004) has been especially contentious. We find the evidence suggesting that shifting African Americans to less-selective schools would increase bar passage rates, particularly for first-time passage, to be fairly convincing. This is especially the case since the low quality of the data would tend to bias estimates away from finding mismatch. On the other hand, an argument could be made that the data are too noisy and provide sufficiently imprecise information on actual law school quality that they preclude one from drawing any concrete conclusions regarding mismatch. Regardless, the law school debate makes clear that this is a question that merits further attention, where more definitive answers could be answered with better data. Our hope is that better data sets soon will become available.³¹

³⁰Williams (2013) also argues that the measurement error in the law school tier variable results in attenuation bias. Hence he instrumented for law school tier with the second choice variable, resulting in massive mismatch effects. However, the estimated effects are so large as to not be plausible.

³¹In *Sander v. State Bar of California*, the California Supreme Court ruled unanimously in favor of Sander's petition

4.2 Labor Market Outcomes

Sander (2004) also argued that affirmative action lowers the earnings of black lawyers, using data from After the JD (AJD). In 2002, the AJD surveyed those who became lawyers (i.e. passed the bar) in 2000.³² As with bar passage and graduation, the mechanism for the earnings result is that attending a more elite school has a positive direct effect on earnings but also a negative indirect effect through law school grades. Sander argues that the negative indirect effect is stronger than the positive direct effect, again assuming that differences in normalized law school grades will disappear after affirmative action is eliminated. That this impact still is present even after selecting on bar passage, which itself appears to be affected by mismatch, is notable.

Rothstein and Yoon (2009) is one of the few studies that also addresses the earnings question, using data from the BPS dataset. The BPS followed a sub-sample of the initial participants four to six months after they graduated from law school. Using the same identification strategies they used for bar passage, Rothstein and Yoon find higher earnings for blacks from attending more selective schools.³³

Beyond the differences in identification strategies, a potential way to reconcile the two sets of results comes from Lehmann (2013). Lehmann develops a model of hiring and promotion when firms face diversity pressures. Using the AJD, Lehmann shows that black lawyers have better early labor market outcomes, consistent with a model where the cheapest way to satisfy diversity pressures is with newly-minted lawyers. After hiring, however, black lawyers receive worse tasks and are less likely to be promoted. Lehmann's results point to the difficulty of comparing black and white outcomes when affirmative action is present in the labor market, as well as pointing out that these

for access to a de-identified version of public law school admission data, stating "We conclude that under the common law right of public access, there is a sufficient public interest in the information contained in the admissions database such that the State Bar is required to provide access to it if the information can be provided in a form that protects the privacy of applicants and if no countervailing interest outweighs the public's interest in disclosure."

³²The AJD has the actual school that the lawyer attended. However, law school grades are self-reported in bins, so adjustments need to be made to renormalize them.

³³They also find that black students at elite schools were also more satisfied with their jobs than matched students of either race at less-elite institutions.

sorts of comparisons may change with the age of the cohort.

4.3 The extensive margin

In contrast to undergraduate institutions, almost all law schools are selective. Hence, an affirmative action ban may result in some minority law school applicants, who would have been admitted prior to the ban, being unable to obtain admission to any law school. Assuming that the number of black applicants would remain unchanged and that their admissions and application behavior would now match those of whites with similar LSAT scores and grades after an affirmative action ban, Wightman (2003) estimates a 14.1% drop in the number of black applicants who would be admitted to law school if an affirmative action ban were passed that was based on applications for law schools admissions in 2001. This is the number Sander (2004) uses in his calculation of the extensive margin.

Both Chambers et al. (2005) and Rothstein and Yoon (2008) argue that this number is too low. Rothstein and Yoon (2008) use a grid model for 1990-91 and show there would be a 60% drop in law school enrollees, with a 90% drop at the most-selective schools. Chambers et al. (2005) extend the grid model in Wightman (2003) to other years, showing that in 2003 and 2004, the most recent years available at the time, the decline would be much higher, at 32.5% in 2004.³⁴ Chambers et al. (2005) also argue that the drop in admissions may translate into lower take up rates, but there are no data to provide support for or against this claim.

There are at least two reasons, however, why estimates of the extensive margin may be overstated. First, the grid model assumes after a ban that black applications will look like white applications. But, the least selective law schools are historically black institutions, which are likely

³⁴Sander (2005b) argues that those two years are outliers driven by a surge of white applications, looking over the previous ten years suggests that the 14.1% number is representative. The magnitude of the extensive margin may depend on whether what was happening in 2003 and 2004 reflects permanent or temporary change in application behavior.

more appealing to black applicants than white applicants.³⁵ How historically black institutions would respond to a surge in black applications following an affirmative ban is an open question.³⁶ Second, universities may respond to affirmative action bans by looking at other ways of increasing diversity by, for example, putting more weight on characteristics that are correlated with particular races.³⁷ Yagan (2014) provides evidence of this, looking at law school applications from an elite college to UC Berkeley and UCLA before and after Proposition 209, which banned the use of racial preferences in admissions. Yagan shows that although admissions probabilities fell substantially for black applicants after Proposition 209, black applicants still had an advantage in admissions over white applicants after controlling for a set of observed credentials. Long and Tienda (2008) show qualitatively similar results in Texas.

5 Undergraduate decisions and outcomes

5.1 The Extent of Racial Preferences

Unlike in the law school setting, affirmative action is present only at a subset of all undergraduate institutions. In *The Shape of the River*, Bowen and Bok (1998) argue that only 20 to 30% of four-year colleges practice racial preferences, as most schools simply are not selective. These results are consistent with Kane (1998) and Arcidiacono (2005), who both estimate models of admissions using nationally-representative data in which students self-report where they have been admitted. Both find that racial preferences are large at the top schools, with little scope for racial preferences at the bottom schools.

At these top schools, however, racial preferences result in substantial differences in academic

³⁵Rothstein and Yoon (2008) predict an 84% drop in minority enrollees at HBCUs, which they acknowledge is too high. However, they argue that the movement towards HBCUs would come in part from other law schools resulting in those schools being even less diverse.

³⁶Of course, having the extensive margin being somewhat removed due to historically black institutions may be unsatisfying, harkening back to ‘separate but equal’ education.

³⁷See Chan and Eyster (2003) and Fryer, Loury, and Yuret (2008) for models of how these adjustments could occur in equilibrium.

backgrounds between majority and minority enrollees.³⁸ Using the College and Beyond data set, which focuses on a set of highly-selective colleges,³⁹ Arcidiacono, Khan, and Vigdor (2011) show within-school SAT score gaps of at least 140 points.⁴⁰

It is typically difficult to get information on the accept-reject decisions of universities. But there are data from the University of California (UC) system, acquired under a Freedom of Information Act request, that can be used to open the black box of university admissions. This individual-level data set covers applicants from the years 1992 to 2006 and follows these individuals through to graduation.⁴¹ While missing many covariates that are likely of importance (such as gender), the data contain a continuous measure of academic preparation. Namely, they include an academic index, which is a weighted average of the student's SAT score and high school grades.⁴² We standardize the academic index so that it has zero mean and standard deviation one for the set of individuals that applied to a UC school during this period.

The data show that half of the minority applicants who applied to UC Berkeley were admitted compared to 35% of white applicants. Minority applicants to UC Berkeley had significantly higher probabilities of being admitted despite academic indices that were almost one standard deviation lower than white applicants. Figure 3 shows the academic index distribution for those admitted and rejected by UC Berkeley for both minority and white applicants. For both minority and white students, admits have academic indices on average around one standard deviation higher than rejects. But the median white reject has an academic index higher than the median minority admit.

In fact, the median minority admit's academic index is at the seventh percentile of the white admit

³⁸Kane (1998) points out that even in the absence of racial preferences minorities would on average have lower test scores within each school given that lower minority test scores would place them more on the margin of being admitted. Given the magnitude of the differences in the entering credential of enrollees, however, it is clear that racial preferences at top schools are large.

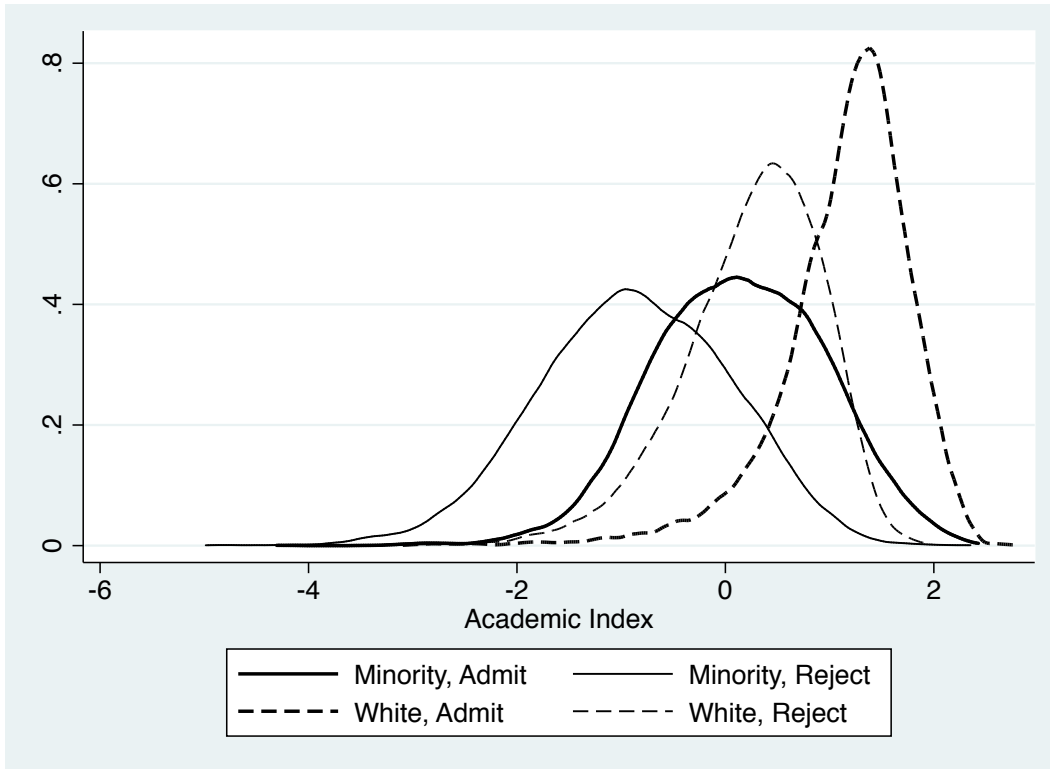
³⁹This is the same data set as used in *The Shape of the River*.

⁴⁰The SAT score here includes only the verbal and math sections.

⁴¹The year the individual applied is given only by its three-year interval. In 1998, Proposition 209 was implemented, which banned the use of racial preferences in California. Hence, two (three) three-year periods are offered before (after) the implementation of Proposition 209. See Antonovics and Sander (2013) for a fuller description of the data.

⁴²The SAT score is rescaled on a range from 0 to 600 and high school grades are rescaled from 0 to 400. The two numbers are then added together.

Figure 3: Distribution of Academic Indexes for Applicants to UC Berkeley by Minority and Accept/Reject Status



Data source: UCOP, years 1995-1997. Academic index is a weighted average of the student’s SAT score and high school GPA and is normalized to be mean zero and standard deviation one for the population of applicants.

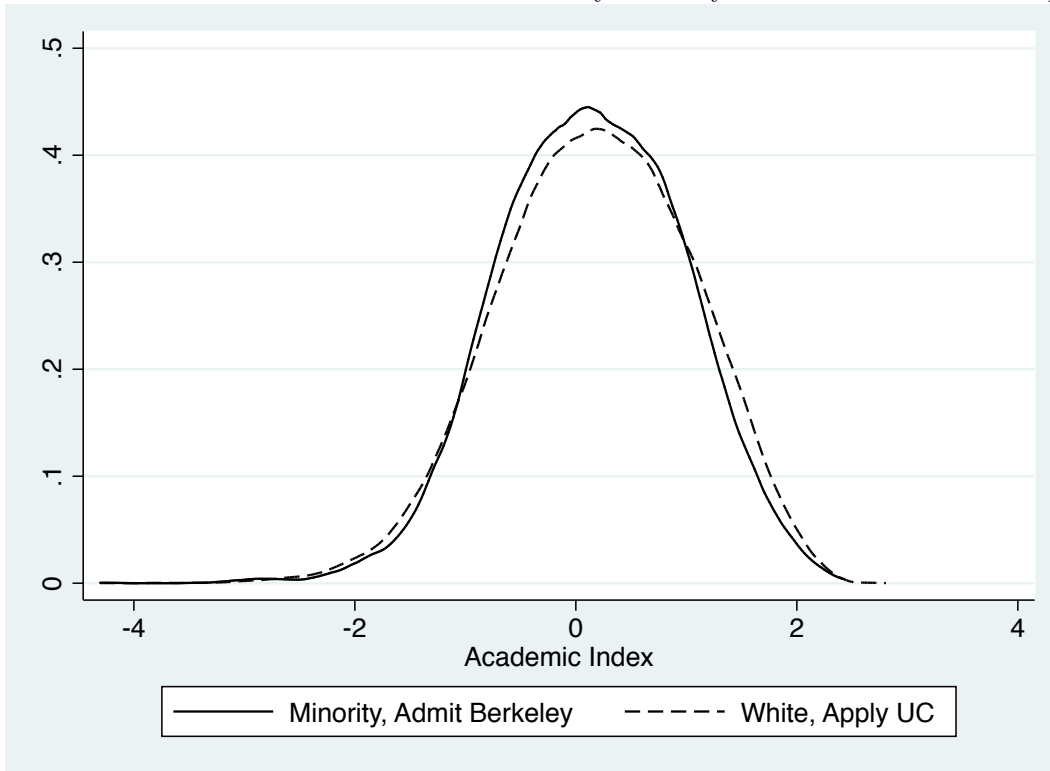
distribution.

To further demonstrate the extent of racial preferences at UC Berkeley during this period, Figure 4 compares the academic index distribution for minority admits at UC Berkeley to the academic index distribution for whites who applied to *any* UC school. The distributions look remarkably similar: randomly drawing from the pool of white students who applied to any UC school would produce an academic index distribution similar to that of minority admits at UC Berkeley.⁴³

As shown in Antonovics and Backes (2014), UC Berkeley had the largest racial preferences in the UC system. After conditioning on a set of observables, minorities had increased chances of

⁴³Looking at enrollees instead of admits shifts the distribution of academic indexes to the left, with the median minority index falling by 0.24 standard deviations. This would place the median minority enrollee at Berkeley below the 9th percentile of white enrollees at Berkeley.

Figure 4: Distribution of Academic Indexes for Minority Berkeley Admits and White UC Applicants



Data source: UCOP, years 1995-1997. Academic index is a weighted average of the student's SAT score and high school GPA and is normalized to be mean zero and standard deviation one for the population of applicants.

admission at the top four UC schools, ranging from 40% at UC Berkeley to 21% at UC San Diego. Minority advantages were smaller at the next set of schools, likely for two reasons. First, there is less scope for racial preferences at the bottom set of schools. For example, for those who applied to UC Santa Cruz in this period, both whites and minorities had an over 80% chance of being admitted. Second, selection on unobservables is likely stronger at more elite schools as a result of minorities being aware of affirmative action. Those applying to the bottom schools probably were aware that their chances of admission were lower at UC Berkeley, despite racial preferences.

Dillon and Smith (2013) provide some additional information on how affirmative action affects sorting using NLSY97 data. Their paper is focused on characterizing the extent of academic mismatch more generally, which they define as differences between student rank on the Armed Services

Vocational Aptitude Battery (ASVAB) test given to NLSY97 respondents and college rank on an index of quality constructed from several measures of institutional characteristics. This analysis has the value of being nationally-representative, although their small sample precludes them from examining mismatch along the college quality spectrum.

While the aim of Dillon and Smith (2013) is measuring and characterizing overall mismatch in higher education, they show some evidence that speaks directly to mismatch that might be caused by affirmative action. In particular, they show that despite the clustering of black students in lower-quality schools, they are much more likely to be “over-matched” and less likely to be “under-matched” than whites.⁴⁴ This is somewhat mechanical, as blacks have lower test scores and hence are naturally more likely to be over-matched than under-matched relative to whites. Surprisingly, Dillon and Smith (2013) also show that black students are no more likely than whites to be over-matched or under-matched given the same test scores and high school grades. However, this finding only holds when focusing on students at four-year schools. When two-year students are included, blacks are more likely to be over-matched than whites conditional on credentials. It is difficult to fully explain this result with the evidence given in their paper, but one likely explanation is that more black students are over-matched at two-year schools and non-selective four-year schools relative to white students.⁴⁵

The findings of Dillon and Smith of no over-match among black students stands in contrast to much of the prior literature. In particular, Bowen and Bok (1998) provide extensive evidence of over-match among African American students. It thus is instructive to understand why these studies differ in their conclusions. One core difference between them is that Dillon and Smith consider the entire distribution of colleges, while Bowen and Bok examine just elite schools. Arcidiacono, Khan,

⁴⁴ “Over-matched” (“under-matched”) is defined as the ASVAB rank being more than 20 points below (above) the college rank.

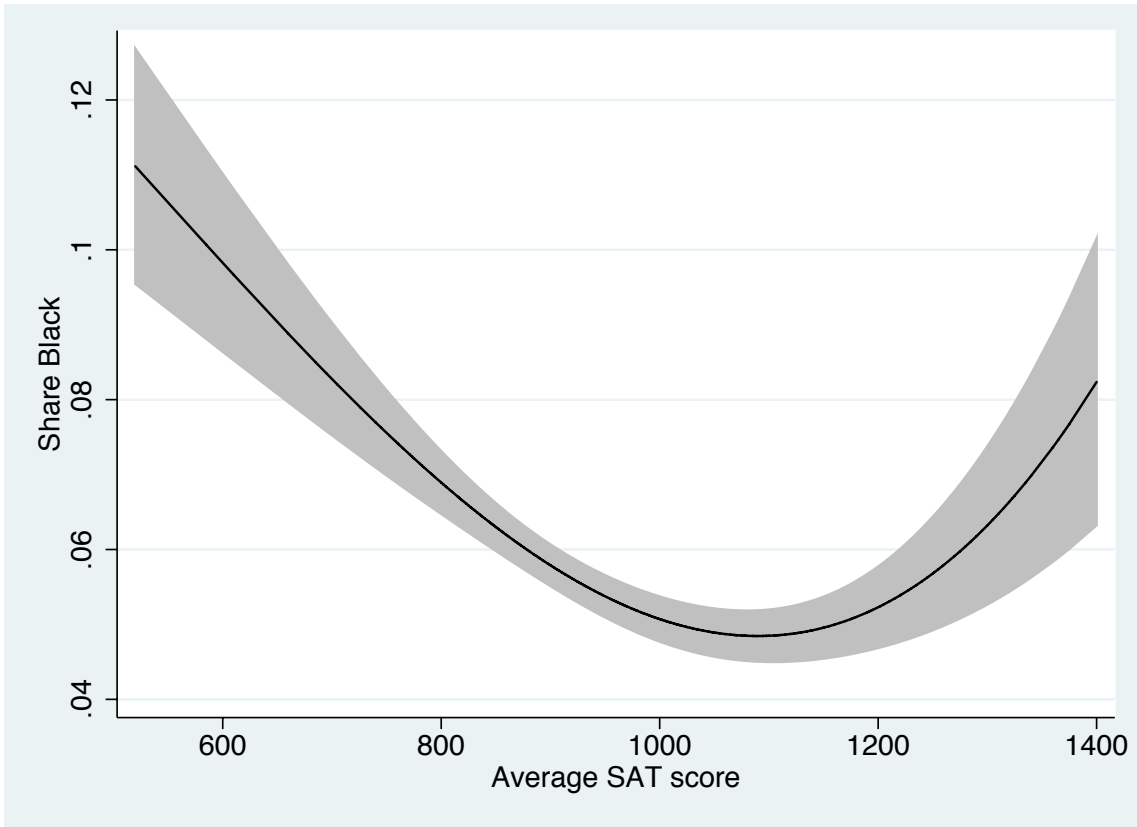
⁴⁵ Because Dillon and Smith categorize a student as overmatched if he or she has an ASVAB score that is 20 percentile points below the mean of the school he or she is attending, it is entirely possible that black students are over-matched at community colleges.

and Vigdor (2011) show how the share of black students varies with the average SAT score of the school using data from U.S. News & World Report for 1991. Figure 5 is reproduced from their paper. The figure shows a U-shaped pattern, where the lowest and highest quality schools, as measured by the average SAT scores of their student body, have the highest share of black students. Because of this feature of how African Americans sort into different-quality postsecondary schools, it is possible for there to be little over-match among blacks across the whole four-year quality distribution while still exhibiting over-match among the more-elite schools. It also is important to highlight that Bowen and Bok use older data (their most recent cohorts matriculated in 1989), and the extent of mismatch may have changed over time. This discussion makes clear that much more work needs to be done to understand the extent of over-match among African Americans across the postsecondary quality distribution.

The U-shaped curve shown in Figure 5 is of interest in its own right, and it is useful to consider what factors might cause this pattern. First, affirmative action at the top schools can lower the share of black students in the schools just below them. To the extent that the fall in the black share at this next set of schools makes them less attractive to black students (for example, they are perceived as not having a critical mass), these students may choose to attend lower-quality schools where, conditional on their credentials, they are undermatched. Another possibility is that the middle-tier schools practice less affirmative action. Third, black students may be less-informed on average about their college possibilities. As illustrated by Hoxby and Avery (2013), those who are economically-disadvantaged seem to make poorer college choices than their wealthier counterparts. Economically-disadvantaged blacks may then choose colleges at the bottom of the quality distribution even though they would have been a better match for a school in the middle of the quality distribution.

Regardless of the reasons for the U-shaped pattern, the removal of affirmative action would affect the sorting of black students into selective colleges. By altering minority college enrollment

Figure 5: College percent Black as a function of average SAT score



Data Source: America's Best Colleges produced by U.S. News & World Report for 1991

behavior, removing or restricting racial preferences could impact minority graduation rates, major choices and the returns to college that they experience in the labor market. Below, we discuss what is known about how affirmative action policies and affirmative action bans affect each of these outcomes, as well as highlight questions that remain for future research.

5.2 Enrollments

Before discussing the effects of affirmative action policies on longer-run outcomes such as college completion and wages, we provide a summary of the evidence on how affirmative action affects minority enrollment behavior, both overall and across the higher education quality distribution. Many of the effects of affirmative action on student outcomes flow through changes in whether and where students enroll, and thus understanding the effect of racial admissions preferences on minority enrollment is an important first step in thinking about these broader outcomes.

Figures 3 and 4 suggest that affirmative action policies act to sort minority students to more selective schools, which implies that repealing them will lower minority representation at these schools. Structural models of how students sort into schools by Arcidiacono (2005) and Howell (2010) have been used to predict how the removal of affirmative action affects both the extensive margin (whether students enroll at all) and the intensive margin (where they enroll). Both of these papers have individuals make their application decisions based on expectations regarding their probabilities of admittance. These models predict substantial reductions in the share of minority students attending the most-selective schools following a ban but, because many schools are non-selective, they predict little effect on the share of minorities enrolled in four-year colleges overall.⁴⁶

Recent affirmative action bans allow us to address how affirmative action affects enrollment decisions directly. The enactment of these bans provides quasi-experimental variation in affirmative

⁴⁶Epple, Romano, and Sieg (2008) calibrate a general equilibrium model of how students sort into colleges when colleges are differentiated only in terms of quality (as in Panel (a) of Figure 1). In their model, schools adjust their admissions decisions in reaction to an affirmative action ban so as to attain more diversity by weighting characteristics correlated with race more heavily. They too show substantial predicted reductions in enrollments at top tier colleges.

action policies, and a large portion of studies on which we will focus in this and subsequent sections examines these policy changes. Table 1 shows the timing of the bans in each state as well as the source of each ban's enactment. The bans come from a mix of sources, ranging from voter initiatives to executive orders and court rulings. The variation in time and location of these bans provides useful variation with which to identify the effect of affirmative action policies on college outcomes. Importantly, these bans only affect public universities, so private schools in these states can continue to practice race-based affirmative action.

Consistent with the structural predictions, the general consensus from empirical analyses of affirmative action bans is that elimination of race-based preferences reduced under-represented minority student enrollment at flagship and elite public schools but had little impact on the overall likelihood of four-year college enrollment. Kain, O'Brien and Jargowsky (2005) examine enrollment behavior among Texas public college attendees between 1991-2002. They show evidence that when affirmative action was banned in 1997, minority enrollment dropped substantially at the state flagship schools – University of Texas at Austin and Texas A&M-College Station. In the two years post-*Hopwood*,⁴⁷ the number of black students enrolling at UT-Austin dropped to 68% of the pre-*Hopwood* level and dropped to 72% of the pre-ban level at Texas A&M. Hispanic enrollment also declined by 6 and 25 percent of the pre-ban levels at UT-Austin and Texas A&M, respectively. However, white enrollment jumped by 7% at these schools, which is consistent with affirmative action bans reducing enrollment levels among under-represented minorities at elite public schools when race-based admission policies are eliminated. Long and Tienda (2008) complement these results using administrative data from the pre- and post-ban eras at Texas A&M and UT-Austin. They estimate probit regressions of admission at these schools and show that post-*Hopwood*, these schools no longer gave independent weight to race and ethnicity in admissions.

Hinrichs (2012) and Backes (2012) exploit the timing of affirmative action bans across several

⁴⁷*Hopwood v. Texas* is the court case that led to the affirmative action ban in Texas.

of the states shown in Table 1 to estimate difference-in-difference models of college enrollment. Due to their use of plausibly exogenous affirmative action variation, these are the most informative estimates of how affirmative action affects flagship enrollment among under-represented minority students. Using data from the American Community Survey and the Current Population Survey, Hinrichs (2012) shows that affirmative action bans had no overall effect on the likelihood minority students attended any college or attended a four-year college. However, with IPEDS data on college enrollment by race and institution, he shows bans led to large declines in black and Hispanic enrollment shares at more elite universities. In particular, enrollment by black students at US News top-50 ranked schools declined by 1.6% and Hispanic enrollment at these schools dropped by 1.8%. Because these declines include private schools that were not subject to these bans, these results suggest the affected minority students did not just shift enrollment to other schools of similar quality that were still practicing affirmative action. Furthermore, he shows that the effects on enrollment were larger at schools that had lower admissions rates and higher SAT scores. Overall, these estimates indicate that banning affirmative action had large impacts on the quality of schools in which under-represented minority students enroll. Backes (2012) shows very similar results using an expanded time frame and a more detailed set of institutional quality measures.

Research on enrollment responses to affirmative action bans also shows that most of the enrollment effect came from changes in admissions behavior: application patterns among minority students changed little following affirmative action bans (Card and Krueger, 2005; Antonovics and Backes, 2013).⁴⁸ Thus, affirmative action bans alter admissions rules such that minority students are less likely to gain admittance to more elite in-state, public schools. This reduces the quality of schools these students attend, although it does not affect the likelihood under-represented minority

⁴⁸Long (2004) compares changes in SAT score-sending rates among white and minority students in the late 1990s and finds some evidence of a relative decline in minority applications to the top in-state institutions, especially in California. However, he also shows similar patterns in applications to private and out-of-state colleges, which is inconsistent with these patterns being driven by affirmative action changes.

students enroll in a four-year college. In what follows, we discuss the potential effects of this type of re-sorting on college graduation, college major, and labor market returns to college investment.

An important caveat to all of these studies is that the measured effects are in an environment where racial preferences are still in place at all private schools and at public schools in states that do not have bans. A universal ban on racial preferences would likely have much different effects, for example, on the share of minorities at flagship universities than a local ban (a ban on racial preferences at state institutions). Minorities admitted to the flagship school under a universal ban will have fewer options than those admitted under a local ban, suggesting minority students admitted to the flagship would be more likely to attend under a universal ban than a local ban.

5.3 Graduation

Affirmative action policies can influence graduation rates through two conflicting mechanisms. As discussed in Section 5.2, race-based admissions policies can have a sizable influence on the distribution of minority students across schools of differing quality. The Texas and California experience with banning affirmative action had a large impact on the representation of minority students at the more selective, in-state universities. The mismatch hypothesis predicts that banning race-based admissions could increase the rate at which they graduate from college. This would occur because students are now better matched academically for the schools in which they enroll.

However, there also is a growing literature (discussed below) suggesting that college quality and collegiate resources that are strongly correlated with measures of college quality exert an independent influence on students' abilities to obtain a college degree. If such college quality effects are large for students affected by race-based admissions policies, then affirmative action policies could raise collegiate attainment among under-represented minority students.

5.3.1 Graduation and College Quality

We begin this section with an overview of the literature on the effect of college quality on BA completion. Table 2 shows six-year graduation rates across different higher education sectors in the United States for the 2011-2012 school year. We split the four-year higher education sector into four mutually exclusive and exhaustive groups: flagship public universities,⁴⁹ non-flagship public universities, highly-selective private schools, and less-selective private schools. The highly selective private schools are the top-65 private universities and colleges according to the 2012 US News and World Report rankings as well as the top-50 liberal arts schools. We have chosen to categorize the higher education market in this way because these sectors represent the broad choices students face when deciding what type of college to attend. The differences across the tiers are substantial: public flagships and elite privates have six-year graduation rates of 70.2% and 86.8%, respectively, compared to between 44-48% in the other two sectors. These differences are particularly pronounced for Hispanic and black graduation rates, which highlights the potential importance of school quality for these groups. In addition, the table shows resource differences across sectors in terms of student-faculty ratios, per-student expenditures and faculty salaries. For all measures, the highly-selective private schools have much higher resource levels than any other sector, with public flagships also having substantially more resources than lower-ranked public and private schools. Thus, affirmative action bans that re-sort minority students from flagship universities to non-flagship public schools will in effect reduce the amount of postsecondary resources to which these students have access, which in turn could lead to worse postsecondary outcomes.

The graduation rate differences across sectors shown in Table 2 also reflect factors that make it difficult to identify the causal effect of college quality on graduation. Due to the nature of how students select into different quality tiers as a function of pre-collegiate academic training, the

⁴⁹The flagship public schools are straightforward to identify. See Lovenheim and Reynolds (2013) for a description of this classification.

most academically qualified students sort into the highest-quality schools. The outcome differences across school types thus could reflect prior academic training, and it therefore would be incorrect to extrapolate these graduation rates to students who were admitted under an affirmative action policy. Additionally, it is possible that collegiate resources have heterogeneous effects on students with different pre-collegiate academic abilities, which forms the basis of the mismatch hypothesis.

The literature on the effect of college quality on graduation rates attempts to control for differences in the pre-existing academic capabilities of students in order to isolate the role of college quality.⁵⁰ Although not without its difficulties in terms of causal identification due to the college selection process, prior work shows consistent evidence that college quality has a large and positive effect on the likelihood a student graduates from a four-year school on average.⁵¹

Bound, Lovenheim and Turner (2010) show evidence from the National Longitudinal Survey of the High School Class of 1972 (NLS72) and the National Educational Longitudinal Study of 1988 (NELS:88) that once one controls for an extensive set of student background characteristics, including high school test scores, large completion differences across college quality tiers still remain. The NLS72 and NELS:88 datasets are nationally-representative and follow the high school classes of 1972 and 1992, respectively, as they transition to college and then into the workforce. Their categorization of schools follows closely with the one presented in Table 2, and their estimates indicate that students beginning college at a top-50 public school or at an elite private school are significantly more likely to obtain a BA within 8 years than are students who begin college at less-selective four-year public schools. Furthermore, this gap grew substantially across cohorts,

⁵⁰In this literature, the terms “college quality,” “college resources” and “college selectivity” are used interchangeably. We follow that convention here.

⁵¹Prior research also has addressed the question of whether starting college in a two-year school versus in a four-year school affects the likelihood of BA receipt. Rouse (1995) shows evidence from community college openings that some students are diverted from a four-year school and are less likely to complete a BA as a result. Using matching methods to help account for the non-random selection of students across school sectors and the non-overlap in observable characteristics between two- and four-year students, Long and Kurlaender (2009), Reynolds (2012) and Doyle (2009) all report large negative effects of two-year attendance on eventual BA receipt. Because most affirmative action policies affect admission and enrollment at four-year schools due to the open-enrollment nature of community colleges, we restrict our attention to the literature that focuses on that sector.

suggesting that college quality is becoming more important over time.

Bound and Turner (2007) use a different and novel source of variation in school quality derived from population variation to identify the role of collegiate resources on graduation that is less prone to the biases associated with the selection-on-observables approach. Their insight is that relatively large cohorts of students represent a demand shock for college. Since the supply of spaces at higher-resource schools tends to be inelastic, and since state funding responds to demand increases with long lags, students born in larger cohorts are less likely to attend a selective school and probably experience lower resources when they are there. They term this phenomenon “cohort crowding.” They show that the elasticity of college graduation with respect to cohort birth cohort size is smaller than the elasticity of college enrollment with respect to cohort birth cohort size. This finding is driven by some combination of students in larger cohorts attending less-selective colleges and with per-student resources declining at less-selective schools when there is a positive demand shock. Regardless of which effect dominates, both factors point towards resources mattering for degree attainment.⁵²

Another source of variation researchers have used to inform the question of how college quality affects college graduation rates is state merit aid systems. State merit aid programs are prevalent and growing in the US, and they offer higher-performing students scholarships to attend college in-state. Cohodes and Goodman (2014) study the Massachusetts Adams Scholarship that provides free tuition (but not fees) for high-performing students in Massachusetts. Using a regression discontinuity design surrounding the state test score eligibility cutoffs, they show that becoming eligible for this scholarship reduced college quality by shifting students from higher-quality private schools to lower-quality public universities. They also show BA attainment rates within four years declined by 1.7

⁵²One objection to their approach is that students from larger cohorts may experience crowding in elementary and secondary schools that makes them less academically prepared for college. Bound, Lovenheim and Turner (2010) conduct a similar analysis in which they control for student pre-collegiate test scores. Their findings are very similar, which suggests reduced pre-collegiate training is not driving the cohort crowding result.

percentage points, which is consistent with declining school quality on average worsening college outcomes.⁵³

While the estimates in Cohodes and Goodman (2014) also could be driven by changes in family budgets and the fact that treated students are more likely to go to school closer to home and maintain their high school peer groups, they show evidence that it is the highest-scoring students whose BA attainment is most affected. Since these are the students who are most likely to experience a decline in college quality due to the scholarship, their results are consistent with an important role for college quality in driving college completion. Cohodes and Goodman (2014) also may be picking up matching effects, as students who are qualified to receive the merit aid program may be the most likely to benefit from attending a higher-quality college.

5.3.2 Graduation and Fit

This research on the graduation effects of college quality examines average effects, but average effects may not be relevant for all students if there is significant mismatch. Several papers in the literature have examined the empirical relevance of “mismatch,” particularly along racial/ethnic lines. Loury and Garman (1995) provide one of the first analyses of the effect of mismatch amongst whites and blacks using data from the NLS72. Using selection-on-observables models that attempt to control for selection using SAT scores, parental income and education, and student demographic variables, they examine whether black and white students are more or less likely to graduate if they are “better matched” to their college in terms of their own SAT score being closer to the college median. Consistent with Bound, Lovenheim and Turner (2010), they find college quality has a sizable average impact on the likelihood of obtaining a BA for both blacks and whites, but

⁵³The finding that state merit aid reduces BA attainment is far from universal in this literature. Fitzpatrick and Jones (2013) and Sjoquist and Winters (2012) find that the implementation of a state merit aid program does not affect graduation rates, while Scott-Clayton (2011) shows positive effects of state merit aid on collegiate attainment. A major difference between these studies is in the quality of schools students would attend in the absence of merit aid. These differences likely drive much of the heterogeneity in findings in this literature, although no work has directly tested this hypothesis.

lower-SAT whites and blacks both are less likely to graduate when they attend a higher-SAT school. This penalty is particularly large for black students. The authors interpret these results as evidence of mismatch, as less academically prepared students are less likely to get a four year degree when they attend a higher resource school, despite the higher level of resources they experience there.

Light and Strayer (2000, 2002) provide some additional evidence on the extent of mismatch using data from the NLSY79. They model college graduation as a two-stage process: in the first stage, students decide whether to go to college and which college type to attend, and in the second stage they make a graduation decision. They estimate these two stages simultaneously and allow for correlations of the error terms across equations. This method explicitly allows for unobservables to influence both enrollment and completion. The results from Light and Strayer (2000) show that the predicted likelihood of graduation among the lowest-quartile AFQT students declines monotonically as college quality increases. As student academic ability rises, this pattern largely reverses.⁵⁴ These results again suggest that policies that induce less-academically prepared students to attend higher-quality schools might be counter-productive in terms of graduation. In a related paper using similar methods, Light and Strayer (2002) find similar patterns. Here, minority status has a positive effect on graduation rates, a positive intercept regardless on the what college the student attends. Light and Strayer (2002) argue that this positive effect may be result of affirmative action programs being successful. However, this ignores the effect of minority status on college quality, the primary channel through which racial preferences affect outcomes. And here, the results indicate match effects are important and generally dominate quality effects. That the minority intercept is positive may instead reflect higher returns to a college degree for African Americans.⁵⁵

⁵⁴The reversal is not perfect. A surprising finding of Light and Strayer (2000, 2002) is that schools in the second lowest quartile are associated with the highest graduation rates, though the standard errors are large.

⁵⁵Both Arcidiacono (2005) and Arcidiacono, Bayer, and Hizmo (2010) find that African Americans receive a larger wage bump from college graduation than whites.

5.3.3 Evidence from Affirmative Action Bans

The recent affirmative action bans make it possible to directly test how affirmative action affects graduation rates. Hinrichs (2012) examines the effect of affirmative action bans for public postsecondary schools that were enacted in California (1998), Florida (2001), Texas (1997) and Washington (1999). Using American Community Survey data from 2005-2007 that links each respondent to whether an affirmative action ban was in place in his state of birth when he was 18, Hinrichs (2012) estimates difference-in-difference models that examine how ban enactments change the BA attainment rate of minority students. He finds no effect: his preferred specification produces a coefficient on the ban that is small in absolute value and is not statistically significant. His estimates are precise enough that he can rule out effects larger than 0.8 of a percentage point (or 7% relative to the sample mean) decline in BA attainment rates. In followup papers, Backes (2012) and Hinrichs (forthcoming) examine the question of how affirmative action bans affect completion using institutional data from IPEDS.⁵⁶ Backes (2012) examines two additional states, Georgia (2002)⁵⁷ and Michigan (2004), focusing on six-year graduation rates at public schools. He finds a significant negative graduation effect on the share of public school graduates who are black, but the magnitude is very small.⁵⁸ Hinrichs (forthcoming) examines both four and six-year graduation rates, where the dependent variable is the share of blacks (or Hispanics) who graduated from a particular type of institution. Considering all institutions, no negative effects were found for either blacks or Hispanics. Backes (2012) and Hinrichs (forthcoming) also look at the share of graduates of particular races at selective colleges. Consistent with the findings on enrollment, the share of graduates who are under-represented minorities at selective colleges falls. These estimates thus point to little impact

⁵⁶A potential issue with using IPEDS for this question is that some minority students may react to a ban by attending college out of state.

⁵⁷There is some ambiguity about how to treat Georgia's affirmative action policy change, as only the University of Georgia eliminated racial admissions preferences.

⁵⁸Note that by focusing on public schools a possible substitution from public to private schools that can still use affirmative action is ignored.

of affirmative action bans on minority graduation rates.

Cortes (2010) analyzes the effect of the Texas affirmative action ban using data from six public universities of differing quality in Texas. In particular, her data contain the two flagship universities in Texas as well as a set of less-selective public schools and two private schools. The identification strategy employed is to compare changes in six-year graduation rates among those in the top decile of their high school class to changes among lower-ranked students in a difference-in-difference framework. The argument for this methodology is based on the fact that the Texas Top 10% rule guaranteed admissions to any public Texas university for students in the top 10% of her high school class,⁵⁹ and students in the top 10% were highly likely to be admitted to Texas A&M or UT-Austin prior to the affirmative action ban. The results indicate that, after the ban, minority students in Texas were between 2.7 and 4.0 percentage points less likely to earn a BA within six years. However, these effects are due predominantly to an increase in graduation rates among those in the top decile (who are assumed to be untreated by the ban) rather than a decline among lower-ranked students. This pattern is more consistent with the top 10% rule positively impacting top-decile students than with the affirmative action ban driving a decline in completion rates for lower-ranked students.

Arcidiacono, Aucejo, Coate, and Hotz (2014) examine how the affirmative action ban in California affected five-year graduation rates, paying particular attention to issues of the match between the student and the school. They show graduation rates for under-represented minorities in the UC system increased by over four percentage points after the ban. They then decompose this change into three parts: improvement in the fit between the student and the school, selection into the UC system as minorities were a smaller share of the student body after the ban, and gains due to universities responding to the ban by changing how academic preparation translates into graduation. Their model allows the returns to academic preparation to vary across campuses, consistent with

⁵⁹We discuss the Texas top 10% rule as well as other related “percent plans” as a replacement for affirmative action in Section 6.

Figure 1. The results show patterns consistent with the bottom panel of Figure 1: the most selective schools have the highest graduation rates for the most prepared students, while the least selective schools in the system had the highest graduation rates for the least prepared students. Reshuffling minority students within the system according to their sorting patterns after the ban results in graduation gains for minority students between 0.1 and 1.2 percentage points, depending on the controls for selection.⁶⁰ Substantially higher gains are found for those in the bottom quartile of the minority preparation distribution, where the resulting improvements in fit improve graduation rates by between 0.8 and 2.2 percentage points. At the same time, these students are also the ones most likely not to be admitted to any school in the system in the presence of a ban.

The results of Arcidiacono, Aucejo, Coate, and Hotz (2014) also suggest that universities may respond to affirmative action bans by focusing more resources on ensuring that the minority students they do have actually graduate. Their lower bound estimates suggest that, conditional on both observed and unobserved ability, students had significantly higher probabilities of graduating after the ban. Antonovics and Backes (2014) show further evidence of endogenous institutional responses to affirmative action bans. In the UC system, one of the responses to the ban was to adjust the admissions weights placed on different background characteristics in diversity-enhancing ways. Universities also may adjust how they recruit students and how they invest in students conditional on arrival, as occurred in Texas with the Century and Longhorn scholarship programs that were implemented post-*Hopwood*.⁶¹ Such changes may provide additional support for minority students at more elite schools after affirmative action bans are implemented.

Fit seems to be more important as the graduation window shortens. Arcidiacono, Aucejo, and Hotz (2013) show that minority four-year graduation rates in the UC system are highest at UC

⁶⁰The smallest results are for the baseline case that only controls for selection on observables. Larger results are found when controls similar to those used in Dale and Krueger (2002) are used. See Section 5.5 for a discussion of Dale and Krueger (2002).

⁶¹See Andrews, Ranchhod and Sathy (2010) for a description of these programs and their effect on minority enrollment.

Santa Cruz despite minorities at UC Santa Cruz having substantially worse academic backgrounds than those at places like UC Berkeley and UCLA. This pattern does not show up in the raw data for five-year graduation rates, suggesting that match effects may affect time to degree as well as completion rates. Arcidiacono and Koedel (2014) examine eight-year graduation rates in the Missouri system. While there is some evidence that fit matters when comparing the most-selective schools in the system to moderately selective schools, there are a set of schools at the bottom of the quality distribution that appear to be worse for all students. This may be the result of quality effects dominating any match effects at the bottom of the school quality distribution or because of the wide graduation rate window considered in this study.

In sum, the evidence suggests that 1) increasing college quality positively affects graduation rates on average, 2) these positive effects turn negative for those with relatively low levels of academic preparation, and 3) affirmative action bans have had little effect on the overall graduation rates of minority students. These results are consistent with offsetting influences of quality effects and match effects. Given that more selective schools have an overall positive effect on graduation, racial preferences that are milder than those currently in place, at least at the very best schools, could result in higher graduation rates for minority students, taking advantage of the higher resources associated with more selective schools without being so aggressive such that match effects degrade educational outcomes. However, the estimated graduation gains from milder preferences at the best schools are small. As Arcidiacono and Koedel (2014) note, substantially higher gains could be achieved from policies that lead to college quality increases among minority students attending the very bottom schools.

5.4 Choice of major

Many studies have documented substantial differences in the monetary returns to different majors.⁶² Relatively few studies have analyzed how the fit between the student and the school affects one's choice of major, however, but it is here where fit may be the most important. This is in part due to the very different demands placed on students in different majors. Majors in STEM (science, technology, engineering, and math) as well as in economics have very different grading distributions (Sabot and Wakeman-Linn 1991, Johnson 2003) and are associated with more study time (Brint, Cantwell, and Saxena 2012, Stinebrickner and Stinebrickner 2011).

Differences in grading standards and study time have disproportionate effects on those who are at the bottom of the preparation distribution. Within colleges, there is massive sorting surrounding who persists in majors in the sciences, with those who have higher SAT scores (and in particular SAT math scores) being much more likely to persist in the sciences (Arcidiacono 2004, Arcidiacono, Aucejo, and Hotz 2013).

Since affirmative action results in minority students being relatively less-prepared than their non-minority counterparts within a given school, differences in expectations across fields can result in minorities shifting away from the sciences over the course of their college career. Using data on Duke students, Arcidiacono, Aucejo, and Spenner (2012) showed that, conditional on gender, African Americans were more likely than white students to have an initial major in the sciences. However, their probability of finishing a major in the sciences was much lower. For example, of those who expressed an initial interest in the sciences, 54% of African American males finished in the humanities or social sciences compared to 8% of white males. Arcidiacono, Aucejo, and Spenner show that this entire racial gap can be accounted for by controlling for academic background: there are no differences in major switching patterns across races once one accounts for differences in

⁶²See Altonji, Blom, and Meghir (2012) for a recent review.

academic background.

What the Arcidiacono, Aucejo, and Spenner study cannot do is say whether the African American students at Duke would have been more likely to persist in science majors had they attended a less-selective school. That is, is it relative preparation or absolute preparation that matters for obtaining a science degree? Smyth and McArdle (2004) use the College and Beyond data⁶³ and estimate models of graduating in a science field, controlling both for the difference between the individual's SAT math score and the average SAT score of their school as well as the average SAT score of the school itself.⁶⁴ Smyth and McArdle (2004) find support for relative preparation mattering for attaining a STEM degree: the total effect of college quality (the direct effect and the effect through relative preparation) on STEM attainment is negative.

Arcidiacono, Aucejo, and Hotz (2013) also explore the issue of institutional fit but allow sufficient flexibility such that better students may see higher probabilities of graduating in the sciences at better schools, consistent with the bottom panel of Figure 1. They use data from the cohorts of students applying to the University of California system between 1995 and 1997. These data show substantial within-school gaps in SAT scores between minority students who start in the sciences and finish in the sciences and minority students who start in the sciences but finish outside of the sciences, gaps that are much smaller for non-minority students.⁶⁵ For example, at UC Berkeley the gap is 105 points for minority students and only 21 points for non-minority students. The probability of graduating with a science degree conditional on starting in the sciences at UC Berkeley is 61% for non-minority students, but just over half that (31%) for minority students.

⁶³This is the same data used in Bowen and Bok's seminal book *The Shape of the River* and in Dale and Krueger (2002).

⁶⁴Smyth and McArdle's study was motivated by Elliot et al. (1996), who examine data from eleven colleges of varying selectivity. At each of the schools in their data, individuals who graduate are divided by their tercile of the school's SAT math score distribution. They find that the share of science degrees awarded to each tercile was very similar across institutions despite large differences in the average SAT math scores of the schools, with the top tercile generally producing over half the science graduates. While more selective schools also produced more science graduates overall, the net effect is still that attending a less-selective institution makes persisting in the sciences more likely. A caveat with this study is that it focuses on college graduates, and dropout rates may differ across the institutions.

⁶⁵Finishing in the sciences and graduation probabilities both refer to five-year graduation rates.

Graduating in the sciences with a low SAT score is difficult at UC Berkeley, but would these same students have a higher probability of graduating in the sciences from a school like UC Riverside? The answer appears to be yes. Denote the median SAT score for minorities at Berkeley as A . The minority students below this score who are interested in the sciences have a five-year graduation rate in the sciences of 13%. Now consider minority students at Riverside, again with an initial interest in the sciences, who have SAT scores below A .⁶⁶ Their five-year graduation rate in the sciences is 20%, despite likely having lower SAT scores overall and likely having a worse academic record.⁶⁷

Arcidiacono, Aucejo and Hotz estimate a model of major choice, controlling for some observed preparation measures as well as Dale and Krueger (2002) controls, where schools differ in their major-specific returns to academic preparation. Hence, as in Figure 1, it is possible for the most-prepared students to be best matched at the most-selective schools and less-prepared students to be best matched at less-selective schools. Estimates of the model reveal exactly this pattern. After accounting for selection, the vast majority of minority students, who are coming in with significantly worse preparation than their non-minority counterparts, would have higher graduation probabilities in the sciences had they attended a less-selective school in the UC system. On the other hand, non-minority students are generally well-matched: a much smaller share would see higher graduation probabilities in the sciences had they attended a less-selective school.⁶⁸

Note that the studies discussed above have focused on selective schools. Arcidiacono (2004, 2005) both analyze the choice of college and major using the NLS72 data set and come to a different conclusion on how college quality and the choice of major interact. Namely, increasing college quality makes majoring in the sciences more attractive. There are two reasons why we do not put as much weight on these studies in terms of how college quality interacts with major choice. First, they placed

⁶⁶87% of entering minority science majors at Riverside have SAT scores below A .

⁶⁷Note that the overall graduation rate for this group is higher at Berkeley. Berkeley does an excellent job at graduating minority students, just not in the sciences.

⁶⁸Indeed, Arcidiacono, Aucejo, and Hotz find that non-minority students in the UC system would have had lower graduation rates in the sciences had they been assigned to schools like their minority counterparts.

considerable structure on how college quality and major choice interacted by constraining college quality to make particular majors more or less attractive for *all* students as opposed to allowing for the complementarities discussed in Section 2. As the discussion of Arcidiacono, Aucejo, and Hotz (2013) suggests, mismatch only becomes an issue when one’s preparation is severely behind one’s peers.⁶⁹ Second, there are no controls for intended major, and those who intend to major in the sciences are on average stronger students and therefore more likely to attend higher-quality colleges.

But it is possible that the difference in results are is partly driven by the national data used in Arcidiacono (2004, 2005). Arcidiacono and Koedel (2014) find that the lowest-quality four-year institutions in the Missouri system have very few if any STEM majors. It may be the case that colleges quality effects dominate at the bottom of the distribution, whereby better schools are better for everyone, while matching effects are more important at the top of the distribution. The affirmative action debate has typically focused on racial preferences at the very best institutions, but preferences may be most beneficial further down the college quality distribution.

5.5 Labor Market Outcomes

Ultimately, when assessing the impact of affirmative action on minority students, we care about how their long-run life outcomes are affected. An important component of such outcomes is labor market performance, such as earnings, wages, and the likelihood of employment. Affirmative action will potentially affect labor market outcomes through its influence on the myriad of collegiate outcomes considered above: quality of schools attended, the likelihood of graduation, and choice of major. Affirmative action policies also could have direct effects on labor market outcomes that work through employer beliefs. As shown by Coate and Loury (1993), affirmative action can lead to employers developing “negative stereotypes” about minority workers that lower their returns to

⁶⁹Note then that it is not just a matter of being a small fish in a big pond as Gladwell (2013) suggests, but how well one is prepared for the material being taught.

a given educational investment. In equilibrium, employers will believe minority workers to be less productive, will pay them less, and thus will lower the incentive for these workers to invest in labor market skills (including education). If employers develop beliefs about worker productivity that are a function of school attended and race, then such an equilibrium could arise due to affirmative action in the higher education market. Indeed, the concern that employers will treat under-represented minority workers from elite schools differently because of affirmative action is a core argument opponents of affirmative action adduce against these policies. Furthermore, affirmative action could have important distributional consequences for minority workers, with affirmative action bans potentially influencing wages among very high-skilled minorities.

The results from affirmative action bans discussed in Section 5.2. highlight that a primary effect of affirmative action on minority students is to induce higher rates of attendance at more elite schools. With little effect on overall graduation rates, the effect of affirmative action in admissions on earnings will flow predominantly through the choice of major and college quality.

We begin our discussion by considering the returns to college quality in general. This is a growing area of research that has made many recent advances due to growing data quality that allow one to link students' K-12 education records to higher education records and future earnings. As in the estimation of college quality effects on higher BA attainment rates, the struggle in all studies on the returns to college quality is to deal with the fact that students with higher pre-collegiate academic ability that is independently valued in the labor market are more likely to sort into higher-quality schools. Put differently, student sorting across the quality distribution of the higher education market is highly non-random, and even when the characteristics of students at different types of schools overlap we must still be concerned about why they wound up in different institutions and whether these differences are correlated with unobserved attributes that will drive labor market outcomes. The literature to date has used three methods to address this endogeneity

concern: selection on observables, regression discontinuity using admissions rules, and matching based on students with similar application and admissions histories. We discuss each of these in turn below, and we also point interested readers to the thoughtful review provided in Hoxby (2009).

The largest group of studies uses student observable characteristics to control for the differences across students in underlying earnings potential across the higher education quality distribution. In early work on this question, Brewer, Eide and Ehrenberg (1999) use the rich demographic and pre-collegiate test score data in the NLS72 and High School and Beyond (HSB) data sets to estimate wage and earnings premiums for those who attend schools from different higher education sectors.⁷⁰ Their results indicate that students attending a high-quality public university (defined as being most or highly competitive in the Barron's ratings) earn as much as 25% more than those attending a low-quality one (defined as being less- or non-competitive). In addition to the concern that there is selection on unobservable characteristics, such as student motivation, there is a concern that OLS models will produce misleading estimates because of limited overlap in the observable characteristics of students attending schools of differing quality.

If there was full sorting on background characteristics, then the only way to adjust for differences in background characteristics across schools would be to impose a potentially restrictive functional form assumption. This issue is taken up directly in Black and Smith (2004). They use the National Longitudinal Study of 1979, which similar to NLS72 and HSB, contains detailed background and pre-collegiate cognitive skill measures, and they estimate matching models that allow them to both assess the degree of common support across the distribution of college quality and to examine how sensitive the results are to excluding the non-overlapping set.⁷¹ Their results demonstrate a surprising amount of overlap in the propensity score distributions of the likelihood of being in

⁷⁰Note that the student need not graduate from college to be included in the sample.

⁷¹Black and Smith (2004) measure college quality using an index based on average faculty salaries in 1997, average SAT score of the entering class in 1990 and the average first-year retention rate in 1990. Black and Smith (2006) use a more extensive set of institutional characteristics to measure college quality, again showing significant college quality effects.

the top versus the bottom quartile of college quality as a function of the observables. They find a sizable effect of college quality: going from the bottom to the top quartile of college quality increases wages by 14% for men and 8% for women.⁷² Although these estimates are difficult to compare with Brewer, Eide and Ehrenberg (1999) due to differences in the way they measure school quality and not separating public and private schools, they both indicate substantial returns to attending a higher-quality college.

Arcidiacono (2005) is the only study that links racial preferences in admissions to future earnings. Using data from the NLS72, he estimates a structural model where students first decide the set of colleges to which they will apply given expectations about their probabilities of admission, financial aid conditional on admission, enrollment and major choice decisions, and future earnings. Arcidiacono's results on earnings show small returns to college quality as measured by the average SAT math score of the student body, particularly after allowing for the errors in the various parts of the model (applications, admissions, earnings, etc.) to be correlated.⁷³ However, as pointed out by Black and Smith (2004, 2006), using only one noisy measure of college quality may bias downward the estimated returns.

Arcidiacono's model is estimated in a regime where racial preferences are present. The structural parameter estimates are then used to forecast how black student decisions over applications, enrollment, and major choice would change if they faced the admission and/or financial aid rules of whites. The effects of removing racial preferences in admissions on expected black earnings, even among those most likely to attend a selective college, were negative but quite small. This is due to racial preferences in admissions having little effect on enrollment and the low estimated returns to college quality.⁷⁴

⁷²These are similar to the OLS estimates, and for both men and women the effects grow substantially, to 25% and 16%, when the estimation sample is restricted to the "thick support region," where students are most similar in terms of observables across school types.

⁷³Similar low returns, albeit with the same data set, are found in Arcidiacono (2004).

⁷⁴Arcidiacono (2005) finds slightly larger negative effects from removing racial preferences in financial aid as racial preferences in financial aid affect whether black students attend college at all. He estimates that the returns to college

These results are based on older cohorts who graduated from college in the 1970s and 1980s. It thus is important to consider some more recent evidence. Long (2010b) shows suggestive evidence that the returns to college quality has increased over time. He estimates the effect of college quality in the NLS72, HSB and NELS:88 surveys, and he shows that the effect on earnings of a one standard deviation increase in quality has grown from 2.6% to 4.8% across cohorts. However, this change was not statistically different from zero. Interestingly, the return to college quality for African American students in NELS:88 dropped to zero from the 7.5% estimate in HSB. Thus, the increasing returns do not appear to be uniformly distributed across students, and if the returns to quality have dropped for black students this has important implications for how affirmative action changes have affected the returns to education for this group.

Evidence from the most recent cohorts comes from Texas, where it is possible to link students' K-12 records to higher education outcomes and labor market earnings starting from the 1994 high school graduating cohort (Andrews, Li and Lovenheim, forthcoming). Andrews, Li and Lovenheim (forthcoming) is distinguished by the large volume of student high school test score controls that the longitudinal datasets used in other work do not have and by the large sample sizes that provide more statistical power. They estimate OLS models of the effect of UT-Austin and Texas A&M-College Station graduation on earnings relative to earnings from graduates at other four-year schools in Texas. These estimates are similar to those from prior work, suggesting an earnings premium from UT-Austin of 11.5% and from Texas A&M of 21%. For black and Hispanic students, however, the returns to UT-Austin are small, while the returns to Texas A&M are about 20%. Furthermore, they estimate quantile treatment effects that show how college quality shifts the earnings distribution. At UT-Austin, there is a wide distribution, with earnings at the top of the distribution shifting out much more than at the bottom. The effect of Texas A&M on the earnings distribution is much

are significantly higher for blacks than whites given the same background characteristics. Racial preferences in financial aid have the added advantage that mismatch effects are unlikely to result.

more constant. These results may be consistent with the bottom panel of Figure 1 if the quantiles correlate strongly with academic preparation, which would explain the lower returns at UT-Austin for blacks and Hispanics. Alternatively, the wide distribution at UT-Austin may simply reflect UT-Austin being a risky proposition regardless of where one lies in the preparation distribution.

Overall, the evidence from studies that rely on student observable characteristics to account for selection into schools of differing quality points towards significant effects of college quality on earnings. However, there is much concern that despite the large volume of information contained in the conditioning sets used, there still is residual selection on unobserved characteristics that could exert an independent influence on labor market outcomes. Hoekstra (2009) takes a unique approach to solving this identification problem by using an admission cutoff rule at a large flagship university based on SAT scores and GPA. Students are assigned a score based on a combination of SAT and GPA, and admission rules use a cutoff for each SAT-GPA combination that form the basis of a regression discontinuity (RD) design. He shows that students attending the flagship because they score just above the cutoff earn 25% more than those who do not attend because they are just below. A complication with this study is that he does not observe what happened to students who do not attend the flagship, so the counterfactual is difficult to assess. However, in terms of affirmative action, this paper is relevant for two reasons. First, the effect is focused on the students at the margin of being academically eligible for admission at a flagship university, and under-represented minority students who are impacted by affirmative action rules are more likely to be in this group. Second, the correct counterfactual when thinking about imposing stronger or weaker admissions requirements is what students would do if not admitted. Thus, Hoekstra's results are applicable when attempting to predict the policy effect on earnings of imposing higher admissions requirements for a marginally-admitted group.⁷⁵ However, an important caveat to his findings with respect to

⁷⁵Zimmerman (forthcoming) shows a similarly-sized effect on earnings using an admission cutoff rule for a low-selectivity school in Florida. This finding suggests that earnings returns to college are not limited to the high-quality schools, but we focus this discussion on the flagship and selective schools because these are the schools whose minority

affirmative action is that he focuses only on white students; it is unclear whether one can generalize his findings to other racial and ethnic groups.

The final method researchers have used to identify the effect of college quality on earnings is through a matching technique that compares earnings across students who attend different-quality universities but who applied to and were admitted to the same set of schools. Dale and Krueger (2002) were the first to implement this method using College and Beyond and NELS:88 data, and the motivation is that the student unobservable qualities one might be concerned about in the selection-on-observable studies are contained in the application and admission set. They show that once one conditions on the application and admissions sets, there is no evidence of an earnings return to college quality. This finding is re-confirmed in a follow-up study (Dale and Krueger, 2014), however both papers find sizable college quality effects for students from low-income families. To the extent under-represented minority students are more likely to be from lower-income families, these results suggest college quality may still be an important component of future earnings for this group. This methodology has several appealing attributes, most notably that it is comparing students who are very similar in terms of the types of schools they express a desire to attend and the types of schools to which they are academically eligible to attend.⁷⁶ But, this method also has generated controversy because of the concern over what is driving the residual selection. That is, why do two different students who have identical choice sets choose schools of different quality? Some of the potential explanations could be the types of unobservables this model is designed to overcome (e.g., differences in student motivation). Or, there could be differences in match quality that are highly important in this context, driven by differences in schools' focus in terms of academic area or other attributes students value. Finally, the differences in college quality within matched sets, particularly at the top of the quality distribution, tend to be small. It could be the case that the effects of college

enrollments are most affected by affirmative action.

⁷⁶Indeed, Dale and Krueger (2014) shows positive returns for minorities for the most recent cohort. However, similar positive findings are not found for the older cohort.

quality are non-linear in means, which could complicate comparisons across very similar schools.

Our conclusion from this literature is that, on the whole, the evidence supports the claim that there are significant returns to college quality. Even Dale and Krueger (2002) show sizable effects for low-income students. The estimates are particularly large when examining the elite public schools, which typically are state flagship universities. Table 2 shows that these schools have substantially higher resources and higher graduation rates, both overall and for minority students. It is reasonable to conclude that these higher resources allow these schools to have a larger effect on human capital than lower-resource schools. A prediction that stems directly from this evidence and from the enrollment effects discussed in Section 5.2 is that affirmative action likely raises the earnings of minority students due to its importance in driving minority enrollment in selective universities and the substantial earnings returns that tend to follow.

At the same time, many of the results in this literature do not speak that closely to the question of how admitting lower-ability minority applicants through affirmative action affects their future earnings. More broadly, this literature has not explored how the match between the school and the student affects the returns to college quality. This is an important hole in the returns to college quality research and is thus a ripe area for future research. The methods used by some of the more prominent papers in this literature, namely Dale and Krueger (2002; 2014) and Hoekstra (2009), are particularly ill-suited for this task. In the case of Dale and Krueger, due to the noisiness of college quality measures and the fact the students' first and second choices do not differ in terms of measured quality by that much, any mismatch differences across first and second choice schools are likely to be small. With respect to Hoekstra (2009), his estimates are identified off of academically-marginal students whose outside options are worse in terms of quality. While minority students are more likely to be academically marginal, the existence of affirmative action policies makes their outside options far better than an observationally-equivalent white student.⁷⁷ These studies also generally

⁷⁷Hoekstra focuses only on white students in his analysis, so he cannot test this hypothesis.

focus on earnings not conditional on graduation.⁷⁸ Since college quality overall has a positive effect on graduation rates, this may be the channel through which it influences earnings. But as we saw in Section 5.3.1, higher graduation probabilities from attending a higher quality college may not result if the student is significantly underprepared relative to his peers.

5.6 Summary of the Evidence

This section reviewed the large literature on the varied returns to college quality as well as the evidence on how affirmative action bans affect student outcomes. Given the quantity of research discussed, we briefly summarize the lessons learned from this research and highlight some remaining research questions. First, we find the research rather clear in showing that, on average, there is a positive return to college quality in terms of the likelihood of graduation as well as of subsequent earnings. The identification strategies supporting these conclusions are all imperfect, but sizable effects have been found using a variety of methods, time periods, and data sources.

The literature is much less clear about how the match between the school and students affects these outcomes more generally and how affirmative action impacts these outcomes specifically. The research using affirmative action bans is certainly the most straightforward to interpret, and this work shows no effect of these bans on the likelihood of graduation. This is a surprising result, because these studies demonstrate that affirmative action bans significantly reduce college quality for African American and Hispanic students. These results can be reconciled by the changes in matching effects from the bans canceling out the college quality effects, though more evidence is needed. There is also a lack of understanding of how universities might respond to changes in affirmative action policies. If universities shift resources to under-represented minority students in order to compensate for externally-imposed admission constraints, this could have sizable impacts

⁷⁸An exception is Andrews, Li and Lovenheim (forthcoming), which examines college quality earnings premiums among graduates.

on student outcomes. No research exists that examines institutional responses to affirmative action bans, which is a notable gap in the literature. In addition, there is no work estimating how affirmative action bans influence earnings. Understanding how these bans influence earnings may yield some insight into how affirmative action policies affect long-run outcomes, and this line of inquiry thus is ripe for future study.

As noted above, graduation is an imperfect measure of how affirmative action impacts students because they can adjust to along other margins. Changing majors is a one of the central ways in which students can compensate for for a poor academic match. The evidences suggests that relatively less-prepared students are likely to switch into less-demanding majors. Furthermore, if these students had gone to slightly lower-quality schools, they would have had a better chance of completing the more difficult major. The central question that arises from this research is what are the returns to majors versus school quality, and how do the returns to major and school combinations differ across different student types. It is extremely difficult to identify the returns to college majors and to major-school choice combinations.⁷⁹ Nonetheless, Gemici and Wiswall (2014) suggest that choice of major has become more important over time and this is an important are of future research.

6 Percent Plans as an Alternative to Affirmative Action

In the wake of court decisions and voter referendums eliminating affirmative action in higher education (see Table 1), states have been increasingly turning to “percent plans” in order to achieve diversity goals.⁸⁰ Percent plans guarantee admission to a set of schools to students who are in

⁷⁹This is not the case in other countries where admissions to school-major combinations are sometimes a reflection of only test score imply regression discontinuity methods may be employed. See, for example, Hastings, Neilson, and Zimmerman (2014) and Kirkeboen, Leuven, and Mogstad (2014). The usefulness of discontinuities is limited to local effects when matching effects are present, though.

⁸⁰This section examines percent plans as a method for states to achieve diversity goals in lieu of race-based affirmative action. If the goal is instead to better the schooling outcomes of minority students, they have a much larger set of policies with which to work, including financial aid, income supports for low-income families, college counseling, information dissemination, and investments in minority students’ education at younger ages. There are large literatures examining the effects of each of these policies that are beyond the scope of this review.

the top X percent of their high school class. The idea behind such rules is that high schools tend to be ethnically and racially segregated, and so a percent plan will necessarily expand access to low-income minority students who are in heavily non-white schools. The most prominent percent plan is in Texas. The “Texas Top Ten Percent Rule” went in to effect in 1998, and it guaranteed all students in the top 10% of their high school class automatic admission to *any* public college or university in Texas. California also has an admission policy in which any student in the top 9% of their class or in the top 9% of students overall in the state are guaranteed admission to *some* University of California campus.⁸¹

In this section, we provide a brief overview of the evidence on how these percent plans affect student outcomes in terms of enrollment and graduation, and we discuss evidence on some of the unintended consequences of these rules that policymakers need to consider when deciding on whether to replace existing affirmative action policies with a percent plan. Most of the evidence comes from studying the Texas Top Ten Percent Rule, so we focus our discussion there. Kain, O’Brien and Jargowsky (2005) examine eight cohorts of college students in Texas that span the pre-ban period, the two years the ban was in effect, and the Top Ten Percent Rule implementation period. They show that the enrollment declines among African Americans at UT-Austin and Texas A&M-College Station post-*Hopwood* discussed in Section 5.2 were partially reversed after the Top Ten Percent Rule was implemented. However, black enrollment at the flagship schools still remained 14% below the pre-ban levels. Among Hispanic students, enrollments at UT-Austin fully recovered relative to pre-ban levels and at Texas A&M it remained at about 20% below the pre-ban level. Overall, the evidence suggests that the Texas Top 10% Rule went some way to undoing the declines in minority

⁸¹The percent plan in California originally considered only those in the top 4% of their graduating class or in the top 12.5% of students in the state. These admission rules were changed for the incoming class of 2012. The within-school ranking in California is based on GPA, while the statewide rankings are based on an academic index, which is a weighted average of SAT/ACT score and high school GPA. Florida has a percent plan as well, called the *Talented 20*. Under this plan, which was enacted in 2000, the top 20% of students in each high school class who complete a college preparatory curriculum are guaranteed admission to at least one of the public universities in Florida, but in practice this program does not bind.

enrollment at Texas Flagship Universities after affirmative action was banned, but the recovery was not full, especially for African Americans. The enrollment patterns also show a large increase in white enrollment at these 3 schools, jumping 27% after the implementation of the Top 10% Rule.

Long (2004b) argues that enrollment trends are misleading because they are confounded by growing minority populations within the state. He performs simulations of enrollment probabilities using the NELS:88 survey and shows that percent plans are unlikely to affect minority enrollment for two reasons. First, most minority students in the top decile of their class are admitted to the state flagship anyway. Second, under-represented minority students are not concentrated enough in most areas and they are unlikely to be in the top portion of their high schools for this policy to have large effects on their enrollment patterns.

A source of variation that is increasingly being exploited to identify the effects of percent plans on students is the discontinuity in access to flagship universities that occurs just above the rank cutoff. This methodology is not subject to the criticisms outlined in Long (2004b), but they do have the drawback that the estimates are only local to the specific cutoff. Thus, these studies do not tell us the overall effect of percent plans on the composition of the student body; they only can inform how students respond when given admission guarantees to flagship schools. They are informative about how students respond to admission guarantees, but the estimates are too local to tell us much about the overall impacts of percent plans on the state postsecondary system (or on flagship schools, *per se*).

There are two studies that have used regression discontinuity methods to study the Texas Top Ten Percent Rule, and they come to somewhat different conclusions. Niu and Tienda (2010) use data from the Texas Higher Education Opportunity Project (THEOP), a state-representative, longitudinal survey of Texas public high school students that began in 2002. Class ranks were verified using transcripts among those who persist in the survey to the second follow-up. The main finding

is a 12-14 percentage point increase in the likelihood of Hispanic students who are just over the 10 percent rank cutoff enrolling in one of the two Texas flagships. Students at predominantly minority schools also experience a jump of 14 percentage points in the likelihood of enrolling in a flagship at the cutoff. White students and students in majority white schools are unaffected, as are African American students. This evidence is consistent with the tabulations in Kain, O'Brien and Jargowsky (2005), which shows Hispanic enrollment at Texas flagships increased substantially after the Top Ten Percent rule was implemented relative to enrollment in the ban years.

Daugherty, Martorell and McFarlin (2014) examine how the cutoff not only affects whether students attend the flagship school, but also how it affects attendance patterns at other schools. They match six cohorts of students to a large, un-named district in Texas and estimate regression discontinuity models using administrative data on class rank that is linked to college enrollment outcomes both within and outside of Texas. They show that, while being in the top ten percent increases the likelihood students enroll in a state flagship, it displaces enrollment in private and out-of-state universities. As a result, students just above and below the top ten percent cutoff experience the same level of college quality.⁸² Interestingly, they also find no effect on the likelihood of enrollment, even though the baseline enrollment rate in college is only 30% in the study's district (and is only 58% among those in the top decile of class rank). This finding highlights the fact that percent plans influence where students attend, not whether they attend at all, similar to affirmative action policies more generally.

The findings of this paper are relevant for affirmative action policy, as they suggest that expanding admission preferences at any one school (or among a small number of schools) may simply re-sort students among similar-quality institutions. This is perhaps an explanation for the lack of graduation effects found in the studies examining affirmative action bans discussed in Section 5.3.3.

⁸²Daugherty, Martorell and McFarlin (2014) also highlight that they use class rank at the end of 11th grade, which is the relevant rank for college applications, while Niu and Tienda (2010) use the rank at the end of 12th grade. They show evidence that this difference can account for the different findings across the two studies.

Furthermore, this study seemingly conflicts with the findings in Hoekstra (2009). Hoekstra finds large effects of attending the state flagship on earnings using an admission discontinuity. While a full understanding of why the findings differ is difficult (and is a ripe area for further research), the explanation likely is related to the fact that the students in the Daugherty, Martorell and McFarlin (2014) study are not marginal for admission to the state flagship in terms of academic quality. Indeed, it was highly likely they would have been admitted prior to the Top 10% Rule. In the Hoekstra study, the students who are barely admitted are marginal for enrollment by design, so their alternatives are very likely to be lower-quality institutions. This difference highlights the limited applicability of the results in Hoekstra (2009) for affirmative action, as most under-represented minority students admitted under affirmative action will have access to an array of similarly-competitive schools.

In addition to the direct effects on which types of schools students attend, percent plans can have general equilibrium effects that often are not intended by policy makers. A prime example is found in Cullen, Long and Reback (2012), who show that a small but non-trivial proportion of student transfers across high schools after the Top Ten Percent Rule was implemented were to improve their chances of being in the top 10%. Not only does this strategic behavior reduce the high school quality students experience, it also can undo the effectiveness of the percent plan by displacing minority students with non-minority transfer students. Long, Saenz and Tienda (2010) show that the Top Ten Percent Rule expanded access to flagships geographically, with a large drop in the proportion of minority students from “feeder” schools and a higher percentage coming from rural towns and small cities. Thus, the percent plan in Texas altered the composition of minority students, the effects of which on the academic attainment of these students or on their peers at flagship universities currently is not known. In addition, percent plans can alter grading incentives placed on schools and/or specific teachers, can alter the distribution of resources in schools to students of different ability levels, and can affect student effort incentives for those close to the percentile cutoff. The

existence and extent of these effects on students and schools has not been studied previously, and we view these as important questions for future research.

7 Conclusion

Intense debate continues in the United States over the role of affirmative action in college and graduate school admissions. Although the Supreme Court has ruled that the constitutional basis for affirmative action must be rooted in schools' preferences for a diverse student body, the original impetus for these policies and much of the resulting focus among policy-makers and researchers alike is on whether and how affirmative action "levels the playing field" between under-represented minority and white students. Of central importance, then, is how racial preferences in college and graduate school admissions influence students' long-run educational and labor market outcomes, relative to a counterfactual of race-neutral admissions.

From a policy perspective, understanding the extent and effect of affirmative action policies is of primary importance. The literature clearly shows positive average effects of college quality on a host of outcomes. This suggests that mild racial preferences will have a positive impact on minority outcomes. The issue is whether racial preferences in their current form are so strong that mismatch effects may arise. At more selective undergraduate schools and among virtually all law schools, affirmative action is practiced extensively, particularly for African Americans, and leads to substantial differences in the academic backgrounds of minority and majority students. These differences are particularly striking in law schools, where the median black admit may have credentials that are below that of the first percentile of white admits.

As we argue throughout this paper, the net impact of racial preferences come down to whether the strength of the overall college quality effect is larger or smaller than any match effects. If matching effects are dominant given the current levels of racial preferences, then the dual goals of

affirmative action of generating a diverse student body and supporting the educational attainment of minority students may be at odds with one another. Large matching effects would support the need for more policy consideration of what the optimal amount of affirmative action is at different school types in order to balance the negative matching effects with institutional diversity goals. Conversely, absent matching effects, these policy goals can be complementary.

The evidence suggests that racial preferences are so aggressive that reshuffling some African American students to less-selective schools would improve some outcomes due to match effects dominating quality effects. The existing evidence indicates that such match effects may be particularly relevant for first-time bar passage and among undergraduates majoring in STEM fields. However, shifting minority undergraduates to low-resource non-selective schools ultimately may undo any gains from higher match quality, and shifting minorities out of law schools altogether could lead to worse labor market outcomes among these students than had they been admitted to some law school. Alternatively, schools that wish to practice extensive affirmative action could provide targeted services to these students in order to overcome any mismatch induced by their admission policies, such as offering tutoring and remedial classes. While the evidence on targeted college services is scant,⁸³ it is plausible such interventions could be successful in mitigating any negative match effects. The extent to which schools can successfully target services to less-academically-prepared students is an important question for future research.

The past several decades have witnessed a revolution in data availability, with many states as well as the federal government making data available to researchers that permit one to follow students from their K-12 education, through college and into the workforce. These data sets have immense potential to inform most education policy debates, and using these data to study the

⁸³The existing research has suffered from identification issues surrounding student selection into the college services being offered. Abrams and Jernigan (1984) and Chaney et al. (1998) find positive effects of targeted services to lower-performing students. However, neither study can credibly separate student selection into the program from the effects of the program on student outcomes themselves, and neither focuses on students impacted by affirmative action.

effects of affirmative action on long-run outcomes would be quite fruitful. This paper outlines many of the gaps in our current knowledge and the difficulty in credibly disentangling the relevance of fit versus mismatch as they relate to affirmative action. We are hopeful that these longitudinal administrative data sets will allow us to gain new insights into the effects of affirmative action on minority educational and labor market outcomes in the near future.

Administrative data would also permit more exploration of how racial preferences in college admissions affects other margins, such as pre-college investment decisions. If, in equilibrium, affirmative action leads employers to believe minority students graduating from more elite schools are lower-productivity, affirmative action may induce lower minority effort (Coate and Loury, 1993). Conversely, affirmative action bans could lower minority human capital investment by making it less likely they will be admitted to higher-quality colleges that may have higher returns. Caldwell (2010) and Furstenberg (2010) find evidence that Proposition 209 increased the test racial test score gap, consistent with affirmative action raising effort levels among minority students, a result also supported by Hickman (2013).⁸⁴ In contrast, Antonovics and Backes (2014) do not find a widening score gap after Proposition 209 using arguably better data. But in all cases the data is sufficiently poor that, in our view, the question is unsettled. We view this as an important topic for future research.

⁸⁴This is an emerging literature and none of these studies were published at the time of this writing. Some experimental work has also found evidence suggesting affirmative action can raise investment incentives. See for example Schotter and Weight (1992), Calsamiglia, Franke, and Rey-Biel (2013), and Cotton, Hickman, and Price (2014). How these experiments translate into student investment decisions and on what margins student are affected are open questions.

References

- [1] Abrams, Helene G. and Louise Podojil Jernigan (1984). “Academic Support Services and the Success of High-Risk College Students.” *American Education Research Journal* 21(2): 261-274.
- [2] Alon, Sigal, and Marta Tienda (2005). “Assessing the ‘Mismatch’ Hypothesis: Differentials in College Graduation Rates by Institutional Selectivity.” *Sociology of Education* 78(4): 294–315.
- [3] Andrews, Rodney J., Jing Li and Michael F. Lovenheim (2014). “Heterogeneous Paths Through College: Detailed Patterns and Relationships with Graduation and Earnings.” *Economics of Education Review* 42: 93-108.
- [4] Andrews, Rodney J., Jing Li and Michael F. Lovenheim (forthcoming). “Quantile Treatment Effects of College Quality on Earnings.” *Journal of Human Resources*.
- [5] Andrews, Rodney J., Vimal Ranchhod and Viji Sathy (2010). “Estimating the Responsiveness of College Applications to the Likelihood of Acceptance and Financial Assistance: Evidence from Texas.” *Economics of Education Review* 29(1): 104-115.
- [6] Antonovics, Kate and Richard H. Sander (2013). “Affirmative Action Bans and the Chilling Effect.” *American Law and Economics Review* 15(1): 252-99.
- [7] Antonovics, Kate and Ben Backes (2013) “Were Minority Students Discouraged From Applying to University of California Campuses After the Affirmative Action Ban?” *Education Finance and Policy* 8(2): 208-50.
- [8] Antonovics, Kate and Ben Backes (2014a) “The Effect of Banning Affirmative Action on College Admissions Rules and Student Quality.” *Journal of Human Resources* 49(2).

- [9] Antonovics, Kate and Ben Backes (2014b) “The Effect of Banning Affirmative Action on Human Capital Accumulation Prior to College Entry.” *IZA Journal of Labor Economics*, Vol. 3, Article 5.
- [10] Altonji, Joseph G., Erica Blom and Costas Meghir (2012). “Heterogeneity in Human Capital Investments: High School Curriculum, College Major, and Careers.” *Annual Review of Economics* Vol. 4: 185-223.
- [11] Arcidiacono, Peter (2004). “Ability Sorting and the Returns to College Major,” *Journal of Econometrics*, Vol. 121, Nos. 1-2, 343-375.
- [12] Arcidiacono, Peter (2005) “Affirmative Action in Higher Education: How do Admission and Financial Aid Rules Affect Future Earnings?” *Econometrica*, Vol. 73, No. 5, 1477-1524.
- [13] Arcidiacono, Peter, Esteban Aucejo, Patrick Coate, and V. Joseph Hotz (2014) “Affirmative Action and University Fit: Evidence from Proposition 209”, *IZA Journal of Labor Economics*, Vol. 6, Article 3.
- [14] Arcidiacono, Peter, Esteban Aucejo, Hanming Fang, and Kenneth Spenner (2011) “Does Affirmative Action Lead to Mismatch? A New Test and Evidence”, *Quantitative Economics*, Vol. 2, No. 3, 303-333.
- [15] Arcidiacono, Peter, Esteban Aucejo, and V. Joesph Hotz (2013) “University Differences in the Graduation of Minorities in STEM Fields: Evidence from California”, NBER working paper #18799.
- [16] Arcidiacono, Peter, Esteban Aucejo, Andrew Hussey, and Kenneth Spenner (2013) “Racial Segregation Patterns in Selective Universities”, *Journal of Law and Economics*, Vol. 56, 1039-1060.

- [17] Arcidiacono, Peter, Esteban Aucejo, and Kenneth Spenner (2012) “What Happens After Enrollment? An Analysis of the Time Path of Racial Differences in GPA and Major Choice” *IZA Journal of Labor Economics*, Vol. 1, Article 5.
- [18] Arcidiacono, Peter, Pat Bayer, and Aurel Hizmo (2010) “Beyond Signaling and Human Capital: Education and the Revelation of Ability” *American Economic Journal: Applied Economics*, 2(4): 76-104.
- [19] Arcidiacono, Peter, Shakeeb Khan, and Jacob Vigdor (2011) “Representation versus Assimilation: How do Preferences in College Admissions Affect Social Interactions?” *Journal of Public Economics*, Vol. 95, Nos. 1-2, 1-15.
- [20] Arcidiacono, Peter and Cory Koedel (2014) “Race and College Success: Evidence from Missouri” *American Economic Journal: Applied Economics*, 6(3): 20-57
- [21] Arcidiacono, Peter, Michael Lovenheim, and Maria Zhu (2014) “Affirmative Action in Undergraduate Education.” Working Paper.
- [22] Arcidiacono, Peter and Jacob Vigdor (2010) “Does the River Spill Over? Estimating the Economic Returns to Attending a Racially Diverse College” *Economic Inquiry*, Vol. 47, No. 3.
- [23] Ayres, Ian and Richard Brooks (2005) “Does Affirmative Action Reduce the Number of Black Lawyers?” *Stanford Law Review*, Vol. 57, No. 6, 1807-1854.
- [24] Baker, S., Mayer, A., and S.L. Puller (2011) “Do More Diverse Environments Increase the Diversity of Subsequent Interaction? Evidence from Random Dorm Assignment” *Economics Letters*, 110: 110-112.
- [25] Backes, Ben (2012). “Do Affirmative Action Bans Lower Minority College Enrollment and Attainment?” *Journal of Human Resources*, 47(2): 435-455.

- [26] Black, Dan A. and Jeffrey A. Smith. (2004). “How Robust is the Evidence on the Effects of College Quality? Evidence from Matching.” *Journal of Econometrics* 121: 99-124.
- [27] Black, Dan A. and Jeffrey A. Smith. (2006). “Estimating the Returns to College Quality with Multiple Proxies for Quality.” *Journal of Labor Economics* 24(3): 701-728.
- [28] Bound, John and Sarah Turner (2007). “Cohort Crowding: How Resources Affect Collegiate Attainment.” *Journal of Public Economics* 91: 877–899.
- [29] Bound, John, Michael Lovenheim and Sarah Turner. (2010). “Why have College Completion Rates Declined? An Analysis of Changing Student Preparation and Collegiate Resources.” *American Economic Journal: Applied Economics* 2(3): 129–157.
- [30] Bowen, W.; and Bok, D. (1998). *The Shape of the River: Long-Term Consequences of Considering Race in College and University Admissions*. Princeton, NJ: Princeton University Press.
- [31] Brewer, Dominic J., Eric R. Eide, and Ronald G. Ehrenberg (1999). “Does It Pay to Attend an Elite Private College? Cross-Cohort Evidence on the Effects of College Type on Earnings.” *Journal of Human Resources*. 34(1): 104-123.
- [32] Brint, Steven, Allison M. Cantwell, and Preeta Saxena (2011). “Disciplinary Categories, Majors, and Undergraduate Academic Experiences: Rethinking Bok’s ‘Underachieving Colleges’ Thesis.” *Research in Higher Education*. 53:1-25.
- [33] Camargo, B., Stinebrickner, R., and T. Stinebrickner (2010) “Interracial Friendships in College,” *Journal of Labor Economics* 156(1): 106-129.
- [34] Caldwell, Ronald C., Jr. (2010). “The Effects of University Affirmative Action Policies on the Human Capital Development of Minority Children: Do Expectations Matter?” Working Paper.

- [35] Card, David and Alan 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.
- [36] Calsamiglia, Caterina, Jorg Franke, and Pedro Rey-Biel (2013). “The Incentive Effects of Affirmative Action in a Real-effort Tournament.” *Journal of Public Economics*, 98:15-31.
- [37] Chambers, David L., Timothy T. Clydesdale, William C. Kidder, and Richard O. Lempert (2005). “The Real Impact of Eliminating Affirmative Action in American Law Schools: An Empirical Critique of Richard Sander’s Study.” *Stanford Law Review* 57(6): 1855-1898.
- [38] Chan, Jimmy and Erik Eyster (2003). “Does Banning Affirmative Action Lower College Student Quality?” *American Economic Review*. 93(3): 858-872.
- [39] Chaney, Bradford, Lana D. Muraskin, Margaret W. Cahalan, and David Goodwin (1998). “Helping the Progress of Disadvantaged Students in Higher Education: The Federal Student Support Services Program.” *Education Evaluation and Policy Analysis* 20(3): 197-215.
- [40] Clotfelter, Charles T., Hellen F. Ladd, and Jacob L. Vigdor (2012). “The Aftermath of Accelerating Algebra: Evidence from a District Policy Initiative.” NBER working paper #18161.
- [41] Coate, Stephen and Glenn C. Loury (1993). “Will Affirmative-Action Policies Eliminate Negative Stereotypes?” *American Economic Review*. 83(5): 1220-1240.
- [42] Cohodes, Sarah and Joshua Goodman (2014). “Merit Aid, College Quality and College Completion: Massachusetts’ Adams Scholarship as an In-Kind Subsidy.” *American Economic Journal: Applied Economics* 6(4): 251-285.
- [43] Cortes, Kalena E. (2010) “Do Bans on Affirmative Action Hurt Minority Students? Evidence from the Texas Top 10% Plan.” *Economics of Education Review* 29(6): 1110-1124.

- [44] Cotton, Christopher, Brent R. Hickman, and Joseph P. Price (2014). "Affirmative Action and Human Capital Investment: Evidence from a Randomized Field Experiment." NBER working paper #20397.
- [45] Cullen, Julie Berry, Mark C. Long and Randall Reback (2012). "Jockeying for Position: Strategic High School Choice Under Texas' Top Ten Percent Plan." *Journal of Public Economics* 92: 32-48.
- [46] 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.
- [47] Dale, Stacy, and Alan B. Krueger (2014). "Estimating the Return to College Selectivity over the Career Using Administrative Earnings Data." *Journal of Human Resource* 49(2): 323-358.
- [48] Daugherty, Lindsay, Francisco Martorell and Isaac McFarlin, Jr. (2012). "Percent Plans, Automatic Admissions, and College Outcomes." *IZA Journal of Labor Economics* 3(10).
- [49] Dickson, Lisa M. (2006). "Does Ending Affirmative Action in College Admissions Lower the Percent of Minority Students Applying to College?" *Economics of Education Review* 25(1): 109-119.
- [50] Dillon, Eleanor and Jeffrey Smith (2013). "The Determinants of Mismatch Between Students and Colleges." NBER Working Paper No. 19286.
- [51] Doyle, William R. (2009). "The Effect of Community College Enrollment on Bachelor's Degree Completion." *Economics of Education Review* 28(2): 199-206.
- [52] Durlauf, Steven N. (2008). "Affirmative Action, Meritocracy, and Efficiency." *Politics, Philosophy, & Economics* 7: 131-158.

- [53] Epple, Dennis, Richard Romano, and Holger Sieg (2008). “Diversity and Affirmative Action in Higher Education.” *Journal of Public Economic Theory* 10(4): 474-501.
- [54] Espenshade, Chung, and Walling (2004). “Admission Preferences for Minority Students, Athletes, and Legacies at Elite Universities.” *Social Science Quarterly*, 85(5): 1422-1446.
- [55] Fitzpatrick, Maria D. and Damon Jones (2012). “Higher Education, Merit-Based Scholarships and Post-Baccalaureate Migration.” NBER Working Paper #18530.
- [56] Fryer, Roland G. Jr., Glenn C. Loury, and Tolga Yuret (2008). “An Economic Analysis of Color-Blind Affirmative Action.” *Journal of Law, Economics, and Organization*, 24(2): 319-355.
- [57] Furstenberg, Eric (2010). “College admissions reform and racial SAT score gaps: an empirical analysis”, working paper.
- [58] Gemici, Ahu, and Matt Wiswall (2014). “Evolution of Gender Differences in Post-Secondary Human Capital Investments: College Majors” *International Economic Review*, 55(1):137-85.
- [59] Gladwell, Malcolm (2013). *David and Goliath: Underdogs, Misfits, and the Art of Battling Giants*. Little, Brown, and Company.
- [60] Green, Kenneth C., Scott Jaschik, and Doug Lederman (2011). “The 2011 Inside Higher Ed Survey of College & University Admissions Directors” Inside Higher Ed.
- [61] Hastings, Justine, Christopher Neilson, and Seth Zimmerman (2014). “Are Some Degrees Worth More Than Others? Evidence From College Admission Cutoffs in Chile”, NBER working paper #19241.
- [62] Hickman, Brent R. (2013). “Pre-College Human capital Investment and Affirmative Action: A Structural Policy Analysis of US College Admissions.” Working Paper.

- [63] 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.
- [64] Hinrichs, P. (Forthcoming). “Affirmative Action Bans and College Graduation Rates.” *Economics of Education Review*
- [65] Ho, Daniel E. (2005a) “Why Affirmative Action Does Not Cause Black Students to Fail the Bar”, *Yale Law Journal* Vol. 114, No. 8, 1997-2004.
- [66] Ho, Daniel E. (2005b) “Affirmative Action’s Affirmative Actions: A Reply to Sander”, *Yale Law Journal* Vol. 114, No. 8.
- [67] Hoekstra, Mark (2009). “The Effect of Attending the Flagship State University on Earnings: A Discontinuity-Based Approach.” *Review of Economics and Statistics* 91(4): 717-724.
- [68] Holzer, Harry and David Neumark (2000). “Assessing Affirmative Action.” *Journal of Economic Literature* 38(3): 483-568.
- [69] Howell, Jessica S. (2010). “Assessing the Impact of Eliminating Affirmative Action in Higher Education.” *Journal of Labor Economics* 28(1): 113-166.
- [70] Hoxby, Caroline M. (2009). “The Changing Selectivity of American Colleges.” *Journal of Economic Perspectives*. 23(4): 95-118.
- [71] Hoxby, Caroline and Christopher Avery (2013). “The Missing ‘One-Offs’: The Hidden Supply of High-Achieving, Low Income Students.” *Brookings Papers on Economic Activity* Spring: 1-65.
- [72] Hoxby, Caroline and Sarah Turner (2013). “Expanding College Opportunities for High-Achieving, Low Income Students”, working paper.

- [73] Johnson, Valen (2003). *Grade Inflation: A Crisis in College Education*, Springer-Verlag New York, Inc., New York, NY.
- [74] Kain, John F., Daniel M. O'Brien and Paul A. Jargowsky (2005). "Hopwood and the Top 10 Percent Law: How They Have Affected the College Enrollment Decisions of Texas High School Graduates." Report to the Andrew W. Mellon Foundation: http://www.utdallas.edu/research/tsp-erc/pdf/wp_kain_2005_hopwood_top_10_percent.pdf
- [75] Kane, Thomas J. (1998). "Racial and Ethnic Preferences in College Admission," in *The Black-White Test Score Gap*, ed. by C. Jencks and M. Phillips. Washington, DC: Brookings Institution, 431456.
- [76] Kane, Thomas J. and Cecilia Elena Rouse (1995). "Labor-Market Returns to Two- and Four-Year College." *American Economic Review* 85(3): 600-614.
- [77] Kirkeboen, Lars, Edwin Leuven, and Magne Mogstad (2014). "Field of Study, Earnings, and Self-Selection," NBER Working Paper #20816.
- [78] Lehmann, Jee-Yeon K. (2013). "Job Assignment and Promotion Under Statistical Discrimination: Evidence from the Early Careers of Lawyers", working paper.
- [79] Light, Audrey and Wayne Strayer (2000). "Determinants of College Completion: School Quality or Student Ability?" *Journal of Human Resources* 35(2): 299-332.
- [80] Light, Audrey and Wayne Strayer (2002). "From Bakke to Hopwood: Does Race Affect College Attendance and Completion?" *Review of Economics and Statistics* 84(1): 34-44.
- [81] Long, Bridget Terry (2004). "How Have College Decisions Changed over Time? An Application of the Conditional Logistic Choice Model." *Journal of Econometrics* 121(1-2): 271-296.

- [82] Long, Bridget Terry and Michal Kurlaender (2009). "Do Community Colleges Provide a Viable Pathway to a Baccalaureate Degree?" *Education Evaluation and Policy Analysis* 31(1): 30-53.
- [83] Long, Mark C. (2004a). "College Applications and the Effect of Affirmative Action." *Journal of Econometrics* 121(1-2): 319-342.
- [84] Long, Mark C. (2004b). "Race and College Admission: An Alternative to Affirmative Action?" *The Review of Economics and Statistics* 86(4): 1020-1033.
- [85] Long, Mark C. (2007). "Affirmative Action and its Alternatives in Public Universities: What Do We Know?" *Public Administration Review* 67(1): 311-325.
- [86] Long, Mark C. (2010a). "Affirmative Action at Nearby Colleges: Temporal and Regional Changes." *The B.E. Journal of Economic Analysis & Policy* (Contributions) 10(1), Article 65.
- [87] Long, Mark C. (2010b). "Changes in the Returns to Education and College Quality." *Economics of Education Review* 29(3): 338-347.
- [88] Long, Mark C., Victor Saenz, and Marta Tienda (2010). "Policy Transparency and College Enrollment: Did the Texas Top Ten Percent Law Broaden Access to the Public Flagships?" *The Annals of the American Academy of Political and Social Science* 627(1): 82-105.
- [89] Long, Mark C. and Marta Tienda (2008). "Winners and Losers: Changes in Texas University Admissions Post-Hopwood." *Education Evaluation and Policy Analysis* 30(3): 255-280.
- [90] Lott, J.R. (2011). "Peer Effects in Affirmative Action: Evidence from Law Student Performance." *International Review of Law and Economics* 31(1): 1-15.
- [91] Loury, Linda Datcher and David Garman (1995). "College Selectivity and Earnings." *Journal of Labor Economics* 13(2): 289-308.

- [92] Luppino, Marc and Richard H. Sander (2012). “College Major Competitiveness and Attrition from the Sciences,” working paper.
- [93] Niu, Sunny Xinchun and Marta Tienda (2010). “The Impact of the Texas Top Ten Percent Law on College Enrollment: A Regression Discontinuity Approach.” *Journal of Policy Analysis and Management* 29(1): 84-110.
- [94] Reynolds, C. Lockwood (2012). “Where to Attend? Estimating the Effects of Beginning College at a Two-year Institution.” *Economics of Education Review* 31(4): 345-362.
- [95] Rosenbaum, Paul R. (1984). “The Consequences of Adjustment for a Concomitant Variable That Has Been Affected by the Treatment” *Journal of the Royal Statistical Society. Series A* 147(5): 656-666.
- [96] Rothstein, Jesse and Albert Yoon (2008). “Affirmative Action in Law School Admissions: What Do Racial Preferences Do?” *University of Chicago Law Review* Vol. 75.
- [97] Rothstein, Jesse and Albert Yoon (2009). “Mismatch in Law School.” working paper.
- [98] Sabot, Richard and John Wakeman-Linn (1991). “Grade Inflation and Course Choice.” *Journal of Economic Perspectives* 5(1): 159-170.
- [99] Sander, Richard H. (2004). “A Systematic Analysis of Affirmative Action in American Law Schools.” *Stanford Law Review* 57(2): 367-483.
- [100] Sander, Richard H. (2005a). “Mismeasuring the Mismatch: A Response to Ho.” *The Yale Law Journal* 114(8): 2005-2010.
- [101] Sander, Richard H. (2005b). “A Reply to Critics”, *Stanford Law Review* 57(6): 1963-2016.

- [102] Sander, Richard H. and Jane Bambauer (2012). “The Secret to My Success: How Status, Eliteness, and School Performance Shape Legal Careers.” *Journal of Empirical Legal Studies* 9(4): 893-930.
- [103] Sander, Richard H. and Stuart Taylor Jr. (2012). *Mismatch: How Affirmative Action Hurts Students It’s Intended to Help, and Why Universities Won’t Admit It*, Basic Books.
- [104] Schotter, Andrew, and Keith Weigelt (1992). “Asymmetric Tournaments, Equal Opportunity Laws, and Affirmative Action: Some Experimental Results.” *Quarterly Journal of Economics*, 107(2):511-539.
- [105] Scott-Clayton, Judith (2011). “On Money and Motivation: A Quasi-Experimental Analysis of Financial Incentives for College Achievement.” *Journal of Human Resources* 46(3): 614-636.
- [106] Sjoquist, David J. and John V. Winters (2012). “State Merit-Based Financial Aid Programs and College Attainment.” IZA Discussion Paper #6801.
- [107] Smyth, F.L., and J.J. McArdle (2004). “Ethnic and Gender Differences in Science Graduation at Selective Colleges with Implications for Admission Policy and College Choice.” *Research in Higher Education* 45(5): 353-381.
- [108] Stinebrickner, Todd R. and Ralph Stinebrickner (2011). “Math or Science? Using Longitudinal Expectations Data to Examine the Process of Choosing a College Major.” NBER Working Paper #16869.
- [109] Stinebrickner, Todd R. and Ralph Stinebrickner (2012). “Learning about Academic Ability and the College Dropout Decision.” *Journal of Labor Economics* 30(4): 707-748.
- [110] Wightman, Linda F. (1998). “LSAC National Bar Passage Study”, Law School Admission Council, Newton PA.

- [111] Wightman, Linda F. (2003). “The Consequences of Race-Blindness: Revisiting Prediction Models with Current Law School Data.” *Journal of Legal Education* 53(2): 229-253.
- [112] Williams, Doug (2013). “Do Racial Preferences Affect Minority Learning in Law Schools?.” *Journal of Empirical Legal Studies* 10(2): 171-195.
- [113] Yagan, Danny (2014). “Affirmative Action Bans and Black Admission Outcomes: Selection-Corrected Estimates from UC Law Schools” UC Berkely Working Paper: <http://eml.berkeley.edu/~yagan/AffirmativeAction.pdf>
- [114] Zafar, Basit (2011). “How Do College Students Form Expectations?” *Journal of Labor Economics* 29(2): 301-348.
- [115] Zimmerman, Seth (forthcoming). “The Returns to College Admission for Academically Marginal Students.” *Journal of Labor Economics*.

Table 1: Affirmative Action Bans in Undergraduate Admissions by State

State	Year of Enactment	Source of Ban
Texas	1997*	Court Ruling: <i>Hopwood v. Texas</i>
California	1998	Voter Initiative (Prop. 209)
Washington	1999	Voter Initiative (Initiative 200)
Florida	2001	Executive Order 99-281
Georgia**	2002	Court Ruling: <i>Johnson v. University of Georgia</i>
Michigan***	2006	Voter Initiative (Proposition 2)
Arizona	2010	Voter Initiative (Proposition 107)
New Hampshire	2012	Legislative Act (HB 623)
Oklahoma	2013	Voter Initiative (State Question 759)

* The affirmative action ban in Texas was overturned in 2003 by the 5th Circuit Court of Appeals.

** This ruling only affects the University of Georgia. The race-based admission system used was ruled unconstitutional, but other public institutions in Georgia still consider race in admissions.

*** In July 2011, the Michigan affirmative action ban was ruled unconstitutional by the 6th Circuit Court of Appeals. The case was argued in front of the Supreme Court in October 2013, but no decision has been made on the constitutionality of this law.

Table 2: Median Institutional Characteristics by College Quality Tier

College Characteristic	Flagship	Non-Flagship	Highly-Selective	Less-Selective
	Public	Public	Private	Private
Six-Year Graduation Rate	0.702	0.444	0.868	0.484
Black Six-Year Graduation Rate	0.598	0.333	0.810	0.333
Hispanic Six-Year Graduation Rate	0.670	0.393	0.818	0.417
White Six-Year Graduation Rate	0.715	0.460	0.869	0.514
Student-Faculty Ratio	11.59	13.84	6.36	13.48
Per-Student Expenditure	39,685	15,741	54,028	17,232
Per-Student Instructional Expenditure	12,081	6,054	20,478	5,433
Faculty Salaries	91,177	65,961	96,101	54,520

Source: 2011-2012 IPEDS data. The per-student expenditures and student-faculty ratios use all student enrollments, not just undergraduates. Faculty salaries and student-faculty ratios use all instructional staff, which includes non-tenure track lecturers.