Does Banning Affirmative Action Harm College Student Quality?

Jimmy Chan
Department of Economics
Johns Hopkins University
3400 N.Charles Street
Baltimore, MD 21218
Tel:410-516 7610
Fax:410-516 7600
chan@jhunix.hcf.jhu.edu
and

Erik Eyster
Department of Economics
University of California, Berkeley
549 Evans Hall #3880
Berkeley, CA94720-3880
Tel:510 642 0822
Fax: 510 642 6615
eyster@econ.berkeley.edu

December 14, 1999

Abstract

Banning affirmative action from college admissions decisions cannot prevent an admissions office that cares about diversity from achieving it through channels other than the explicit consideration of race. We construct a model of college admissions where candidates from two groups with different average quali cations compete for a xed number of seats. When an admissions office that cares both about the quality and diversity of its entering class can use group identity as a criterion for admissions, its preferred admissions rule selects the best-quali ed candidates from each group. When it cannot use affirmative action, the admissions office s preferred rule generally does not select the best-quali ed candidates from either group: it randomizes over candidates to achieve diversity, at the expense of within-group selection. A ban always reduces diversity, and may also lower average quality. Moreover, even when a total ban on affirmative action raises average quality, a partial ban may raise average quality even more.

Acknowledgments We thank Bob Anderson, Kim-Sau Chung, Bill Dickens, Aaron Edlin, Joe Harrington, Lorraine Mitchell, Tarun Sabarwal, Ilya Segal, Chris Shannon, and seminar participants at Berkeley, Johns Hopkins, the Southeast Economic Theory Conference, and SITA 1999 for helpful comments. We are especially grateful to Matthew Rabin for his many helpful suggestions and untiring encouragement.

1 Introduction

After several decades of widespread use in American college admissions, affirmative action has come under increasing attack. Recent court rulings and ballot referenda have banned it from admissions decisions at public universities in California, Texas, and Washington, and it is currently being challenged in several other states. But while a court or voters can prevent a college from using affirmative action, neither one can dictate its admissions policy. In principle, a college s regents or trustees may have considerable authority over its admissions policies, but in practice the formulation of an admissions rule is invariably delegated to the college s administration. Since many college administrations care about diversity, they will respond to bans on affirmative action by seeking out new channels to achieve it.[1]

College admissions is a complex process. To select their entering classes, many admissions offices consider a multitude of factors, including candidates high-school records, standardized-test scores, non-academic achievements, socio-economic backgrounds, etc. Because it has considerable latitude in how it weights these various factors, a college admissions office can promote diversity by increasing the importance of those factors most favorable to minority candidates, at the expense of other factors. For example, it might rely less on standardized-test scores and more on demonstrated ability to overcome adversity. Since underrepresented-minority candidates as a group have lower average standardized-test scores than whites and Asians, and are more likely to come from disadvantaged backgrounds, such a change may increase the number of minority candidates admitted. But if candidates credentials were properly weighted under affirmative action, then reweighting them under a ban is inefficient: the new admissions rule does not admit the best-quali ed candidates from any racial group. If the new admissions rule is inefficient enough, then a ban on affirmative action designed to increase student quality may back re, lowering quality instead.[2][3]

This paper uses a simple model of a college s admissions process to explore the consequences of banning affirmative action from college admissions. In our model, the college s regents or trustees, who determine whether the college may use affirmative action, have different preferences over

[1] Most American universities and colleges care about diversity. For example, see statements by the American Council on Education [1998] and the Association of American Universities [1997], both at www.umich.edu/~newsinfo/Admission/admiss.html

[2] In this paper, the term quality means narrowly-de ned academic ability.

[3] In principle, a college s regents could overcome this agency problem by stipulating that candidates be admitted solely on the basis of SAT score. However, doing so may lower the quality of the admitted students substantially. Using detailed admissions records for a selective college s 1989 freshman class, Vars and Bowen [1998] nd that the college s academic ratings, which integrate all objective information used to predict academic performance, predicts college grades much better than SAT scores alone (with R^2 statistics of 0.368 versus 0.276).

admissions rules than its admissions office, which formulates an admissions rule in compliance with the regents affirmative action policy.[4] Speci cally, while both the regents and the admissions office care only about student quality and diversity, the admissions office cares more about diversity than do the regents. We show that under a ban on affirmative action, an admissions office s preferred rule may partially ignore candidates quali cations. If such a rule is inattentive enough to candidates quali cations, then banning affirmative action can lead to a less quali ed entering class. More surprisingly, even when a total ban on affirmative action raises average quality compared to affirmative action, a partial ban may raise average quality even more. Thus, banning affirmative action may not be the optimal policy even for regents who care only about average quality.

Admissions candidates in our model come from two groups: a majority group and a minority group.[5] All candidates take a test, and test scores are informative signals of academic ability. Minority candidates tend to score lower than majority candidates. Within any group, both the admissions office and the regents prefer higher-scoring candidates to lower-scoring candidates. Formally, an admissions rule is an increasing function that assigns each test score a probability of admission. Thus, the admissions office is forbidden to directly admit low-scoring candidates over higher-scoring candidates, but it is allowed to use an admissions rule that partially depends on a random signal uncorrelated with quality or group identity. If the admissions rule depends on the random signal, a high-scoring candidate is still admitted with a probability at least as high as that of a low-scoring candidate, but the highest-scoring candidates are not admitted with probability one.

Under affirmative action, the admissions office can apply a different admissions rule to each group. In that case, its favored admissions rule is a threshold rule: each group is assigned a threshold, and any candidate whose test score meets or exceeds her group s threshold is admitted. The admissions office s taste for diversity determines the group composition of its entering class: if the minority group scores lower than the majority group on the test (in a sense to be made precise later), then the admissions office sets a lower threshold for minority candidates than for majority candidates.

When the regents ban affirmative action, the admissions office must use an admissions rule

[4]The actors in our model need not be a university s regents and its admissions office; for example, a court or voters of a state may play the role of the regents, and a state legislature that of the admissions office. Nevertheless, for expositional clarity, the labels regents and admissions office are used henceforth.

[5]Nothing in the model depends on the majority group outnumbering the minority group; rather, the terms majority and minority are used to clarify exposition and relate the formal model to current debate over university admissions.

that does not depend on group identity. In this case, the admissions office may randomize over candidates to achieve diversity. Although it prefers higher-scoring candidates to lower-scoring candidates *within each group*, the admissions office may not prefer the same over the entire applicant pool. If the admissions office cares enough about diversity, it may prefer low-scoring candidates to high-scoring candidates because a larger share of low-scoring candidates belong to the minority group. But because the admissions office is constrained to use an admissions rule that is increasing in test score, it cannot admit low-scoring candidates with positive probability without admitting *all* higher-scoring candidates with probabilities at least as large. Therefore, it may prefer to randomize, and admit a wide range of candidates with positive probability, than to use a threshold rule that admits only the highest-scoring candidates with probability one, rejecting everyone else. Whether the admissions office randomizes due to a preference for low-scoring candidates depends upon its taste for diversity and the distributions of test scores across the different groups.

Any meaningful preference for diversity must displace candidates from some group to make room for candidates from another. An affirmative-action threshold rule achieves diversity by replacing the least-quali ed majority candidates who would have been admitted absent affirmative action with the most-quali ed minority candidates who would not have been admitted absent affirmative action. Randomization, however, does not do the same. Because the admissions office s preferences over candidates within a given group are increasing in test score, any admissions policy that rejects candidates with high test scores in some group in favor of those with lower test scores in the same group introduces inefficiency: it does not select the best-quali ed candidates from that group. Thus, a probabilistic admissions comes with a cost. Indeed, whenever the admissions office randomizes, both parties would prefer a class selected by affirmative action *with the same racial composition* as the one selected by randomization. A policy lesson is that if the regents believe that the admissions office cares too much about diversity, then they may do better by cajoling the admissions office into reducing affirmative action than by banning affirmative action altogether.

In our model, banning affirmative action always reduces the minority group s share of the entering class. More surprisingly, it may lower the quality of the entering class. Under affirmative action, only the best-quali ed candidates from each group are accepted, but the marginal majority candidate is better quali ed than the marginal minority candidate. Randomization, however, lowers student quality *within each group*. The relative size of these two effects determines which class is of higher quality; for example, if the quality gap between marginal minority and majority candidates under affirmative action is relatively small, while randomization falls far short of admitting the best-qualified candidates from each group, then banning affirmative action lowers student quality. Likewise, banning affirmative action may increase the gap in average test score between matriculants from the two groups. If such a gap causes harmful racial stereotyping, then banning affirmative action may exacerbate stereotyping. Because a ban on affirmative action always increases the number of majority candidates admitted, and affirmative action admits the highest-scoring majority candidates, a ban always lowers the average test score of majority matriculants. However, if the admissions office randomizes enough that minority candidates who are not accepted under affirmative action are accepted with positive probability under a ban, then the average test score for minority matriculants may fall as well. If it falls enough, then banning affirmative action may increase the average test-score gap.

Even when a total ban on affirmative action raises the quality of the entering class, a partial ban may raise quality even more. One form of partial ban is for the admissions office first to fill part of the class using a group-blind rule, and then fill the balance of the class however it likes. If the admissions office's discretionary seats are sufficiently few in number, then it will use them to admit the highest-scoring minority candidates not already admitted. These minority candidates are better qualified than the randomly selected candidates they replace. But there is a countervailing effect: because under a partial ban the admissions office knows that it can admit some number of high-scoring candidates no matter how much it randomizes, it may randomize more than it does under a total ban. We show that in some cases the first effect dominates, and, hence, the quality of the entering class is higher under a partial ban than a total ban.

The next section presents a brief overview of affirmative action in college admissions and recent bans. Section 3 introduces a formal model of a college's admissions process. Section 4 characterizes admissions rules under affirmative action; Section 5 characterizes admissions rules under a ban; and Section 6 compares the two regimes. Section 7 examines limited affirmative action, and Section 8 concludes.

2 Reactions to Bans on Affirmative Action

In 1995, the Regents of the University of California banned race-conscious admissions. In the 1996 case *Hopwood v. Texas*, a panel from the Fifth Circuit Court of Appeals, which covers Louisiana, Mississippi, and Texas, struck down the admissions system at the law school at the University of Texas at Austin; the court's ruling is interpreted to forbid race-conscious admissions at all public universities in Louisiana, Mississippi, and Texas. In the same year, the voters of California approved Proposition 209, which prohibits public colleges and universities from using race in any state admissions or financial-aid decision. In 1998, Washington state voters passed a proposition much like California's 209. Currently, lawsuits similar to *Hopwood* are pending in Maryland, Michigan, Oklahoma, and Washington.[6]

At public universities in both California and Texas, bans on affirmative action were followed by significant declines in minority enrollment at the most prestigious campuses. The ban at the University of California went into effect with the selection of Fall 1998 freshmen. At Berkeley, African Americans fell from 7.6 to 3.9 percent of the entering class, while Chicanos and Latinos fell from 13.9 to 8.5 percent.[7] At all University of California campuses, African Americans fell as a share of matriculants by 27.6 percent, and Latinos by 10.0 percent. At the University of Texas at Austin, where the ban was implemented a year earlier, African Americans fell from 4.1 to 2.7 percent of the entering class, while Chicanos and Latinos fell from 14.5 to 12.6 percent.

In response to declines in minority enrollment, universities in California and Texas have changed their admissions policies in ways favorable to minority candidates.[8] In 1998, Berkeley adopted a new admissions policy that substantially alters its method of measuring academic achievement. University of California system-wide policy requires that each campus fill at least half of its freshman class solely on the basis of academic achievement, which Berkeley previously measured by an Academic Index Score, a mathematical formula based on high-school GPA, SAT I, and SAT II (achievement test) scores. Berkeley's new admissions policy replaces the old Academic Index Score with a broader measure of academic achievement that includes factors such as the type and number of high-school classes taken, grades in individual courses, and performance relative to high-school classmates. However, unlike the Academic Index, the new measure of academic qualification does not assign any specific weight to any variable; instead, admissions committee members have the discretion to rate applications based on their overall impressions of candidates' credentials. Together, American Indians, African Americans, Chicanos and Latinos made up 14.4 percent of the freshman class for Fall 1999, up from 12.8 percent in 1998. (See Table 1.)

In March of 1999, the Regents of the University of California adopted a proposal that, starting in the Fall of 2001, grants eligibility to the University of California to every student in the top four percent of her high-school class, where class rank is determined by Academic Index Score. (California's master plan for higher education calls for 12.5 percent of high-school graduates to be

[6] See Alger [1998] for an overview of current lawsuits related to affirmative action.

[7] The figures for UC Berkeley include only those who report their ethnicity. See footnote 11 for details.

[8] Universities in California and Texas have made other changes to increase minority enrollment, including increasing minority recruitment.

eligible for admission to the University of California. However, UC-eligibility does not guarantee admission to every UC campus.) Thus, whereas a high-school senior used to be UC-eligible if she was in the top 12.5 percent of high-school seniors statewide, starting in 2001 she will be UC-eligible if she is in the top 4 percent of her graduating class or in the top 8.5 percent of high-school seniors statewide not in the top 4 percent of their graduating classes. Geiser [1998] reports simulations showing that the change will increase the share of UC-eligible candidates who are African-American, Chicano, or Latino by approximately 10 percent.

Boalt Hall, the University of California at Berkeley s law school, attracted national media attention when following implementation of the ban in 1997 its entering class of 268 included only one African American.[9] For the next academic year, administrators made a number of changes to their admissions policy. One was to no longer assign candidates an academic index number, which previously had been a function of undergraduate GPA, weighted by the quality of the candidate s undergraduate institution, and LSAT score. The admissions office no longer adjusts candidates GPAs to account for the quality of their undergraduate institutions. Also, the admissions committee no longer reports candidates exact LSAT scores to admissions committee members; instead, the admissions committee partitions the range of LSAT scores into intervals, and admissions committee members are told which interval contains a candidate s score. In the year following these and other changes, African American enrollment increased from 1 to 8, and Chicano and Latino enrollment from 14 to 23. (See Table 2.)

The Texas state legislature responded to *Hopwood* in 1997 by passing a law that requires each campus of the University of Texas (e.g., UT Austin or Texas A&M) to admit any candidate who graduated in the top ten percent of her high-school class (in the previous two years), where rank is determined solely by high-school GPA. Following the change, African-American enrollment at UT Austin increased from 2.7 to 3.0 percent, and Chicano and Latino enrollment increased from 12.6 to 13.2 percent. (See Table 3.) Holley and Spencer [1999] argue that minority enrollment will increase even further in the future when the new rule is better publicized.[10]

These recent changes in admissions policies share two features. First, they all were adopted after bans on affirmative action led to signi cant drops in minority enrollment. Second, they all deemphasize the use of standardized-test scores. In UC Berkeley and UT Austin undergraduate admissions, more weight has been put on high-school rankings. With the exception of UT Austin undergraduate admissions, more weight has also been put on less quanti able aspects of candi-

[9] The California ban was implemented a year earlier in graduate admissions than in undergraduate admissions.
[10] See Wilgoren [1999] for details of the impact of the 10% rule in Texas.

dates quali cations. Given that underrepresented minority groups have lower standardized-test scores than whites and Asians, and that high schools are partially racially segregated, both changes are likely to increase minority enrollment. Tables 1-3 report enrollment gures for new rst-year registrants at Boalt Hall, UC Berkeley, and UT Austin.[11] At all three institutions, minority enrollment rose with implementation of the new admissions policies.[12] While these policies may not have been designed entirely to increase minority enrollment, in several cases administrators have attributed increases in minority enrollment in part to the new policies.

Thus, the recent changes in admissions policies at some of the elite public universities and professional schools in California and Texas support one of the main predictions of our model, namely that admissions offices may respond to bans on affirmative action by altering their admissions standards in ways favorable to minority candidates. But these methods are likely to be inefficient: they do not admit the best-quali ed candidates from any ethnic group. For example, the top-ten-percent law in Texas may force colleges to reject candidates in the second decile of an excellent high school to make room for candidates in the top decile of a mediocre high school, even when the former are more quali ed than the latter: without knowing the optimal weight on standardized-test scores, we can be fairly con dent that it is not zero. Likewise, Boalt Hall s new policy of not reporting exact LSAT scores is inefficient because it forces the admissions rate to be constant over intervals of test scores: even within an ethnic group, the admissions committee may not be able to admit higher-scoring candidates with higher probability than lower-scoring candidates.[13]

Despite their inefficiencies, we do not know that these new policies are so inefficient as to lower average student quality relative to affirmative action. Since standardized-test scores are imperfect

[11] All three institutions give applicants the option of not reporting their ethnicity. At UC Berkeley, the number of students refusing to report their ethnicity grew signi cantly after the ban, from about 6% of the freshman class in 1997 to 15.3% in 1998 and 8.8% in 1999. Since UC Berkeley is barred by UC regulations from tracking down non-reporting students, there is no conclusive data on their ethnic composition. We strongly suspect, however, that the increase in the non-response rate is due mainly to a change in the application form in 1998, and is not directly related to the ban on affirmative action. Before 1998, item 17 on the application form asked applicants to identify themselves as a member of one of several ethnic groups listed on the form. In 1998, when the ethnicity question was moved to item 130 and applicants had to look up an ethnic-group code from a separate pamphet, the non-response rate increased more than 150 percent. In 1999, when the ethnicity question remained item 130 but ethnic groups were once again listed on the application form, the non-respeonse rate fell more than 40 percent, leaving it less than 50 percent higher than in 1997. Thus, we report the size of ethnic groups at UC Berkeley by their shares of of new rst-year registrants reporting their ethnicity. Most new registrants at Boalt Hall and UT Austin report their ethnicity voluntarily, and the fraction of the students who refused to report did not change signi cantly after the ban. For these two institutions, we report the size of ethnic groups by their shares of all new registrants.

[12] Recall that the new admissions policies were implemented at Boalt Hall and UT Austin in 1998 and at UC Berkeley in 1999.

[13] The law school s rationals for this change is that differences of one to three LSAT points are not signi cant. While this largely may be true, the new practice clearly makes certain one-point differences in score *very* signi cant, namely those that move a candidate from one element of the partition to the next.

predictors of academic ability, the fact that admissions offices put less weight on them now than before does not imply that they put too little weight on them. Our theory, however, predicts that to the extent that changes in admissions policies were spurred by the belief that standardized-test scores are unfavorable to minority candidates, new admissions rules will tend to rely too little on test scores, when compared to optimal rules under affirmative action.

Table 1: First-Time Freshman at UC Berkeley by Ethnicity
(Percent of Fall Registrants Reporting Ethnic Data, Number in Parentheses)

	1995	1996	1997	1998	1999
African American	6.8	6.6	7.6	3.9	3.8
	(222)	(233)	(257)	(126)	(130)
American Indian	1.9	1.5	0.7	0.4	0.7
	(63)	(52)	(23)	(14)	(23)
Asian American	39.0	40.8	43.4	49.0	47.1
	(1,268)	(1,432)	(1,468)	(1,565)	(1,595)
Chicano and Latino	16.3	15.6	13.9	8.5	9.9
	(531)	(549)	(472)	(271)	(335)
White	31.3	31.0	30.1	34.1	33.6
	(1,018)	(1,090)	(1,018)	(1,090)	(1,138)
Foreign	3.2	2.7	2.1	2.5	3.0
	(105)	(96)	(72)	(81)	(100)
Other	1.4	1.7	2.2	1.5	1.9
	(47)	(60)	(76)	(48)	(63)

Source: Office UC Berkeley, Admissions & Enrollment

*1999 Preliminary Data as of September 9, 1999

Table 2: First-Year Registrants at UC Berkeley School of Law by Ethnicity

(Percent of Fall Registrants, Number in Parentheses)

	1995	1996	1997	1998	1999
African American	7.9	7.6	0.4	3.0	2.6
	(21)	(20)	(1)	(8)	(7)
American Indian	1.9	1.5	0.0	0.7	0.7
	(5)	(4)	(0)	(2)	(2)
Asian	12.0	14.4	12.7	14.1	11.2
	(32)	(38)	(34)	(38)	(30)
Chicano and Latino	13.5	10.6	5.2	8.6	5.9
	(36)	(28)	(14)	(23)	(16)
Non-Minority	64.7	65.8	81.7	73.6	79.6
	(172)	(173)	(219)	(198)	(214)

Source: UC Berkeley School of Law, 1999 Annual Admissions Report

Note: Non-Minority Includes Registrants Not Reporting Ethnicity

Table 3: First-Time Freshman at UT Austin by Ethnicity

(Percent of Fall Registrants)

	1995	1996	1997	1998
African American	4.9	4.1	2.7	3.0
American Indian	0.4	0.5	0.5	0.5
Asian American	14.2	14.7	15.9	16.8
Chicano and Latino	14.7	14.5	12.6	13.2
White	64.2	64.7	66.8	65.2
Foreign	1.5	1.5	1.5	1.2
Unknown	0.0	0.0	0.0	0.0
Total Number	6,352	6,430	7,085	6,744

Source: UT Austin, Office of Institutional Studies

(www.utexas.edu/academic/ois/stathb.98-99/s17.1.html)

3 A Model of College Admissions

3.1 Candidates

A college must select an entering class of size c from a pool of $C > c$ candidates. Each candidate comes from one of two groups, the majority group, W, or the minority group, N. We use $G \in \{N, W\}$ to denote the number of candidates in group G as well as the group s name.

Each candidate has taken a standardized test and received a test score t lying in $[\underline{t}, \bar{t}]$, the range of the test. We de ne $\lambda(t)$ to be the expected academic ability or quality of a candidate with test score t; λ is assumed not to depend on group identity.[14]

The following assumption captures the notion that higher-scoring candidates are more quali ed than lower-scoring candidates.

Assumption 1 $\lambda(t)$ *is strictly increasing and continuously differentiable.*

The distribution of test scores from candidates in group G is described by the cumulative distribution function F_G and the probability density function f_G, which is continuous and strictly positive everywhere on $(\underline{t}, \bar{t})$.[15] Let $g(t) \equiv G f_G(t)$, (loosely speaking) the number of candidates from G scoring t. An important feature of the model is that minority candidates have lower test scores than majority candidates. This captures the empirical regularity that minority-group members tend to score lower than majority-group members on many standardized tests, including those most often used in college admissions decisions. This disparity may exist because minority candidates are more likely than majority candidates to attend schools that are insufficiently funded, because they suffer discrimination, or simply because they are on average less prepared than majority candidates. (See Jencks and Phillips [1998] for descriptions and implications of racial test-score gaps, as well as Steele and Aronson [1998] for experimental evidence on stereotyping as a cause of these gaps.) We capture this notion that the minority group scores lower on the test than the majority group with the following assumption:

Assumption 2 $\frac{f_W(t)}{f_N(t)}$ *is continuously differentiable and satis es the monotone likelihood ratio condition (MLRC); that is, for each* $t \in (\underline{t}, \bar{t})$, $\frac{f_W(t)}{f_N(t)}$ *is strictly increasing in* t.

[14] There are obviously many reasons why this assumption might not hold. For example, it may be that a minority candidate scoring 1300 on the SAT is almost sure to be of high ability, but a majority candidate with the same score is not, perhaps because majority candidates often take SAT preparatory classes that enable high- and low-ability candidates alike to score well. Since in our model a preference for diversity is tantamount to a preference for minority candidates over majority candidates with the same test score, none of our results would change if the admissions office though minority candidates better quali ed than majority candidates with the same test scores.

[15] Throughout, we ignore integer constraints and think of f_G as the empirical density of G s test scores.

The MLRC means the greater a candidate s test score, the greater the probability that she belongs to the majority group. Let $\psi(t) \equiv \frac{n(t)}{n(t)+w(t)}$, the share of candidates scoring t belonging to the minority group; Assumption 2 implies that ψ is decreasing in t. Finally, let

$$\Psi(t_1, t_2) \equiv \frac{\int_{t_1}^{t_2} n(t)dt}{\int_{t_1}^{t_2} n(t) + w(t)dt},$$

the share of candidates scoring in $[t_1, t_2]$ who belong to the minority group.

Assumption 2 implies that F_W rst-order stochastically dominates F_N; that is, $F_W(t)$ □ $F_N(t)$ for all t. The following lemma will prove useful throughout.[16]

Lemma 1 *If f and g are two density functions on $[\underline{t}, \bar{t}]$ such that f rst-order stochastically dominates g, then for any weakly increasing and non-negative function h and each $t_1 \in [\underline{t}, \bar{t}]$, $\int_{t_1}^{\bar{t}} h(t)f(t)dt \geq \int_{t_1}^{\bar{t}} h(t)g(t)dt$.*

If f rst-order stochastically dominates g, then the expectation of any increasing function h under f is greater than under g; Lemma 1 simply notes that if h is non-negative, and we replace it with zero to the left of any $t_1 \in [\underline{t}, \bar{t}]$, then this new function is also increasing, so its expectation under f is greater than under g. Since f_N and f_W satisfy the MLRC, for each $[t_1, t_2] \subset [\underline{t}, \bar{t}]$, the conditional densities $\frac{f_N(t)}{\int_{t_1}^{t_2} f_N(t)dt}$ and $\frac{f_W(t)}{\int_{t_1}^{t_2} f_W(t)dt}$ on $[t_1, t_2]$ also satisfy the MLRC; hence, $\frac{f_W(t)}{\int_{t_1}^{t_2} f_W(t)dt}$ rst-order stochastically dominates $\frac{f_N(t)}{\int_{t_1}^{t_2} f_N(t)dt}$. Because $h(t) = t$ is strictly increasing, Lemma 1 implies that for each $[t_1, t_2] \subset [\underline{t}, \bar{t}]$, conditional on scoring in $[t_1, t_2]$, majority candidates score higher on average than minority candidates. Thus, the MLRC implies that over any subinterval of the test-score range (including the entire range), majority candidates score higher on average than minority candidates.

For each $G \in \{N, W\}$, let

$$\Lambda_G(t_1, t_2) \equiv \frac{\int_{t_1}^{t_2} \lambda(t) f_G(t)dt}{\int_{t_1}^{t_2} f_G(t)dt},$$

the average quality of group G candidates scoring in $[t_1, t_2]$. Likewise, let

$$\Lambda(t_1, t_2) \equiv \frac{\int_{t_1}^{t_2} \lambda(t)(n(t) + w(t))dt}{\int_{t_1}^{t_2} (n(t) + w(t))dt},$$

the average quality of all candidates scoring in $[t_1, t_2]$. Because λ is strictly increasing, Lemma 1 implies that $\Lambda_W(t_1, t_2) > \Lambda(t_1, t_2) > \Lambda_N(t_1, t_2)$: within any range of test scores, candidates from the majority group tend to be of higher quality than candidates from the minority group.

[16] All proofs are in the appendices.

3.2 Admissions

The college is formally governed by a board of regents, which decides whether or not to allow affirmative action in admissions. However, since they lack the necessary time or expertise, the regents delegate formulation of the actual admissions rule to an admissions office. When formulating its admissions rule, the admissions office knows the number of candidates and test-score distribution from each group, namely N, W, f_N and f_W, as well as λ, the measure of candidates quality.[17] All candidates apply, so candidates decisions to apply are exogenous to the model; in particular, their application decisions are independent of the admissions rule. Once candidates have applied, the admissions office mechanically applies its admissions rule. Finally, all candidates offered admission matriculate. Hence, candidates make no decisions in our model.[18]

An admissions rule assigns to each candidate a probability of admission based on her group affiliation and test score. We require that admissions rules satisfy two conditions. First, within each group, the probability of admission must be (weakly) increasing in test score, meaning that the higher a candidate s test score, the more likely she is to be admitted. Second, the admissions rule must satisfy a capacity constraint, namely that the number of admitted candidates equal c.

De nition 1 *A function $r \equiv (r_N, r_W)$ is an admissions rule if*

1. *for each $G \in \{N, W\}$, $r_G(t) : [\underline{t}, \bar{t}] \to [0, 1]$ is right-continuous and weakly increasing in t, and*
2. $\int_{\underline{t}}^{\bar{t}} r_N(t)n(t) + r_W(t)w(t)dt = c$.

If affirmative action is banned, then r_N and r_W are constrained to be identical.[19]

As we shall see, given the admissions office s preferences, it may not prefer monotone admissions rules. But we feel that a variety of factors make non-monotone rules socially undesirable or infeasible. Examples include fairness people would not think a rule that favored low-scoring candidates over high-scoring candidates were fair if they thought the test were fair and incentive compatibility constraints candidates might deliberately do poorly on the test if they thought they could improve their chances of admission.

[17] The results would be qualitatively similar if the number of applicants from different groups were unknown, but drawn from known probability distributions.

[18] Since candidates with stronger academic credentials are more likely to be admitted by other universities, the university s yield rate is likely to be decreasing in t. Incorporating this stylized fact into our model would not affect any of our results because in our model the yield rate affects only the capacity constraint, and not the admissions office s preferences over candidates with different test scores.

[19] Right continuity of the admissions rule is inessential to the qualitative results, but facilitates expression of the propositions.

Our version of monotone admissions rules is not the only natural de nition of monotonicity in an admissions context. In reality, people may want an admissions rule that is not only monotone *ex ante* (i.e., the higher a candidate s test score, the higher her chance of admission), but also monotone *ex post* (i.e., each admitted candidate s test score exceeds all unadmitted candidates scores).[20] Nevertheless, we feel that our *ex ante* (or stochastic) de nition of monotonicity is the right one. In a world of multiple admissions tests, it is not obvious how best to combine them into a single test on which to base admissions, in which case there is no easily-identi able test on which an admissions rule should be *ex post* monotone. An *ex ante* monotone rule in our model corresponds to a suboptimal weighting of multiple admissions tests.

Throughout the paper, upper-case letters $(N(r), W(r))$ denote the number of candidates from a certain group admitted under rule r. For example, $N(r)$ is the expected number of candidates in N admitted under rule r.

3.3 Preferences

Both the regents and the admissions want to maximize the total quality of admitted candidates and minimize $\left|\frac{N}{C} - \frac{N(r)}{c}\right|$, the difference in group composition between the applicant pool and the entering class under rule r.[21] We de ne the admissions office s preferences, U^{AO}, as

$$U^{AO}(r) = \int_{\underline{t}}^{\overline{t}} \lambda(t)(r_N(t)n(t) + r_W(t)w(t))dt - \alpha\left|\frac{N}{C} - \frac{N(r)}{c}\right|,$$

where α is a positive number representing the admissions office s taste for diversity. When $\alpha = 0$, the admissions office cares only about quality; when α approaches in nity, it cares only about diversity.

There are several reasons why the regents and the admissions office may care about diversity. The rst is that they may believe that N s generate externalities bene ting all students. Several studies indicate that to be the case (e.g., Bok [1982], Committee on Admissions and Enrollment [1989], and Bowen and Bok [1998]). Second, the social value added by a college degree to a minority candidate of a given academic ability may exceed that to a majority candidate of the same ability (Conrad and Sharpe [1996] and Bowen and Bok [1998]). Finally, if minority candidates have been discriminated against in the past, the admissions office may want to admit them in the

[20] Clearly rules that are *ex post* monotone are also *ex ante* monotone.

[21] In reality, the admissions office may seek to minimize $\left|\frac{N}{C} - \delta\frac{N(r)}{c}\right|$ for $\delta \in (0,1)$, meaning that it would like minority candidates to constitute a larger share of the entering class than of the candidate pool. For example, a state university might want the group composition of its entering class to mimic that of its high-school seniors, but candidates may come disproportionately from the majority group. Such an admissions office with a taste for diversity α is equivalent to an admissions office with a taste for diversity $\delta\alpha$ in our model.

interest of social justice.[22] Whatever their reasons, elite universities clearly do think that diversity is important. (For an interesting example, see Geiser s [1998] discussion of simulation results on various methods of de ning eligibility to the University of California, in which he compares different rules by their effects on diversity and on an index of academic quali cations.) In general, though, there is no reason to think that the admissions office s utility is linear in diversity. However, this simpli cation greatly facilitates the expression of optimal admissions rules without affirmative action. (See Proposition 2 below.)

A central assumption for our model is that the regents care less about diversity than the admissions office. For simplicity, we assume that the regents care *only* about quality, and not at all about diversity. Hence, their preferences, U^R, over admissions rules can be represented as

$$U^R(r) = \int_{\underline{t}}^{\overline{t}} \lambda(t)(r_N(t)n(t) + r_W(t)w(t))dt.$$

In reality, the regents preferences may be more complex than those we assume. In particular, they may care about whether a rule itself is fair, and not simply about which candidates the rule admits; if they think using group identity in an admissions rule is unfair, then they may want to ban affirmative action regardless of the effect on student quality. However, determining whether an admissions rule is fair is not an easy task, for even without affirmative action, a group may fare differently under different admissions rules.[23] For example, more minority candidates may be admitted under an admissions rule based solely on class rank than one based solely on SAT score. As we shall see, an admissions office may react to a ban on affirmative action by choosing an admissions rule more favorable to minority candidates. It is unclear whether the regents would deem this type of implicit affirmative action any more fair than explicit affirmative action. But if they do, then they may want to ban affirmative action no matter what the effect on student quality.

If the regents believe that the admissions office cares too much about diversity, then they may want to ban affirmative action or take other measures to counteract the admissions office s taste for diversity. One option which may appeal to them, as later we shall see, is to dictate that the entering class be lled by the candidates with the highest test scores. We do not allow the regents to do

[22] Strictly speaking, the last two points suggest that the admissions office might prefer minority candidates with a given test score to majority candidates with the same test score, rather than care about diversity *per se*. We do not carefully distinguish between these two motives because in our model the minority group is always under-represented (and in reality, the minority group is typically underrepresented); hence, preferring minority candidates to majority candidates with the same test score is tantamount to a taste for diversity in our model.

[23] An interesting example is that the National Merit Scholarship competition weights verbal scores on the SAT twice as heavily as math scores, while most college admissions rules weight the two equally. We are unaware of any consensus on which of the two weightings is fair, though one of us does remember having strong feelings on the matter while in high school.

this. In a richer model than ours, where candidates have several standardized-test scores, grades, recommendations, and the like, these various factors somehow must be combined into a single test score by which to compare candidates. Constructing the single test that is the most powerful in distinguishing candidate quality from a number of different tests with different distributions is a difficult problem, and one which we believe the admissions office should have far more success in solving than the regents themselves. Thus, we restrict the regents to only two options: permitting or banning affirmative action. After they decide, the admissions office picks an admissions rule to solve the constrained maximization program

$$\max_{r \in R} \quad \int_{\underline{t}}^{\bar{t}} \lambda(t)(r_N(t)n(t) + r_W(t)w(t))dt - \alpha \left| \frac{N}{C} - \frac{N(r)}{c} \right| \qquad \text{s.t.} \qquad N(r) + W(r) = c,$$

where R is the set of allowable admissions rules. If the regents ban affirmative action, then R includes only admission rules that do not depend on group affiliation.

Since we are unsure of the exact form of the regents preferences, in the next two sections we only analyze the admissions office s rational response to a ban on affirmative action, and not the regents decision to ban affirmative action itself. As de ned above, the regents preferences coincide with class quality, so the reader may wish to interpret statements about the regents welfare under the two different affirmative-action regimes as statements only about class quality.

4 Affirmative Action

When the regents allow affirmative action, the admissions office can use two separate admissions rules, one for each group. In this case, the admissions office s problem can be decomposed into two parts: rst, given a xed group composition, the admissions office picks the optimal admissions rule for each group to achieve that composition; next, it chooses the optimal group composition based on its preference for diversity.

We rst consider the problem of designing an admissions policy to admit δc candidates from group N and $(1-\delta)c$ candidates from W, where $\delta \in [0, 1]$. Because the group composition is xed, the admissions office s preference for diversity does not matter, and its problem reduces to nding the optimal admissions rule achieving this composition:

$$\max_{r \in R} \quad \int_{\underline{t}}^{\overline{t}} \lambda(t)(r_N(t)n(t) + r_W(t)w(t))dt$$

$$\text{s.t.} \quad \int_{\underline{t}}^{\overline{t}} r_N(t)n(t)dt = \delta c$$

$$\int_{\underline{t}}^{\overline{t}} r_W(t)w(t)dt = (1-\delta)c.$$

Not surprisingly, the most efficient way to achieve any expected group composition is a simple threshold rule.

De nition 2 *An admissions rule r is a threshold rule if there exist t_N and t_W in $[\underline{t}, \overline{t}]$, such that*
$$r_N(t) = \begin{cases} 0 & for \quad t \in [\underline{t}, t_N) \\ 1 & for \quad t \in [t_N, \overline{t}] \end{cases} \quad and \quad r_W(t) = \begin{cases} 0 & for \quad t \in [\underline{t}, t_W) \\ 1 & for \quad t \in [t_W, \overline{t}] \end{cases}$$

The admissions office assigns the thresholds t_N and t_W to the groups N and W, respectively. A candidate is admitted when her test score meets or exceeds her group s threshold, and rejected otherwise. We use R_{th} to denote the set of threshold rules and $r_{th}(t_N, t_W)$ to denote the threshold rule with cutoff levels t_N and t_W. With affirmative action, given a xed group composition, the admissions processes for the two groups are entirely unrelated and can be treated separately. We focus on W, but the same argument applies to N as well.

Admitting all W candidates with test score between t and $t + dt$ raises the value of the objective function by $w(t)\lambda(t)dt$ and lls $w(t)dt$ of the vacancies. Thus $\lambda(t)$ can be interpreted as the value per unit of the capacity constraint of group W candidates with test score t. Since by Assumption 1 $\lambda(t)$ is increasing in t, the admissions office strictly prefers a candidate with a higher test score to one with a lower test score. As a result, it admits students in descending order of test score, starting from $\overline{t}$, until all vacancies are lled. Since the optimal admissions rule to achieve *any* group composition is a threshold rule, the optimal admissions rule to achieve the optimal group composition must also be a threshold rule.

Whether the admissions office adopts affirmative action (by setting $t_N < t_W$) depends on its preference for diversity. If $\alpha = 0$ (the admissions office does not care about diversity), then since the marginal value of a candidate from either group scoring t is $\lambda(t)$, the admissions office uses the common-threshold rule $r_{th}(t_c, t_c)$, where a candidate is accepted if and only if she scores at least t_c, irrespective of her group identity. From Lemma 1,

$$\frac{N(r_{th}(t_c, t_c))}{N} < \frac{W(r_{th}(t_c, t_c))}{W},$$

which implies that

$$\frac{N(r_{th}(t_c, t_c))}{c} < \frac{N}{C}.$$

If the admissions office cares very much about diversity (i.e., α is large), it chooses the rule that admits the groups in proportion to their shares of the applicant pool (i.e., $\frac{N(r)}{c} = \frac{N}{C}$). Let $r_{th}(\tilde{t}_N, \tilde{t}_W)$ be this rule. In this paper, we focus on values of α between the two extremes:

Assumption 3 $0 < \frac{\alpha}{c} < \lambda(\tilde{t}_W) - \lambda(\tilde{t}_N)$.

In this case, the admissions office sets a higher threshold for group W, but group N is still under-represented.

Let $r_{th}(t_N, t_W)$ be the optimal admissions rule when affirmative action is allowed.

Proposition 1 *If Assumption 1 holds, then the optimal admissions rule is a threshold rule. If Assumptions 2 also holds, then* $\frac{N(r_{th})}{c}$ □ $\frac{N}{C}$. *If Assumption 3 also holds, then the admissions office adopts affirmative action (i.e.,* $t_N < t_W$*), but the minority group is strictly under-represented. Moreover, the thresholds* t_N *and* t_W *satisfy* $\lambda(t_W) - \lambda(t_N) = \frac{\alpha}{c}$.

The second part of Proposition 1 says that the minority group is weakly underrepresented with affirmative action. Since the majority group is on average more quali ed than the minority group, if the minority group were overrepresented, then the admissions office could simultaneously raise the entering class s quality and improve its diversity by admitting more group W candidates. The third part of Proposition 1 says that the thresholds are chosen such that the difference in quality of marginal candidates from the two groups, $\lambda(t_W) - \lambda(t_N)$, equals the marginal bene t of improving diversity, $\frac{\alpha}{c}$. By Assumption 3, this implies that t_W is strictly higher than t_N, and $\frac{N(r_{th})}{c}$ is strictly less than $\frac{N}{C}$.

The main implication of Proposition 1 is that affirmative action is the most efficient way to achieve any given level of diversity. Given Assumption 2, the minority group would be underrepresented under the common-threshold rule. To achieve a higher level of diversity, it is inevitable that some better-quali ed W candidates be displaced by some less-quali ed N candidates. Af- rmative action minimizes the cost by replacing the least-quali ed group W candidates with the most-quali ed group N candidates who otherwise would not be admitted.

Since the admissions office cares more about diversity than do the regents, its preferred admissions rule does not coincide with the regents . In fact, in our model, the regents prefer the common-threshold rule $r_{th}(t_c, t_c)$ to $r_{th}(t_N, t_W)$, the one chosen by the admissions office. However, this does not imply that the regents would be better off by banning affirmative action, because, rather than choosing $r_{th}(t_c, t_c)$, the admissions office may react to a ban by choosing some non-threshold rule which adversely affects within-group selection.

5 A Ban on Affirmative Action

When the regents ban affirmative action, the admissions office must use an admission rule that treats the two groups identically. Let R_{NA} be the set of such rules. To simplify notation, we use $r(t)$ to represent the probability that a candidate with test score t is admitted under admissions rule r. The admissions office s task is to choose r to maximize

$$\int_{\underline{t}}^{\bar{t}} r(t)\lambda(t)\left(n(t)+w(t)\right)dt - \alpha\left|\frac{N}{C} - \frac{\int_{\underline{t}}^{\bar{t}} r(t)n(t)dt}{c}\right|,$$

subject to the capacity constraint. Recall that $r(t)$ is required to be increasing in t, meaning that the higher a candidate s test score, the higher her chance of gaining admission. By Lemma 1, the minority group is underrepresented under any increasing admissions rule. Thus, we can ignore the absolute-value sign in the objective function and drop the constant term $\frac{\alpha N}{c}$ to rewrite the admissions office s problem as follows:

$$\begin{aligned} \max_{r\in R_{NA}} \quad & \int_{\underline{t}}^{\bar{t}} r(t)\lambda(t)\left(n(t)+w(t)\right)dt + \frac{\alpha}{c}\int_{\underline{t}}^{\bar{t}} r(t)n(t)dt \\ \text{s. t.} \quad & \int_{\underline{t}}^{\bar{t}} r(t)\left(n(t)+w(t)\right)dt = c. \end{aligned}$$

Without affirmative action, the admissions processes for the two groups are combined: if the admissions office admits minority candidates scoring t, it also must admit majority candidates with the same score.

Let

$$\gamma(t) \equiv \frac{\lambda(t)(n(t)+w(t)) + \frac{\alpha}{c}n(t)}{n(t)+w(t)} = \lambda(t) + \frac{\alpha}{c}\psi(t),$$

the value per unit of the capacity constraint of candidates scoring t; that is, $\gamma(t)$ is how much the admissions office likes candidates scoring t. Unlike $\lambda(t)$, the quality of candidates scoring t, $\gamma(t)$ incorporates the admissions office s taste for diversity: *ceteris paribus*, the higher $\psi(t)$, the share of candidates scoring t belonging to the minority group, the more the admissions office likes candidates scoring t.

When γ is everywhere increasing in t, the optimal admissions rule is a threshold rule. However, γ need not be everywhere increasing. To see this, note that $\gamma'(t) = \lambda'(t) + \frac{\alpha}{c}\psi'(t)$. Since λ, the quality of candidates, is increasing by Assumption 1, whereas ψ, the share of candidates from the minority group, is decreasing by Assumption 2, whether γ increases at t depends on how fast ψ changes relative to λ at t. If in the neighborhood of some test score t, λ changes very little while ψ changes somewhat more, then γ may be decreasing at t.[24] Suppose that for some $t' < t''$, $\gamma(t') > \gamma(t'')$: the admissions office prefers candidates at t' to higher-scoring candidates at t''. If the size of the entering class is insufficient to admit candidates scoring t' using a threshold rule, the admissions office may adopt a random admissions rule. Furthermore, if for each $t \in [t', t'']$, $\gamma(t) < \gamma(t')$, the admissions office will not admit those scoring t with strictly higher probability than those scoring t', and thus an optimal rule is flat between t' and t''. An admissions rule with flat spots allows the admissions office to admit the lower-scoring candidates it likes without admitting any more higher-scoring candidates than necessary under the monotonicity restriction.

Below we show that at least one of the admissions office's preferred rules under a ban belongs to the following simple class of random rules:

Definition 3 (Two-Step Admissions Rule) *An admissions rule $r \in R_{NA}$ is a two-step admissions rule if there exist t_1 and t_2, $\underline{t} \le t_1 \le t_2 \le \overline{t}$, such that*

$$r(t) = \begin{cases} 0 & \text{for} \quad t \in [\underline{t}, t_1) \\ p & \text{for} \quad t \in [t_1, t_2) \\ 1 & \text{for} \quad t \in [t_2, \overline{t}] \end{cases}$$

where $p \equiv \frac{c - \int_{t_2}^{\overline{t}} n(t) + w(t) dt}{\int_{t_1}^{t_2} n(t) + w(t) dt}$.

Let $r_{2s}(t_1, t_2)$ denote a two-step rule with cutoffs t_1 and t_2, and R_{2s} denote the set of all two-step rules.[25] Under the two-step rule $r_{2s}(t_1, t_2)$, the admissions office admits all candidates scoring at least t_2, conducts a lottery over candidates scoring between t_1 and t_2 that gives each of them an equal chance, p, of admission, and rejects all candidates scoring below t_1. Despite the name, a two-step rule need not have two distinct steps: when t_1 equals t_2, the rule reduces to a common-threshold rule; when t_2 equals $\overline{t}$, the rule becomes a single-step rule under which all candidates with test scores greater than or equal to t_1 have an equal chance of admission. We use $r_{th}(t_c)$ to denote the threshold rule with cutoff t_c and $r_{1s}(t_1)$ to denote the one-step rule with cutoff t_1.

[24] For example, suppose t ranges from 0 to 100, $\lambda(t) = \frac{1}{3}(t-50)^3$, and $\psi(t) = 1 - \frac{t}{100}$. The function λ is concave for low test scores ($t \le 50$) and convex for high test scores ($t \geq 50$). In this case, γ has a local maximum at $t_a = 50 - \frac{1}{10}\sqrt{\frac{\alpha}{c}}$ and a local minimum at $t_b = 50 + \frac{1}{10}\sqrt{\frac{\alpha}{c}}$.

[25] Because t_1 and t_2 determine a unique p, we omit p from our notation.

Consider the case where γ has two critical points: a local maximum at t_a and a local minimum at t_b, where $t_a < t_b$. (See Figure 1.)

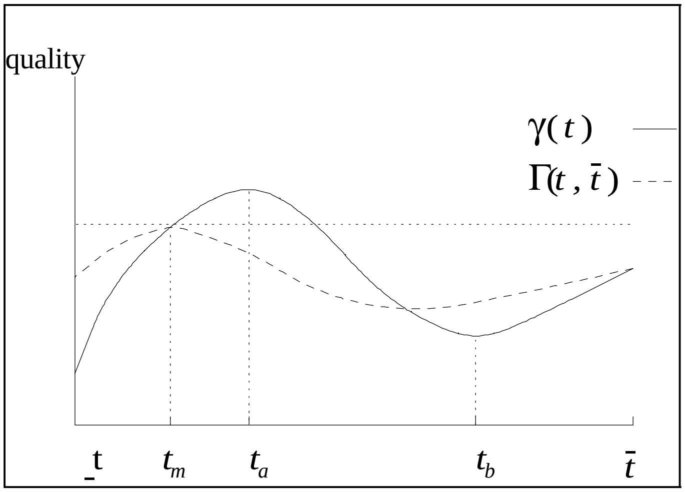

Figure 1

De ne the function $\Gamma : \{(t_1, t_2) : t_1 \in [\underline{t}, \overline{t}],\ t_2 \geq t_1\} \to \mathbb{R}$ as follows:

$$\Gamma(t_1, t_2) = \begin{cases} \frac{\int_{t_1}^{t_2} \gamma(t)(n(t)+w(t))dt}{\int_{t_1}^{t_2} n(t)+w(t)dt} & \text{for} \quad t_1 < t_2 \\ \gamma(t_1) & \text{for} \quad t_1 = t_2. \end{cases}$$

$\Gamma(t_1, t_2)$ is the average value of γ over (t_1, t_2). If $\Gamma(t', \overline{t}) > \Gamma(t'', \overline{t})$, then the admissions office prefers a randomly selected candidate scoring in $[t', \overline{t}]$ to one scoring in $[t'', \overline{t}]$. Suppose, as in Figure 1, that $\Gamma(t, \overline{t})$ reaches its global maximum at t_m. When the entering class is too small to admit all candidates scoring at least t_m, the admissions office conducts a lottery over these candidates, giving each of them an equal chance of admission. When the size of the entering class is large enough to admit all of candidates scoring above t_m, the admissions office lls any remaining seats with candidates scoring below t_m, in descending order of test score, as γ is increasing below t_m. Thus, if the number of seats is larger than the number of candidates scoring above t_m, the optimal admissions rule is a threshold rule. But if the number of seats is smaller than the number of candidates scoring above t_m, then the optimal admissions rule is a one-step rule, and each such candidate is admitted with probability $p \in (0, 1)$.

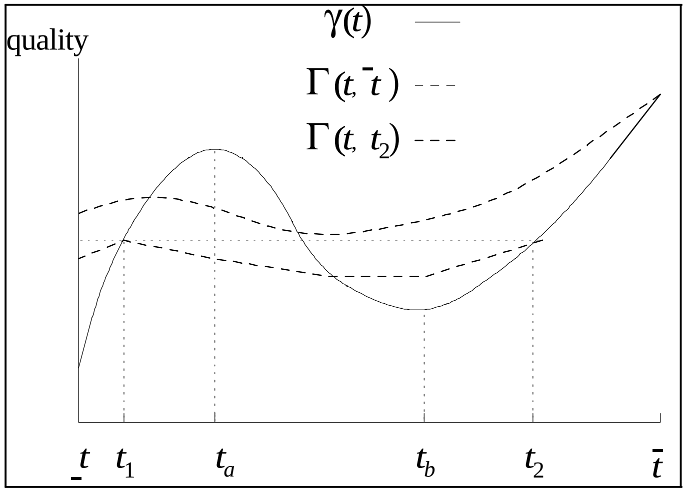

Figure 2

Figure 2 differs from Figure 1 in that $\Gamma(t, \bar{t})$ reaches its maximum at the boundary $\bar{t}$. Let $f(t')$ be the interior maximizer of the function $\Gamma(t, t')$. Consider the function $g(t') \equiv \gamma(t') - \Gamma(f(t'), t')$. From Figure 2, it is clear that $g(\bar{t}) > 0$ and $g(t_b) < 0$. Since g is continuous, there must exist some $t_2 \in (t_b, \bar{t})$ such that $g(t_2) = 0$. Furthermore, it can be shown that for any $t' \in (t_2, \bar{t})$, $g(t') > 0$; that is, the interior maximum of $\Gamma(t, t')$ is less than $\gamma(t')$. This means that for any $t' \in [t_2, \bar{t}]$, the admissions office prefers a candidate scoring t' to one randomly drawn from any $[t'', t') \subset [\underline{t}, t')$. Thus, the admissions office rst admits candidates scoring $t \geq t_2$ in descending order of test score. To ll any remaining seats, the admissions office uses the procedure described in the last example: it rst randomly admits candidates scoring $t \in [t_1, t_2]$, and then lls any remaining seats with candidates not yet admitted, in descending order of test score. The optimal admissions rule is therefore a threshold rule if $\int_{t_2}^{\bar{t}} n(t) + w(t)dt \geq c$ or if $\int_{t_1}^{\bar{t}} n(t) + w(t)dt \square c$, and a proper two-step rule if $\int_{t_1}^{\bar{t}} n(t) + w(t)dt > c > \int_{t_2}^{\bar{t}} n(t) + w(t)dt$.

The above examples illustrate a general result: whenever γ is not monotone, there is an interval $[t_1, t_2] \subset [\underline{t}, \bar{t}]$ such that the admissions office prefers a candidate randomly selected from that interval to one selected from any interval $[t', t_2] \subset [t_1, t_2]$. When the number of seats lies between the number of candidates scoring above t_2 and the number of candidates scoring above t_1, the admissions office chooses a random admissions rule .

The admissions office is forced to randomize because r is constrained to be increasing in t. Without the monotonicity constraint, the admissions office always admits candidates who have the highest γ, and the admissions rule is deterministic. But when γ is not monotone, the admissions may admit some candidates and reject others with higher test scores, and r will not be monotone in t. By forcing the admissions office to accept higher-scoring candidates whenever it accepts lower-scoring candidates, the monotonicity constraint increases the likelihood that the admissions office's preferred rule under a ban is a threshold rule.

The assumption that the admissions office's utility function is linear in diversity, coupled with the result that the minority group is underrepresented under any admissions rule, means that the admissions office's preferences over individual candidates do not depend on the overall group composition of the entering class. Thus, the admissions office's preferences over entering classes can be decomposed into preferences over individual candidates without referring to a specific admissions rule. This, however, would not be true if the admissions office's preferences over diversity were not linear. In that case, the admissions office may still randomize to promote diversity, but its preferred rule might not take any simple form.

Thus far, we have not ruled out the optimality of non-two-step rules. In the appendix, we prove that when γ has no more than two critical points, the optimal admissions rule is unique, and it must be a two step rule.[26] Generally, when γ has more than two critical points, there may be multiple optimal rules. However, we can show that at least one optimal rule must be a two-step rule.[27]

Proposition 2 *There exists a two-step admissions rule that is optimal in R_{NA}. Moreover, when γ has no more than two critical points, the optimal admissions rule is unique.*

Here we provide an outline of the proof. First, we show that optimal rules exist. Next, we show that if γ has no more than $2k-3$ critical points, an optimal rule must be a step rule containing at most k steps. Finally, we show that for any optimal rule, there exists a two-step rule that the admissions office likes just as well. Suppose that a k-step rule under which each candidate scoring between t_j and t_{j+1} has probability $p_j > 0$ of gaining admission is optimal, where $j \in \{1, ..., k-1\}$. If $\Gamma(t_1, t_2) \leq \Gamma(t_2, t_3)$, the admissions office can (weakly) increase its utility by reducing p_1 and increasing p_2 until either $p_1 = 0$ or $p_2 = p_3$. On the other hand, if $\Gamma(t_1, t_2) > \Gamma(t_2, t_3)$, the admissions office can increase its utility by increasing p_1 and reducing p_2 until $p_1 = p_2$. In either case, it has come up with a $(k-1)$-step rule at least as good as its original k-step rule. Repeating the process yields a two-step rule that is weakly superior to the original k-step rule. Intuitively, the admissions office must be indifferent over all steps admitted with probability less than one,

[26] The set of γ's critical points is $\{t : \gamma'(t) = 0\}$. Strictly speaking, a one-step rule is only unique up to $r(\bar{t})$.

[27] We conjecture, but have not proven, that generically two-step rules are uniquely optimal in the sense that fixing all other parameters of the model, the set of values of α in $(0, c(\lambda(\tilde{t}_W) - \lambda(\tilde{t}_N)))$ for which non-two-step rules are optimal is closed and at most countably inifnite.

otherwise it could make itself better off by increasing the probability of acceptance for its preferred step; but if it is indifferent between two adjacent steps, it can merge them into a single step.

Proposition 2 allows us to focus exclusively on two-step rules. Figures 1 and 2 illustrate the necessary conditions for optimal two-step rules. For the one-step rule $r_{1s}(t^*)$ to be optimal, $\Gamma(t, \bar{t})$ must reach its maximum at t^*. For the proper two-step rule $r_{2s}(t_1, t_2)$ to be optimal, $\gamma(t)$ must be increasing at t_2 and $\Gamma(t, t_2)$ must reach the maximum at t_1. Since $\Gamma(t, t')$ is the average of γ over (t, t'), it is increasing in t when $\Gamma(t, t') > \gamma(t)$, decreasing in t when $\Gamma(t, t') < \gamma(t)$, and thus it reaches the maximum when it intersects γ from below. The necessary conditions for two-step rules are summarized by the following lemma:

Lemma 2 *If the one-step rule $r_{1s}(t_1)$ is optimal in R_{NA} , then $\gamma(t_1) = \Gamma(t_1, \bar{t})$ and $\gamma'(t_1) \geq 0$. If the proper two-step rule $r_{2s}(t_1, t_2)$ is optimal in R_{NA}, then $\gamma(t_1) \geq \Gamma(t_1, t_2) = \gamma(t_2)$, $\gamma'(t_1) \geq 0$, and $\gamma'(t_2) \geq 0$.*

Note that for an optimal proper two-step rule, $\gamma(t_1) > \Gamma(t_1, t_2)$ only if $t_1 = \underline{t}$; if $t_1 > \underline{t}$ and $\gamma(t_1) > \Gamma(t_1, t_2)$, then the admissions office can increase its utility by lowering t_1. It follows from Proposition 4 below that for a one-step rule, $t_1 \neq \underline{t}$, so $\gamma(t_1) = \Gamma(t_1, \bar{t})$.

The necessary conditions re ect the trade-off between quality and diversity. They can be rewritten as

$$\Lambda(t_1, t_2) - \lambda(t_1) \geq \frac{\alpha}{c} \left(\psi(t_1) - \Psi(t_1, t_2) \right), \text{ and}$$

$$\lambda(t_2) - \lambda(t_1) \square \frac{\alpha}{c} \left(\psi(t_1) - \psi(t_2) \right),$$

where $t_2 = \bar{t}$ for an optimal one-step rule; the rst condition holds with equality for a one-step rule or a proper two-step rule with $t_1 \neq \underline{t}$, and the second condition holds with equality for a proper two-step rule with $t_1 \neq \underline{t}$. The rst condition states that the diversity gain from admitting a candidate with test score t_1 over the average candidate in (t_1, t_2) is no larger than the quality loss, and the second condition states that the diversity gain from admitting a candidate with test score t_1 over a candidate scoring t_2 is no larger than the quality loss. Intuitively, admitting candidates with lower test scores improves minority enrollment but lowers average student quality; the two conditions show that at the margin, the quality loss is exactly offset by the diversity gain.

Banning affirmative action can cause the admissions office to inefficiently use its test. For example, by using the proper two-step rule $r_{2s}(t_1, t_2)$, the admissions office admits every candidate scoring in (t_1, t_2) with the same probability, which amounts to ignoring test score in that range. Because quality is increasing in test score, the less the admissions office relies on its test to admit candidates, the less quali ed the entering class. The loss in quality increases with the size of the gap $\lambda(t_2) - \lambda(t_1)$ which, from the necessary conditions, is proportional to $\psi(t_1) - \psi(t_2)$, the change in group composition between t_1 and t_2. Roughly speaking, the quality loss from randomization tends to be larger when group composition changes rapidly with test scores.

The following proposition underscores the inefficiency in randomization: both parties would prefer an appropriately constructed affirmative-action threshold rule (like the ones in the last section) to any non-threshold rule $r \in R_{NA}$, including r^*, the admissions office s optimal rule under the ban on affirmative action.

Proposition 3 *For each $r \in R_{NA}$ that involves randomization, there exists an affirmative-action admissions rule $\tilde{r}$ that both the admissions office and regents prefer to r.*

The proof of the proposition follows directly from the proof of Proposition 1. Consider an affirmative-action threshold rule, $\widetilde{r}$, achieving the same diversity as r. We know that $\widetilde{r}$ must generate better within-group selection than r, so total student quality is higher under $\widetilde{r}$ than r; thus, the regents prefer $\widetilde{r}$ to r. But since both rules achieve the same group composition, the admissions office also prefers $\widetilde{r}$ to r.

Although Proposition 3 says that there exists an affirmative-action rule that both parties prefer to the admissions office s preferred rule under a ban, it does not say that by lifting their ban the regents can induce the admissions office to implement $\widetilde{r}$. In Section 7, however, we show that under some conditions an implementable partial ban on affirmative action can dominate a complete ban.

6 Comparing Regimes

When the admissions office reacts to a ban by using the common-threshold rule, banning affirmative action reduces diversity and raises average student quality. But when γ is not increasing in t, the admissions office may respond to a ban by choosing a random admissions rule rather than the common-threshold rule. Compared to the common-threshold rule, an optimal random rule improves diversity but lowers the average quality of the entering class. Thus, randomization works against the interest of the regents. In this section, we compare the characteristics of the entering classes under affirmative action and under optimal two-step rules. Proposition 4 says that the admissions office would never react to a ban by randomizing enough to maintain the optimal level of diversity under affirmative action.

Proposition 4 *Under Assumptions 1-3, banning affirmative action strictly lowers minority enrollment.*

A complete proof is provided in the appendix. Here we show that the proposition holds when the optimal rule under a ban is a two-step rule.[28] The rst-order conditions for optimal affirmative action and optimal two-step rules imply that

$$\lambda(t_2) - \lambda(t_1) \leq \frac{\alpha}{c}(\psi(t_1) - \psi(t_2)) < \lambda(t_W) - \lambda(t_N).$$

Hence, the interval $[t_N, t_W]$ cannot be a subset of $[t_1, t_2]$: either t_1 is larger than t_N, or t_2 is smaller than t_W. If t_1 is larger than t_N, then no minority candidate not admitted under affirmative action is admitted under the optimal two-step rule, while some of those admitted under affirmative action are not admitted under the two-step rule; hence, minority enrollment goes down. If t_2 is smaller than t_W, then all majority candidates admitted under affirmative action are admitted under the two-step rule, while some of those not admitted under affirmative action are admitted under the two-step rule. In this case, majority enrollment goes up, meaning that minority enrollment must go down. Intuitively, the proposition holds because the admissions office would never replace a majority candidate admitted under affirmative action by a minority candidate not admitted under affirmative action.

Given that the gap in quality between t_1 and t_2 is smaller than that between t_N and t_W, one might suspect that the effects of randomization are sufficiently small that banning affirmative action always raises total quality. While the total quality of the entering class always increases when (t_1, t_2) is contained in (t_N, t_W), the following example shows that when (t_1, t_2) is not contained in (t_N, t_W) a ban may lower quality.[29]

Example 1 *Let* $[\underline{t}, \overline{t}] = [0, 80]$, $n(t) = e^{-.17(t+18)}$, $\psi(t) = 10^{-6} + (1 - 10^{-6})e^{-\left(\frac{t+18}{47}\right)^4}$, $\lambda(t) = \frac{(t-17)^3}{80} - 2210$, $\frac{\alpha}{c} = 340$, $\frac{c}{C} = .8$, *and* $\frac{N}{C} = .1955$.

Figure 3 shows the distribution of candidates for the two groups: $n(t)$ is represented by the solid line, and $w(t)$ by the dotted line.

[28] Recall that some optimal rule may contain more than two steps.

[29] When $(t_1, t_2) \subset (t_N, t_W)$, the average quality of those candidates displaced by the ban affirmative action is $\delta\Lambda_N(t_N, t_1) + (1 - \delta)\Lambda_N(t_1, t_2)$ for some $\delta \in [0, 1]$, which is no greater than $\Lambda_N(t_1, t_2)$, while the average quality of those candidates replacing them is $\beta\Lambda_W(t_1, t_2) + (1 - \beta)\Lambda_W(t_2, t_W)$ for some $\beta \in [0, 1]$, which is no smaller than $\Lambda_W(t_1, t_2)$. By Lemma 1, $\Lambda_W(t_1, t_2) > \Lambda_N(t_1, t_2)$, so average quality improves by banning affirmative action.

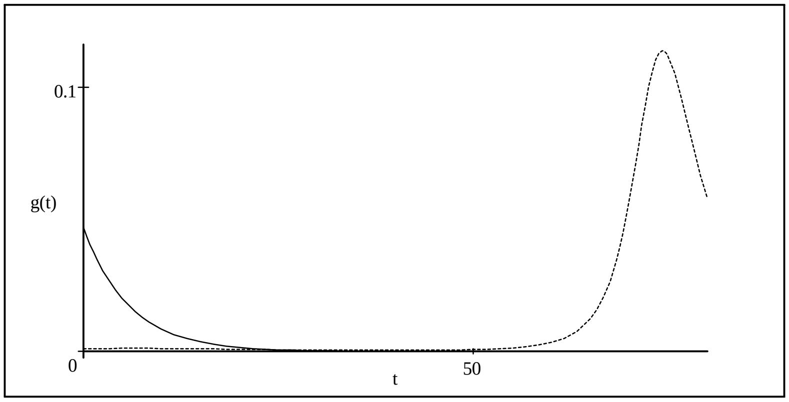

Figure 3

Under affirmative action, $(t_N, t_W) = (12.9, 47.0)$: minority candidates make up 2.7 percent of the entering class, and total quality is 0.94. Figure 4 shows $\gamma(t)$.

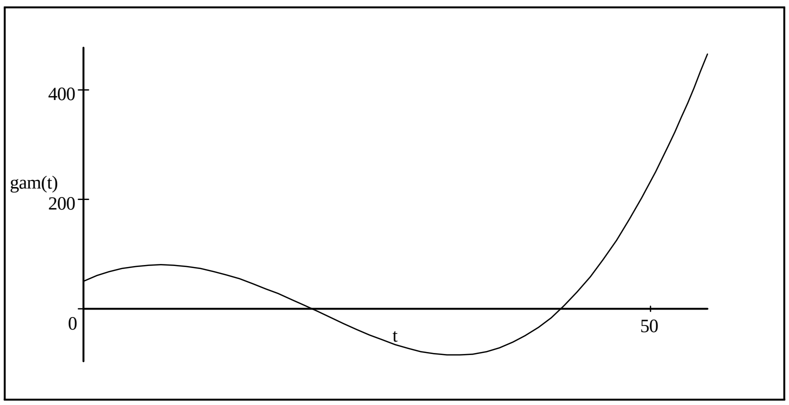

Figure 4

The admissions office responds to a ban by using $r_{2s}(1.0, 44.7)$, where $p = 0.11$: minority candidates make up 2.3 percent of the entering class, and total quality is 0.72. Thus banning affirmative action lowers total quality. Under the common-threshold rule, $t_c = 17.8$, minority candidates make up 1.2 percent of the entering class, and total quality is 2.3.

A ban on affirmative action causes the admissions office to replace minority candidates scoring in (t_N, t_2) with majority candidates scoring in (t_2, t_W) and fraction p of all candidates scoring in (t_1, t_2). Thus, the total quality of the entering class declines under a ban if

$$\Lambda_N(t_N, t_2) > \delta\Lambda_W(t_2, t_W) + (1 - \delta)\Lambda(t_1, t_2),$$

where $\delta \equiv \frac{\int_{t_2}^{t_W} w(t)dt}{p\left(\int_{t_1}^{t_2} n(t)+w(t)dt\right)+\int_{t_2}^{t_W} w(t)dt}$. Note that $\Lambda_W(t_2, t_W)$, the average quality of majority candidates in (t_2, t_N), is higher than $\Lambda_N(t_N, t_2)$, the average quality of minority candidates in (t_N, t_2).

But $\Lambda(t_1, t_2)$, the average quality of candidates in (t_1, t_2), can be smaller than $\Lambda_N(t_N, t_2)$, as $t_1 < t_N$. Thus, whether the regents are better off banning affirmative action depends on δ, the fraction of the slots going to majority candidates in (t_2, t_W), and $\Lambda(t_1, t_2)$, the average quality of candidates between (t_1, t_2). In Example 1, since λ drops off rapidly and the density of candidates increases sharply as t goes under t_N, $\Lambda(t_1, t_2)$ is less than $\Lambda_N(t_N, t_2)$. Moreover, since there are virtually no majority candidates in (t_2, t_W), δ is small. That is, the admissions office essentially replaces minority candidates scoring in (t_N, t_2) with candidates in (t_1, t_2). As a result, the total quality declines when affirmative action is banned.

The above example was constructed speci cally to show that the total quality of the entering class may fall when affirmative action is banned. In general, the effect of randomization may be less extreme, but its existence by no means depends on the particular functional forms we use in the example. To provide an illustration of less severe randomization, we calibrate an example of our model by approximating the data on applicants to ve selective schools in 1989 reported in Bowen and Bok [1998]. Bowen and Bok provide mean SAT scores for African American and white candidates (1098 and 1284, respectively), and their Figure 2.2 illustrates the distributions of SAT scores for the two groups. For our example, we take African Americans to be the minority group, and whites to be the majority group. Using Figure 2.2, we estimate the standard deviations for the two groups to be 150 and 130 points, respectively; we assume each group s test score is truncated normal (which is broadly consistent with the data) with its reported mean and our estimated variance, and take $[\underline{t}, \overline{t}] = [400, 1600]$. Finally, we assume ten percent of the applicant pool belongs to the minority group, and the rest to the majority group.[30] We pick a λ that is concave when t is small, and convex when t is larger. That λ has a at part is important for γ to be non-monotone, but the exact functional form of λ is not crucial. Given λ, we set $\frac{\alpha}{c} = 50$, so the gap between the thresholds for the two groups under affirmative action is roughly consistent with Kane s [1998] estimate.

Example 2 *Let* $[\underline{t}, \overline{t}] = [400, 1600]$, $f_N(t) = \frac{k_1}{150\sqrt{2\pi}} e^{-\frac{1}{2}\left(\frac{t-1098}{150}\right)^2}$, $f_W(t) = \frac{k_2}{130\sqrt{2\pi}} e^{-\frac{1}{2}\left(\frac{t-1284}{130}\right)^2}$, $\lambda(t) = 10^{-14}(t-1350)^7$, $\frac{\alpha}{c} = 50$, $\frac{c}{C} = .31$, $\frac{N}{C} = .1$, *for* $k_1 = 1.00041$ *and* $k_2 = 1.00759$.

[30] Bowen and Bok report that their data set includes over 40,000 applications, of which 2,300 came from candidates who self-identi ed as black. Presumably a majority, but certainly not all, of the remaining applications come from whites, so we choose the ratio of whites to blacks in our example to be nine to one. We remind the reader that this example is meant to illustrate that randomization can occur with plausible test score distributions, and it is not meant to simulate the admissions process at any school.

Figure 5 shows test-score distributions for the two groups, where the minority group is the solid line, and the majority group the dashed line.

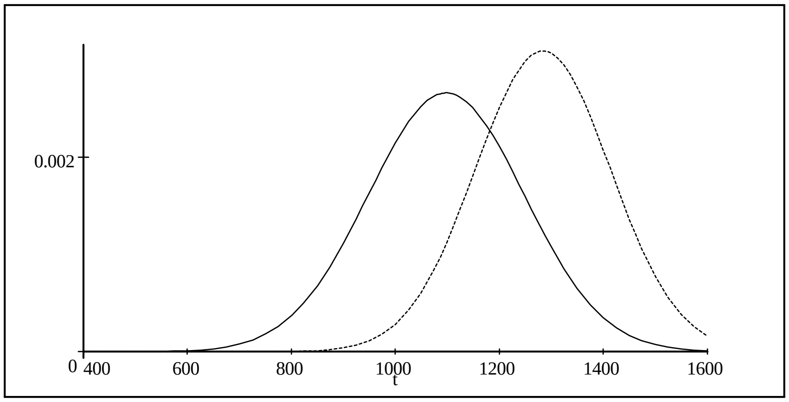

Figure 5

The affirmative action solution is $(t_N, t_W) = (1175, 1346)$: the marginal majority candidate scores 171 points higher than the marginal minority candidate. Minority candidates make up 9.8 percent of the entering class, which is slightly less than their ten-percent share of the candidate pool.

Figure 6 shows $\gamma(t)$.

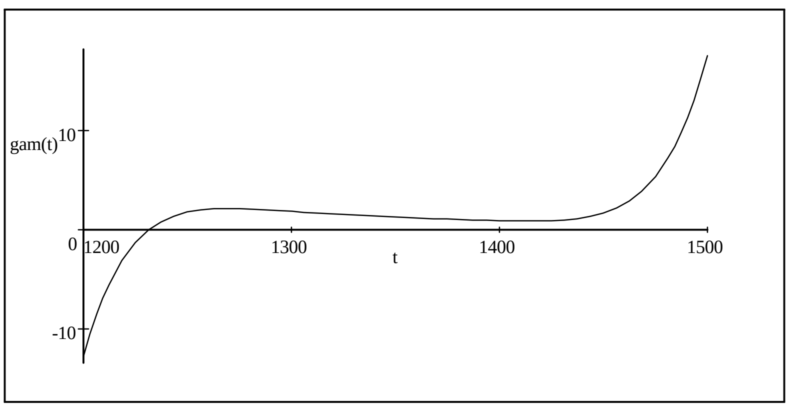

Figure 6

When affirmative action is banned, the admissions office uses the two-step rule $r_{2s}(1245, 1447)$, where $p = 0.46$. Minority candidates now comprise only 2.6 percent of the entering class. However, were the admissions office to use the common-threshold rule (where $t_c = 1336$), the minority group would make up only 1.8 percent of the entering class. Banning affirmative action increases the quality of the entering class: under affirmative action, total quality is 47.0, while it is 49.0 under a ban. (Under the common-threshold rule, total quality is 49.3.) Relative to affirmative action, the admissions office s optimal two-step rule displaces all minority candidates scoring in $(1175, 1245)$, a fraction $1-p$ of those scoring in $(1245, 1447)$, and a fraction $1-p$ of majority candidates scoring in $(1346, 1447)$, replacing them with fraction p of majority candidates scoring in $(1245, 1346)$. Even though most of these displaced candidates come from the majority group and score in $(1346, 1447)$, λ is sufficiently at in that range that the regents care very little about replacing majority candidates in that range with majority candidates in $(1245, 1346)$. What makes the regents better off under a ban is that minority candidates scoring in $(1175, 1245)$ are displaced: λ is very steep for t sufficiently far from 1350.

Thus far, we have focused on the admissions office s trade-off between quality and diversity. Some commentators have also argued that affirmative action has widened the average test-score gap between majority and minority matriculants at elite colleges. (See Jencks and Phillips [1998] for discussion of several facets of racial test-score gaps.) This may harm minority matriculants in two ways. First, some minority candidates who would not have been admitted without affirmative action may have difficulty ful lling school requirements and drop out; they may be better off attending less selective colleges where they can compete successfully. But even those minority candidates who would have been admitted without affirmative action may suffer if they are stereotyped as having lower test scores than they do. As a result, affirmative action may actually harm the minority group as a whole. (See Loury [1987] for this argument.)

In our model, because affirmative action sets a higher threshold for the majority group than for the minority group, the average test-score gap is wider under affirmative action than under the common-threshold rule. Thus, if the admissions office responds to a ban by using the common-threshold rule, banning affirmative action narrows the test-score gap. However, when the admissions office responds to a ban by randomizing, the gap in average test scores may rise, exacerbating any stereotyping problem. Proposition 5 characterizes how the effect of a ban on this gap depends on the admissions office s optimal rule.

Proposition 5

1. *If the admissions office s preferred rule in R_{NA} is the common threshold rule $r_{th}(t_c)$ or the single-step rule $r_{1s}(t_1)$, then the average test-score gap shrinks under a ban on affirmative action.*

2. *If the admissions office s preferred rule in R_{NA} is the proper two-step rule $r_{2s}(t_1, t_2)$, then (i) when $t_1 \geq t_N$, the average test-score gap shrinks under a ban on affirmative action, and (ii) when $t_1 < t_N$, the average test-score gap can either grow or shrink under a ban on affirmative*

action.

From Proposition 4 we know that banning affirmative action increases the majority group s representation in the entering class. Since affirmative action admits the highest-scoring majority candidates, banning affirmative action must lower the average test score of candidates admitted from the majority group. The rst case of Proposition 5 illustrates a common intuition about banning affirmative action: if the admissions office reacts to a ban on affirmative action by using a simple threshold rule, the average test score for the minority group rises, while that for the majority group falls, causing the gap to shrink. If the admissions office uses the single-step rule $r_{1s}(t_1)$, then since it prefers admitting candidates scoring above t_W to those scoring below t_N, $t_1 > t_N$. Thus, the average test score for minority matriculants must rise, and since the average test score of majority matriculants falls, the test-score gap shrinks.

Part 2 describes what happens to the test-score gap when the admissions office uses a proper two-step rule. If $t_1 \geq t_N$, then, just as under a single-step rule, the minority group s average test score rises, causing the gap to shrink. More interestingly, in the second case, the effect of a ban on the test-score gap is ambiguous: while the average test score for majority group candidates falls, the average test score for the minority group may also fall. If the latter falls enough, then the test-score gap might actually increase under a ban on affirmative action. Intuitively, if the admissions office randomizes enough, and admits enough low-scoring candidates, since the minority group is disproportionately composed of these low-scoring candidates, its average test score may fall more than the majority group s.

In Example 2, $t_1 > t_N$, so the second part of the proposition states that the gap in average test scores for the two groups shrinks when the regents ban affirmative action. Under affirmative action, the gap is 155 points, while it is only 48 under a ban. Under the common-threshold rule, the gap is only 20 points. In Example 1, however, banning affirmative action increases the gap in average test scores. Under affirmative action the gap is 55 points, but it rises to 66 points when affirmative action is banned. In this example, the mean test score of admitted minority candidates falls from 18 to 7 when affirmative action is banned, while the mean test score of admitted majority candidates falls negligibly.

7 A Partial Ban on Affirmative Action

The last section showed that banning affirmative action can cause the quality of the entering class to either rise or fall. In this section, we show that even when student quality rises with a ban, depending on the type of rule the admissions office uses, the regents may prefer some type of partial ban to a total ban. Speci cally, the regents may prefer to allow the admissions office to ll part of the entering class with a rule that uses group identity, but require that the balance of the class be lled using a group-blind rule. (Throughout this section, we use the terms partial ban and limited affirmative action interchangeably, and refer to the size of the fraction of the seats that the admissions office can ll however it wants as the size of limited affirmative action.)

Limited affirmative action can take many forms. One particularly simple form is for the regents to allow the admissions office to ll a fraction $x \in [0, 1]$ of the seats in the entering class however it wants, but require that the remaining $(1 - x)c$ seats be lled using an admissions rule that ignores group identity.[31] By setting $x = 0$ the regents can ban affirmative action, and by setting $x = 1$ they can allow complete affirmative action. In this section, we show that even when the regents care only about the quality of the entering class, and prefer a total ban on affirmative action to no ban, they may prefer some affirmative action (setting $x > 0$) to none (setting $x > 0$).

Throughout this section, we consider only the case where x is small. In this case, we can assume without loss of generality that the admissions office uses a two-step rule to ll the $(1 - x)c$ seats.[32] Because when $x = 0$ the minority group is strictly underrepresented, when x is small the admissions office will ll its xc affirmative-action slots entirely with minority candidates. As a result, the admissions office maximizes the objective function

$$p \int_{t_1}^{t_2} \left(\lambda(t) + \frac{\alpha}{c}\right) n(t) + \lambda(t)w(t)dt + \int_{t_2}^{\overline{t}} \left(\lambda(t) + \frac{\alpha}{c}\right) n(t) + \lambda(t)w(t)dt + (1-p) \int_{\widehat{t}}^{t_2} \left(\lambda(t) + \frac{\alpha}{c}\right) n(t)dt$$

subject to the capacity constraints

$$p \int_{t_1}^{t_2} n(t) + w(t)dt + \int_{t_2}^{\overline{t}} n(t) + w(t)dt = (1-x)c$$

$$(1-p) \int_{\widehat{t}}^{t_2} n(t)dt = xc.$$

With the exception of $\widehat{t}$, the notation is the same as before: candidates scoring above t_2 are accepted with probability one; candidates scoring in $[t_1, \widehat{t})$ are accepted with probability p; and candidates scoring in $[\widehat{t}, t_2)$ are accepted with probability one if they belong to the minority group, and probability p if they belong to the majority group. Throughout this section, we assume that there is an unique, optimal two-step rule when $x = 0$. Let t_1^*, t_2^*, p^* denote this solution, and $t_1(x)$, $t_2(x)$, $\widehat{t}(x)$ and $p\,(x)$ denote the unique solution to the problem for x in the neighborhood of zero.

[31] The order in which the admissions office lls the class does not matter for the formal analysis.

[32] See Lemma 7 in Appendix B for details.

Note that t_1 and t_2 are functions of x. As before, when setting its admissions rule, the admissions office must choose between quality and diversity: more randomization admits more minority candidates, but rejects more high-scoring candidates. But by allowing the admissions office to admit the xc highest-scoring minority candidates not admitted under the main admissions scheme, limited affirmative action lessens the negative quality effect of randomization. As a result, the admissions office may adopt a more random admissions rule under limited affirmative action than under a total ban.[33] In the context of two-step rules, more randomization means that $\frac{dt_1}{dx}$ is likely to be negative the more limited affirmative action, the larger the range of candidates admitted with positive probability and $\frac{dt_2}{dx}$ likely to be positive the more limited affirmative action, the smaller the range of candidates admitted with probability one.

When x is small, the optimal two-step rule is close to the optimal two-step rule for $x = 0$, and so the regents utility changes continuously

with x in the neighborhood of $x = 0$. Therefore, if the regents utility increases with x in the neighborhood of $x = 0$, then they must be better off under a partial ban than under a complete ban. The following lemma decomposes the change in regents welfare with respect to x into three parts.

Lemma 3 *The rate of change of the regents utility with respect to x when $x = 0$ is given by:*

$$\begin{aligned}\left.\frac{dU^R}{dx}\right|_{x=0} &= c(\lambda(t_2^*) - \Lambda(t_1^*, t_2^*)) + p^*(\Lambda(t_1^*, t_2^*) - \lambda(t_1^*))(n(t_1^*) + w(t_1^*)) \left.\frac{dt_1}{dx}\right|_{x=0} + \\ &\quad (1 - p^*)(\Lambda(t_1^*, t_2^*) - \lambda(t_2^*))(n(t_2^*) + w(t_2^*)) \left.\frac{dt_2}{dx}\right|_{x=0}.\end{aligned}$$

The rst term represents the direct substitution effect of limited affirmation action, and the second and third terms represent the indirect effects through t_1 and t_2. Since $\lambda(t_2^*)$, the quality of a minority candidate admitted through affirmative action, is higher than $\Lambda(t_1^*, t_2^*)$, the average quality of the candidate she replaces, the direct substitution effect of limited affirmative action is always positive. The second and third terms are negative because of the randomization effect. Whether the regents want to use some limited affirmative action depends on which effect is stronger.

Consider the case where the optimal rule is a one-step rule in the neighborhood of $x = 0$. In that case, $t_2 = \bar{t}$ and $\frac{dt_2}{dx} = 0$. The rst-order conditions of the admissions office s problem imply

[33] Note that this statement is true only when x is small. When x is sufficiently large, the admissions can acheive diversity entirely through affirmative action, and therefore has no need to randomize.

that

$$\gamma(t_1) = \Gamma(t_1, \overline{t}) - (\Lambda_N(\widehat{t}, \overline{t}) - \lambda(\widehat{t}))\frac{\int_{\widehat{t}}^{\overline{t}} n(t)dt}{\int_{t_1}^{\overline{t}} n(t) + w(t)dt}.$$

This condition differs from the necessary condition for the optimality of a one-step rule without affirmative action (Lemma 2) by the second term on the right-hand side (which vanishes when x tends to zero). By raising t_1, the admissions office raises p, thereby increasing the probability of admitting a minority candidate scoring between $\widehat{t}$ and $\overline{t}$ by $\frac{\int_{\widehat{t}}^{\overline{t}} n(t)dt}{\int_{t_1}^{\overline{t}} n(t)+w(t)dt}$. Because this candidate otherwise would be admitted through affirmative action, the admissions office uses the candidate s limited affirmative action slot to admit a minority candidate with test score $\widehat{t}$. But this lowers the quality of candidates admitted through limited affirmative action by $\Lambda_N(\widehat{t}, \overline{t}) - \lambda(\widehat{t})$, which makes raising t_1 less attractive to the admissions office than it would be under a complete ban. Using the implicit function theorem, we can verify that $t_1(x)$ is decreasing in x in the neighborhood of $x = 0$:

$$\frac{dt_1}{dx} = -\frac{c}{(1-p)}\frac{\lambda'(\widehat{t})}{\gamma'(t_1)}\frac{\int_{\widehat{t}}^{\overline{t}} n(t)dt}{\int_{t_1}^{\overline{t}} n(t) + w\,(t)\,dt} < 0.$$

The last inequality holds as $\gamma'(t_1) > 0$. The fact that $\frac{dt_1}{dx} < 0$ con rms the intuition that limited affirmative action leads to extra randomization: the more high-scoring minority candidates the admissions office can admit for sure, the lower the cost of randomization. Note that $\frac{dt_1}{dx}$ is proportional to $\frac{\int_{\widehat{t}}^{\overline{t}} n(t)dt}{\int_{t_1}^{\overline{t}} n(t)+w(t)dt}$, the probability that a randomly picked candidate between t_1 and t_2 is a minority candidate who would otherwise be admitted through affirmative action. As x approaches 0, $\widehat{t}$ tends to $\overline{t}$, and hence $\frac{dt_1}{dx}$ tends to 0. Intuitively, the admissions office wants to randomize more (i.e. decrease t_1) to avoid admitting high-scoring candidates in its main admissions rule who otherwise would be admitted through affirmative action. But with only a few limited-affirmative-action slots, the chance that a randomly picked candidate scoring between t_1 and $\overline{t}$ is a minority candidate scoring above $\widehat{t}$ is exceedingly small, so the admissions office has little extra incentive to do any randomization beyond what it does under a total ban. In that case, the direct effect of limited affirmative action it replaces the average candidate between t_1 and $\overline{t}$ with a minority candidate scoring near $\overline{t}$ outweighs the cost of additional randomization.

Proposition 6 *If $r_{1s}(t_1^*)$ is uniquely optimal in R_{NA}, then there exists $\overline{x} > 0$ such that for each $x \in (0, \overline{x})$, the regents prefer limited affirmative action of size x to a total ban.*

Now, consider the case where a two-step rule is optimal in the neighborhood of $x = 0$. In this case, the admissions office uses its xc affirmative-action slots to admit the highest-scoring minority candidates not admitted with probability one under the main admissions scheme, namely those scoring just below t_2. As a result, the rst-order conditions for the admissions office s problem are

$$\gamma(t_1) = \Gamma(t_1, t_2) - (\Lambda_N(\widehat{t}, t_2) - \lambda(\widehat{t}))\frac{\int_{\widehat{t}}^{t_2} n(t)dt}{\int_{t_1}^{t_2} n(t) + w(t)dt}$$

and

$$\gamma(t_1) = \gamma(t_2) - (\lambda(t_2) - \lambda(\widehat{t}))\psi(t_2).$$

The rst condition is the analog of the rst-order condition under a one-step rule, with t_2 replacing $\overline{t}$. The second captures the effect of limited affirmative action on t_2. By admitting a candidate scoring t_2, the admissions office admits a minority candidate who otherwise would be admitted through affirmative action with probability of $\psi(t_2)$. Thus, the net gain of admitting candidates scoring t_2 is $\gamma(t_2)$ minus $(\lambda(t_2) - \lambda(\widehat{t}))\psi(t_2)$, the expected effect on the quality of the xc candidates admitted through affirmative action. The second rst-order condition says that the value of admitting a candidate with test score t_2, taking the effect on the quality of the xc limited-affirmative-action candidates into account, should equal the value of admitting a candidate with test score t_1.

Using the implicit function theorem, we can obtain $\frac{dt_1}{dx}$ and $\frac{dt_2}{dx}$. Once again $\frac{dt_1}{dx}|_{x=0} = 0$, but now $\frac{dt_2}{dx}|_{x=0} = \frac{c\lambda'(t_2^*)\psi(t_2^*)}{(1-p^*)n(t_2^*)\gamma'(t_2^*)}$ is strictly positive. Recall that the additional incentive to randomize is proportional to the expected effect it has on the average quality of candidates admitted under affirmative action. Moving t_1 to the left (and reducing p correspondingly) is unlikely to change the candidates admitted under affirmative action when x is small. But moving t_2 to the right de nitely does affect the quality of candidates admitted under affirmative action, since all of the displaced minority candidates will be admitted through affirmative action, replacing less quali ed minority candidates (i.e., those with test score $\widehat{t}$). This explains why $\frac{dt_2^*}{dx}$ does not tend to zero as x tends to zero. Substituting the expressions for $\frac{dt_1}{dx}|_{x=0}$ and $\frac{dt_2}{dx}|_{x=0}$ into Lemma 3 gives:

$$\left.\frac{dU^R}{dx}\right|_{x=0} = (\lambda(t_2^*) - \Lambda(t_1^*, t_2^*))(1 - \frac{\lambda'(t_2^*)}{\gamma'(t_2^*)}) < 0,$$

The last inequality holds because $\gamma'(t_2^*) = \lambda'(t_2^*) + \frac{\alpha}{c}\psi'(t_2^*) < \lambda'(t_2^*)$; that is, because the fraction of minority candidates scoring t decreases with t. Unlike with a one-step rule, here a little bit of affirmative action makes the regents worse off. When the admissions office can admit xc candidates however it likes, it reduces the number of candidates admitted with probability one by more than xc, with the balance replaced by candidates with test scores between t_1 and t_2, thus reducing the overall quality of admitted candidates. This result is formally stated in the following proposition.

Proposition 7 *If $r_{2s}(t_1^*, t_2^*)$ is uniquely optimal in R_{NA} and $\gamma(t) - \psi(t_2^*)\lambda(t)$ is increasing in t for all $t \geq t_2^*$, then there exists $\overline{x} > 0$ such that for each $x \in (0, \overline{x})$, the regents prefer a total ban to limited affirmative action of size x.*

If $\gamma(t) - \psi(t_2^*)\lambda(t)$ is not increasing for all $t \geq t_2^*$, the admissions office may reduce the probability of admissions for candidates with some test score $t' > t_2^*$, rather then for those with test score t_2^*. In that case, the optimal admissions rule may not be a two-step rule, and, formally, Proposition 7 does not apply. (See Appendix B for details.) Nevertheless, because replacing candidates with $t' > t_2^*$ by those with test scores t_2^* makes the regents even worse off, we believe that the conclusion of Proposition 7 is likely to hold even when the condition fails.

Our formalization of limited affirmative action is not the only one possible. For example, the regents could allow the admissions office to set a different threshold for each group, but require that the number of minority candidates scoring between the thresholds be smaller than some upper bound. We have chosen our current approach mainly because it is easier to analyze. Finally, we would like to point out that our model probably underestimates the bene ts of limited affirmative action. In reality, the admissions office s preferences are likely to be concave in diversity, rather than linear. If so, then having access to limited affirmative action lowers the admissions office s taste for diversity at the margin and, thus, reduces its incentive to randomize.

8 Conclusion

American universities and colleges control their own admissions policies. Since most elite colleges and universities consider a diverse student body an important part of their missions, they will continue to promote diversity even when affirmative action is banned. In this paper, we explore the consequences of one such method, namely crafting an admissions rule favorable to minority (and low-scoring) candidates by partially ignoring standardized-test scores and other traditional measures of academic ability.

In our model of college admissions, the admissions office always admits the best-quali ed candidates from each group when it has access to affirmative action. However, when affirmative action is banned, it may choose a random admissions rule that partially ignores test scores. Random admissions rules are inefficient because they do not select the best candidates from any ethnic group: for every random admissions rule, there is a affirmative-action admissions rule that improves quality with the same level of diversity. Moreover, random admissions rules may be inefficient enough that banning affirmative action to improve student quality back res, both lowering quality and increasing the gap in average test scores between majority and minority matriculants.

Most of the public debate over affirmative action has focused on whether affirmative action is better than a ban. But most of the horror stories about affirmative action actually concern the size of the gap between marginal minority and majority candidates. We interpret our results on a partial ban as saying that even when the regents and admissions office cannot compromise on the size of affirmative action, there do exist simple, implementable forms of limited affirmative action which both may prefer to a ban.

Recent changes in admissions policies in California and Texas suggest that the phenomena we describe are real. We do not know, however, whether bans on affirmative action have substantially reduced student quality, nor do we know whether they have increased the gap in average test scores between groups and exacerbated stereotyping. But nothing in the structure of the admissions office s problem precludes the possibility that the agency problems we explore are severe enough to lower quality and raise the test-score gap.

While our model has focused on college admissions, its basic theme is likely to play out in other arenas. For example, many re and police departments have been under court order to increase diversity, but cannot use explicit affirmative action to achieve it. (See Lott [1998] for examples.) As a result, several have dropped tests of physical strength, speed, etc. While doing so may increase diversity, it may also reduce the average quality of new police officers from each ethnic group.

One nal lesson is that race-blind or race-neutral admissions policies are chimerical. A rule that treats two candidates who differ only in race identically may have been designed with racial concerns in mind. This last point may be particularly relevant to labor economics and other elds where researchers look for discrimination by testing whether race has predictive power independent of other variables (like test score) in the model. The right question often is not the relatively easy one of whether candidates with the same quali cations from different groups are treated identically, but the more difficult one of whether candidates various quali cations are properly weighted.

9 Appendix A

This appendix describes the general properties of optimal admissions rule when affirmative action is banned. Before proving Proposition 2, we prove three intermediate lemmas. Lemma 4 shows that optimal solutions exist. Lemma 5 shows that when γ has only a nite number of critical points the optimal admissions rule is a step function, and Lemma 6 provides a relation between the number of steps in the optimal rule and the number of γ s critical points.

Lemma 4 *There exists an optimal solution to the admissions office s problem.*

Proof Let $U^{AO}(r) = \int_{\underline{t}}^{\overline{t}} r(t)\gamma(t)(n(t)+w(t))dt$. It is clear that $U^{AO}(r)$ is bounded from above. Let $U^* \equiv \sup\{U^{AO}(r) : r \in R_{NA}\}$, where R_{NA} is the set of admissions rules that treat both groups identically. Let $\{r_i\}$ be a sequence of admissions rule such that $\lim_{i\to\infty} U^{AO}(r_i) = U^*$. Since r_i is increasing for all i, and $0 \,\square\, r_i(t) \,\square\, 1$ for all i and t, by Helly s selection theorem, there is a subsequence of r_i, r_i^*, that converges almost surely to some right-continuous, increasing function r^*. Since r^*, γ, n, and w are all bounded, by the bounded convergence theorem, $\int_{\underline{t}}^{\overline{t}} r^*(t)(n(t)+w(t))dt = \lim_{i\to\infty} \int_{\underline{t}}^{\overline{t}} r_i^*(t)(n(t)+w(t))dt = c$, and $U^{AO}(r^*(t)) = \lim_{i\to\infty} U^{AO}(r_i^*(t)) = U^*$. The rst equation implies that r^* is a well-de ned admissions rule, and the second implies that r^* is optimal. $\square$

Lemma 5 *Let $K_\gamma = \{t \in [\underline{t},\overline{t}] : \gamma'(t) = 0\}$, the set of critical points of the derivative of γ. If K_γ is nite, then no optimal $r \in R_{NA}$ strictly increases over any non-empty, open subset of $[\underline{t},\overline{t}]$.*

Proof If some optimal rule r strictly increases on some non-empty, open subset of $[\underline{t},\overline{t}]$, it strictly increases on at least one non-empty, open interval (a,b). Since K_γ is nite, $K_\gamma \cap (a,b)$ is a strict subset of (a,b). Because $(a,b) - (K_\gamma \cap (a,b))$ is non-empty and open, it contains at least one non-empty interval (c,d), on which γ is strictly monotone. Suppose γ increases on (c,d), and de ne

$$\hat{r}(t) = \begin{cases} r(t) & \text{for} \quad t \notin (c,d) \\ r(c) & \text{for} \quad t \in (c,\hat{t}) \\ r(d) & \text{for} \quad t \in [\hat{t},d), \end{cases}$$

with $\hat{t} \in (c,d)$ de ned such that $\int_c^d \hat{r}(t)(n(t)+w(t))dt = \int_c^d r(t)(n(t)+w(t))dt$. Then

$$U^{AO}(\hat{r}) - U^{AO}(r) = \int_c^d (\hat{r}(t) - r(t))(n(t)+w(t))\gamma(t)dt,$$

which, because γ increases on (c,d), is positive by Lemma 1, a contradiction. The case where γ decreases on (c,d) is similar and therefore omitted. $\square$

Lemma 6 *If a k-step rule, r_{ks}, is optimal in R_{NA}, then K_γ contains at least $2k-3$ points.*

Proof Let $t_1, t_2, ..., t_k$ be the thresholds of the k steps and $p_1, p_2, ..., p_k$ be the step probabilities; that is, p_j is the probability of acceptance for candidates in (t_j, t_{j+1}). The crux of the argument is that for each $j \in \{2, ..., k-1\}$, $\gamma(t_j) = \Gamma(t_j, t_{j+1}) = \gamma(t_{j+1})$. Suppose that for some $j \in \{2, ..., k-1\}$, $\gamma(t_{j+1}) > \Gamma(t_j, t_{j+1})$. Because γ is continuous, there exists some $\epsilon > 0$ such that $\Gamma(t_{j+1} - \epsilon, t_{j+1}) > \Gamma(t_j, t_{j+1})$. By reducing t_{j+1} to $t_{j+1} - \epsilon$ and lowering p_j to satisfy the capacity constraint, the admissions office changes its utility by the term

$$(p_{j+1} - p_j) \left(\int_{t_{j+1}-\epsilon}^{t_{j+1}} (n(t) + w(t)) dt \right) (\Gamma(t_{j+1} - \epsilon, t_{j+1}) - \Gamma(t_j, t_{j+1} - \varepsilon)) > 0,$$

which contradicts r_{ks} being optimal. The case where $\gamma(t_{j+1}) < \Gamma(t_j, t_{j+1})$ and the argument for $\gamma(t_j) = \Gamma(t_j, t_{j+1})$ are similar and therefore omitted. Finally, note that since for each $j \in \{2, ..., k-1\}$, $\gamma(t_j) = \Gamma(t_j, t_{j+1}) = \gamma(t_{j+1})$, γ must have at least two critical points in each step (t_j, t_{j+1}). If it had none, then $\gamma(t_j) \neq \gamma(t_{j+1})$. If it had only one, then for each $t \in (t_j, t_{j+1})$, $\gamma(t) > \gamma(t_j) = \gamma(t_{j+1})$ or $\gamma(t) < \gamma(t_j) = \gamma(t_{j+1})$, and either way $\Gamma(t_j, t_{j+1}) \neq \gamma(t_j)$, a contradiction.

For $j = 1$, if $t_1 > \underline{t}$, $\gamma(t_1) = \Gamma(t_1, t_2)$ by the argument above. If $t_1 = \underline{t}$, then $\gamma(t_1) \geq \Gamma(t_1, t_2)$, otherwise increasing t_1 raises utility. In this case, γ must have a least one critical point in (t_1, t_2). To see why, suppose that γ is monotone on (t_1, t_2). If it is increasing, raising t_1 to $t_1 + \epsilon$ and lowering p_1 to satisfy the capacity constraint increases utility. If it is decreasing, then $\Gamma(t_1, t_2) > \gamma(t_2)$, and raising t_2 to $t_2 + \epsilon$ and raising p_1 to satisfy the capacity constraint increases utility. Both are contradictions, so γ has a least one critical point in (t_1, t_2). Since each of the $k-2$ interior steps has at least two critical points, and the left-most step has at least one critical point, γ must have at least $2(k-2) + 1 = 2k - 3$ critical points. □

Proof of Proposition 2 First consider a k-step rule that is optimal among k-step rules. Let $t_1, t_2, ..., t_k$ be the thresholds of the k steps and $p_1, p_2, ..., p_k$ be the step probabilities. From the proof of Lemma 6, we know that for each $j \in \{1, ..., k-1\}$, $\gamma(t_j) = \Gamma(t_j, t_{j+1}) = \gamma(t_{j+1})$. This implies that for all $j \in \{1, ..., k-2\}$, $\Gamma(t_j, t_{j+1}) = \Gamma(t_{j+1}, t_{j+2})$, meaning that the admissions office is indifferent between candidates drawn from each of the rst $k-1$ steps. Thus, starting from any optimal k-step rule, by setting the probabilities of admission in the rst $k-1$ steps equal, we can merge the rst $k-1$ steps into one step without affecting the utility of the admissions office . As a result, for every optimal k-step rule, r_{ks}, there is a two-step rule, r_{2s}, such that $U^{AO}(r_{2s}) = U^{AO}(r_{ks})$.

More generally, for any feasible admissions rule r, which by de nition is integrable, there exists a sequence of k-step rules $\{r_{ks}\}_{k=1}^{\infty}$ such that $U^{AO}(r) = \lim_{k\to\infty} U^{AO}(r_{ks})$. From above, we know that there exists a sequence of two-step rules, $\{r_{2s}^k\}_{k=1}^{\infty}$, such that $U^{AO}(r_{ks})$ □ $U^{AO}(r_{2s}^k)$. Moreover, because of Lemma 4, we can assume without loss of generality that $\{r_{2s}^k\}$ converges to some well-de ned two-step rule, r_{2s}^*. It follows that $U^{AO}(r) = \lim_{k\to\infty} U^{AO}(r_{ks}) \square \lim_{k\to\infty} U^{AO}(r_{2s}^k) = U^{AO}(r_{2s}^*)$, meaning that for any feasible admissions rule, r, there is a two-step rule, r_{2s}^*, that is at least as good. Hence, an optimal two-step rule must exist.

We now show that the optimal rule is unique when γ has only two critical points. We know from Lemma 6 that in this case an optimal rule has at most two steps. First note that there could not be two optimal proper one-step rules. For if there were, then the function $\Gamma(t, \bar{t})$ would have at least two interior maxima, which is impossible when there are only two critical points. Suppose there are two optimal rules, denoted by $r = r_{2s}(t_1, t_2)$ and $\hat{r} = r_{2s}(\widehat{t_1}, \widehat{t_2})$, respectively. If $t_2 = \widehat{t_2}$, then it must be that $t_1 = \widehat{t_1}$; otherwise, following the logic above, there would be more than two critical points in $[\underline{t}, t_2]$. Thus, if $r \neq \hat{r}$, then $t_2 \neq \widehat{t_2}$. Without loss of generality, we can assume that $t_2 > \widehat{t_2}$. If $\Gamma(t_1, \widehat{t_2}) < \Gamma(\widehat{t_2}, t_2)$, then r cannot be optimal, as the admissions office would be better off replacing candidates scoring between $t_1$and $\widehat{t_2}$ by those scoring between $\widehat{t_2}$ and t_2. If $\Gamma(t_1, \widehat{t_2}) > \Gamma(\widehat{t_2}, t_2)$, then $\hat{r}$ cannot be optimal, as the admissions office would be better off replacing candidates scoring between $\widehat{t_2}$ and t_2 by those scoring between t_1 and $\widehat{t_2}$. Finally, if $\Gamma(t_1, \widehat{t_2}) = \Gamma(\widehat{t_2}, t_2)$, then the admissions office is indifferent between candidates scoring between $\widehat{t_2}$ and t_2 and those scoring between t_1 and $\widehat{t_2}$, meaning that some three-step rule would also be optimal, a contradiction. Thus, the optimal rule must be unique. □

10 Appendix B

This appendix contains proofs of the remaining lemmas and propositions.

Proof of Lemma 1 Integrating by parts yields

$$\int_{t_1}^{\bar{t}} h(t)dF(t)dt = [h(t)F(t)]_{t_1}^{\bar{t}} - \int_{t_1}^{\bar{t}} F(t)dh(t) = h(\bar{t}) - F(t_1)h(t_1) - \int_{t_1}^{\bar{t}} F(t)dh(t).$$

Likewise,

$$\int_{t_1}^{\bar{t}} h(t)dG(t)dt = h(\bar{t}) - G(t_1)h(t_1) - \int_{t_1}^{\bar{t}} G(t)dh(t).$$

Thus,

$$\int_{t_1}^{\bar{t}} h(t)dF(t)dt - \int_{t_1}^{\bar{t}} h(t)dG(t)dt = (G(t_1) - F(t_1))\, h(t_1) + \int_{t_1}^{\bar{t}} (G(t) - F(t))\, dh(t) \geq 0,$$

since h is non-decreasing and non-negative. □

Proof of Proposition 1 Suppose the admissions office wants to admit $(1-\delta)c$ candidates from group W. Let $U_W^{AO}(r) = \int_{\underline{t}}^{\overline{t}} \lambda(t)r(t)w(t)dt$ denote the admissions office s utility from the $(1-\delta)c$ majority candidates admitted under rule r. Let r_{th} be the threshold rule satisfying the capacity constraint and r be some other admissions rule. Because $\int_{\underline{t}}^{t} r_{th}(t)w(t)dt$ □ $\int_{\underline{t}}^{t} r(t)w(t)dt$ for all $t \in [\underline{t}, \overline{t}]$, Assumption 1 and Lemma 1 imply that

$$U_W^{AO}(r_{th}) = \int_{\underline{t}}^{\overline{t}} \lambda(t)r_{th}(t)w(t)dt \geq \int_{\underline{t}}^{\overline{t}} \lambda(t)r(t)w(t)dt = U_W^{AO}(r),$$

and, thus, r_{th} is optimal.

If we ignore the absolute-value sign in its objective function, the admissions office solves

$$\max_{t_N, t_W} \quad \int_{t_N}^{\overline{t}} (\lambda(t) + \frac{\alpha}{c})n(t)dt + \int_{t_W}^{\overline{t}} \lambda(t)w(t)dt$$

subject to the capacity constraint. If the optimal solution to the modi ed problem has the minority group underrepresented, then it is also the solution to the original problem. First-order conditions of the modi ed problem imply that $\lambda(t_N) + \frac{\alpha}{c} = \lambda(t_W)$. Recall that $r_{th}(\widetilde{t}_N, \widetilde{t}_W)$ is the threshold rule that achieves proportionate representation. Since λ is strictly increasing in t, it follows from Assumption 3 that $\widetilde{t}_W > t_W$ and $\widetilde{t}_N < t_N$. Thus, the minority group is underrepresented under the optimal rule, justifying the assumption made above. Finally, since $\alpha > 0$, $t_W > t_N$. □

Proof of Proposition 3 In text.

Proof of Proposition 4 Under affirmative action, all minority candidates scoring above t_N and all majority candidates scoring above t_W are admitted. Thus, for the number of minority candidates to increase (and the number of majority candidates to decrease) under a ban, it must be that r', the admissions rule adopted when affirmative action is banned, admits some candidates scoring below t_N with positive probability and some candidates scoring above t_W with probability less than one. From Proposition 1, $\lambda(t_N) + \frac{\alpha}{c} = \lambda(t_W)$. It follows that for all t_1 □ t_N and for all $t_2 \geq t_W$,

$$\gamma(t_1) = \lambda(t_1) + \frac{\alpha}{c}\psi(t_1) < \lambda(t_N) + \frac{\alpha}{c} = \lambda(t_W) < \lambda(t_2) + \frac{\alpha}{c}\psi(t_2) = \gamma(t_2).$$

That is, the admissions office strictly prefers candidates scoring above t_W to those scoring below t_N. Hence, the admissions office can improve on r' by replacing the candidate with the lowest test score admitted with positive probability, which by assumption is below t_N, by the candidate with the highest test score admitted with probability less than one, which by assumption is above t_W.

Thus, r' cannot be optimal. $\square$

Proof of Proposition 5 From Proposition 3, banning affirmative action strictly increases majority representation. Since the $W(r_{th})$ majority group candidates admitted under $r_{th}(t_N, t_W)$ are the $W(r_{th})$ highest-scoring candidates, the $W(r) > W(r_{th})$ candidates admitted without affirmative action must have a lower average test score.

Part 1 and the rst part of Part 2 follow from $t_c > t_N$ or $t_1 > t_N$ (see the proof of Proposition 4 for the latter), which implies that the average test score for admitted N candidates rises, making the gap fall. $\square$

Proof of Lemma 3 The regents utility is

$$U^R = p(x)\int_{t_1(x)}^{t_2(x)} \lambda(t)(n(t)+w(t))dt + \int_{t_2(x)}^{\bar{t}} \lambda(t)(n(t)+w(t))dt + (1-p(x))\int_{\hat{t}(x)}^{t_2(x)} \lambda(t)n(t)dt.$$

Note that $\widehat{t}(0) = t_2^*$ and hence $\frac{dU^R}{dx}|_{x=0}$ can be written as

$$\begin{aligned}\left.\frac{dU^R}{dx}\right|_{x=0} &= \left.\frac{dp}{dx}\right|_{x=0}\int_{t_1^*}^{t_2^*} \lambda(t)(n(t)+w(t))dt - p^*\lambda(t_1^*)(n(t_1^*)+w(t_1^*))\left.\frac{dt_1}{dx}\right|_{x=0} - \\ &\quad (1-p^*)\lambda(t_2^*)(n(t_2^*)+w(t_2^*))\left.\frac{dt_2}{dx}\right|_{x=0} + (1-p^*)\lambda(t_2^*)n(t_2^*)\left(\left.\frac{dt_2}{dx}\right|_{x=0} - \left.\frac{d\widehat{t}}{dx}\right|_{x=0}\right).\end{aligned}$$

Differentiating the two capacity constraints gives

$$\left.\frac{dp}{dx}\right|_{x=0}\int_{t_1^*}^{t_2^*} n(t)+w(t)dt - p^*(n(t_1^*)+w(t_1^*))\left.\frac{dt_1}{dx}\right|_{x=0} - (1-p^*)(n(t_2^*)+w(t_2^*))\left.\frac{dt_2}{dx}\right|_{x=0} = -c$$

$$(1-p^*)n(t_2^*)\left(\left.\frac{dt_2}{dx}\right|_{x=0} - \left.\frac{d\widehat{t}}{dx}\right|_{x=0}\right) = c.$$

Substituting these into the expression for $\frac{dU^R}{dx}$ above gives the desired result. $\square$

Lemma 7 *Let r_x^* denote an optimal admissions rule when the size of affirmative action is xc. (i) If r_0^* is a one-step rule, then there exists $\bar{x} > 0$ such that for all $x \in (0, \bar{x})$, r_x^* is a one-step rule. (ii) If r_0^* is a proper two-step rule and $\gamma(t) - \psi(t_2^*)\lambda(t)$ is strictly increasing for $t \geq t_2^*$, then there exists $\bar{x} > 0$ such that for all $x \in (0, \bar{x})$, r_x^* is a proper two-step rule.*

Proof Assume throughout that x is sufficiently small that the admissions lls all xc seats with minority candidates. First consider the case where r_0^* is a one-step rule denoted by $r_{1s}^*(t_1^*)$. Recall that Proposition 4 implies that a completely random lottery is never optimal, so $t_1^* > \underline{t}$ and $\Gamma(\underline{t}, \bar{t}) < \Gamma(t_1^*, \bar{t})$. The second-order condition implies that $\gamma'(t_1^*) > 0$. It follows that there exists some $t' < t_1^*$ such that $\gamma'(t) > 0$, for all $t \in [t', t_1^*]$, and that $\gamma(t') > \Gamma(t, t')$ for all $t \in [\underline{t}, t')$. Note that the effect of the affirmative-action slots on the main admission rule is proportional to x, the number of the slots. If x is sufficiently small so that the admissions office never admits a candidate scoring below t' over one randomly drawn from $[t_1^*, \bar{t}]$, we claim that r_x^* must be a two-step rule. First note that r_x^* must be at after t_1^*. If not, the admissions office can raise both the value of candidates admitted through the main admissions rule and the value of those admitted through affirmative action by equalizing the probability of admission for candidates with $t \geq t_1^*$. By construction, no candidate scoring below t' is admitted. Since γ is strictly increasing between t' and t_1^*, when candidates scoring in that range are admitted, it must be in descending order of test score, subject to the constraint that $r_x^*(t) \square r_x^*(t_1^*)$. Thus, r_x^* is an one-step rule.

Suppose r_0^* is a two-step rule denoted by $r_{2s}^*(t_1^*, t_2^*)$. Since $\Gamma(t_1^*, t_2^*)$ is the maximum value of $\Gamma(t, t_2^*)$ for $t \in [\underline{t}, t_2^*]$, following the same argument as above we can show that when x is sufficiently small, r_x^* must be a one-step rule when $t \in [\underline{t}, t_2^*]$. To complete the proof, we need to show that $r_x^*(t)$ is either equal to $r_x^*(t_1^*)$ or 1 for all $t \in [t_2^*, \bar{t}]$.

Let $t_y = \sup\{t : r_x^*(t) < 1\}$ and $t_x = \inf\{t : r_x^*(t) > r_x^*(t_1^*)\}$. If r_x^* is not a two-step rule, then $t_y > t_x$. In that case, we show that the admissions office can improve on r_x^* by replacing a candidate scoring t_x by one scoring t_y, a contradiction. (By construction this change does not violate the monotonicity constraint.) The direct effect of replacing a candidate scoring t_x with one scoring t_y is to raise the quality by $\gamma(t_y) - \gamma(t_x)$. But the net effect on the admissions office s welfare also depends on the indirect effect on the quality of those admitted through affirmative action. There is a chance $\psi(t_y)$ that the newly-admitted candidate scoring t_y is a minority candidate who would have been admitted through affirmative action. If so, then another minority candidate scoring $\widehat{t}^*$, the highest test score among minority candidates not admitted through affirmative action under r_x^*, will now be admitted. If $\widehat{t}^* > t_x$, then the total quality of the affirmative action candidates decreases by $\psi(t_y)(\lambda(t_y) - \lambda(\widehat{t}))$. If $\widehat{t}^* \square t_x$, then the total quality of the affirmative action candidates decreases by $\psi(t_y)(\lambda(t_y) - \lambda(t_x)) - (\psi(t_x) - \psi(t_y))(\lambda(t_x) - \lambda(\widehat{t}^*))$. In either case, the decline in quality is less than $\psi(t_y)(\lambda(t_y) - \lambda(t_x))$. Thus, the net gain in welfare is greater than

$$
\begin{aligned}
& (\gamma(t_y) - \gamma(t_x)) - \psi(t_y)(\lambda(t_y) - \lambda(t_x)) \\
= \quad & ((\gamma(t_y) - \psi(t_2^*)\lambda(t_y)) - (\gamma(t_x) - \psi(t_2^*)\lambda(t_x)) + (\psi(t_2^*) - \psi(t_y))(\lambda(t_y) - \lambda(t_x)).
\end{aligned}
$$

Note that $(\gamma(t_y) - \psi(t_2^*)\lambda(t_y)) - (\gamma(t_x) - \psi(t_2^*)\lambda(t_x)$ is positive, as$(\gamma(t) - \psi(t_2^*)\lambda(t)$ is increasing in t when $t > t_2^*$. Hence, the net gain in welfare must be strictly positive. Thus $t_x = t_y$, and r_x^*must be a two-step rule. □

Proof of Proposition 7 By Lemma 7, without loss of generality we can assume that the admissions office uses a one-step rule. The admissions office s problem is to pick p, t_1,and $\widehat{t}$ to maximize $p\int_{t_1}^{\overline{t}} \gamma(t)(n(t)+w(t))dt + (1-p)\int_{\widehat{t}}^{\overline{t}}(\lambda(t)+\frac{\alpha}{c})n(t)dt$ subject to the capacity constraints: $p\int_{t_1}^{\overline{t}} n(t)+w(t)dt = (1-x)c$ and $(1-p)\int_{\widehat{t}}^{\overline{t}} n(t)dt = xc$. The rst-order conditions of the problem are

$$\gamma(t_1) = \Gamma(t_1,\overline{t}) + \left(\lambda(\widehat{t}) - \Lambda(\widehat{t},\overline{t})\right)\frac{\int_{\widehat{t}}^{\overline{t}} n(t)dt}{\int_{t_1}^{\overline{t}} n(t)+w(t)dt}$$

and the two capacity constraints. It follows from the implicit function theorem that

$$\begin{bmatrix} \gamma'(t_1) & -\lambda'(\widehat{t})\frac{\int_{\widehat{t}}^{\overline{t}} n(t)dt}{\int_{t_1}^{\overline{t}} n(t)+w(t)dt} & 0 \\ 0 & -(1-p)n(\widehat{t}) & -\int_{\widehat{t}}^{\overline{t}} n(t)dt \\ -p(n(t_1)+w(t_1)) & 0 & \int_{t_1}^{\overline{t}} n(t)+w(t)dt \end{bmatrix}\begin{bmatrix} \frac{dt_1}{dx} \\ \frac{d\widehat{t}}{dx} \\ \frac{dp}{dx} \end{bmatrix} = \begin{bmatrix} 0 \\ c \\ -c \end{bmatrix}.$$

When x is small, $\widehat{t}$ is close to $\overline{t}$. Thus, we can ignore $(\int_{\widehat{t}}^{\overline{t}} n(t)dt)^2$ in the neighborhood of $x=0$, and write $\frac{dt_1}{dx}$ as

$$\frac{dt_1}{dx} = -\frac{c}{1-p}\frac{\lambda'(\widehat{t})}{\gamma'(t_1)}\frac{\int_{\widehat{t}}^{\overline{t}} n(t)dt}{\int_{t_1}^{\overline{t}} n(t)+w\left(t\right)dt} < 0.$$

Since $\widehat{t}$ tends to $\overline{t}$ as x tends to 0, $\frac{dt_1}{dx}|_{x=0} = 0$. From Lemma 3, $\frac{dU^R}{dx}|_{x=0} = c\left(\lambda(t_2^*) - \Lambda(t_1^*,t_2^*)\right) > 0$.
□

Proof of Proposition 8 Again, we can assume without loss of generality that the admissions office chooses a two-step rule, in which case it maximizes

$$p\int_{t_1}^{t_2}\gamma(t)(n(t)+w(t))dt + \int_{t_2}^{\overline{t}}\gamma(t)(n(t)+w(t))dt + (1-p)\int_{\widetilde{t}}^{t_2}(\lambda(t)+\frac{\alpha}{c})n(t)dt$$

subject to the constraints

$$p\int_{t_1}^{t_2} n(t)+w(t)dt + \int_{t_2}^{\overline{t}} n(t)+w(t)dt = (1-x)c$$
$$(1-p)\int_{\widehat{t}}^{t_2} n(t)dt = xc.$$

The rst-order conditions are

$$\begin{aligned}\gamma(t_1) &= \Gamma(t_1,t_2) + \left(\lambda(\widehat{t}) - \Lambda(\widehat{t},t_2)\right)\frac{\int_{\widehat{t}}^{t_2} n(t)dt}{\int_{t_1}^{t_2} n(t)+w(t)dt}, \\ \gamma(t_1) &= \gamma(t_2) + (\lambda(\widehat{t}) - \lambda(t_2))\psi(t_2),\end{aligned}$$

and the two capacity constraints. From the implicit function theorem,

$$\begin{bmatrix} \gamma'(t_1) & -\gamma'(t_2)+\lambda'(t_2)\psi(t_2) & -\lambda'(\widehat{t})\psi(\widehat{t}) & 0 \\ \gamma'(t_1) & 0 & \frac{-\lambda'(\widehat{t})\int_{\widehat{t}}^{t_2} n(t)dt}{\int_{t_1}^{t_2} n(t)+w(t)dt} & 0 \\ 0 & (1-p)n(t_2) & -(1-p^*)n(\widehat{t}) & -\int_{\widehat{t}}^{t_2} n(t)dt \\ -p(n(t_1)+w(t_1)) & -(1-p)\left(n(t_2)+w\left(t_2\right)\right) & 0 & \int_{t_1}^{t_2} n(t)+w(t)dt \end{bmatrix} \begin{bmatrix} \frac{dt_1}{dx} \\ \frac{dt_2}{dx} \\ \frac{dt}{dx} \\ \frac{dp}{dx} \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ c \\ -c \end{bmatrix},$$

which implies that $\frac{dt_1}{dx}|_{x=0} = 0$, and $\frac{dt_2}{dx}|_{x=0} = \frac{c\lambda'(t_2^*)\psi(t_2^*)}{(1-p^*)n(t_2^*)\gamma'(t_2^*)}$. Substituting these expression into Lemma 3 yields

$$\left.\frac{dU^R}{dx}\right|_{x=0} = c(\lambda(t_2^*) - \Lambda(t_1^*, t_2^*))(1 - \frac{\lambda'(t_2^*)}{\gamma'(t_2^*)}) < 0.\square$$

References

[1] Alger, Jonathan. Affirmative Action in Higher Education: A Current Legal Overview, <http://www.aaup.org/aalglovr.htm>, 1998.

[2] Bok, Derek. *Beyond the Ivory Tower: Social Responsibilities of the Modern University*. Cambridge, MA: Harvard University Press, 1982.

[3] Bowen, William and Bok, Derek. *The Shape of the River: Long-Term Consequences of Considering Race in College and University Admissions*. Princeton, NJ: Princeton University Press, 1998.

[4] Committee on Admissions and Enrollment, Berkeley Division, Academic Senate, University of California. *Freshman Admissions at Berkeley: A Policy for the 1990s and Beyond*. Berkeley, CA: The Senate, 1989.

[5] Committee on Admissions and Enrollment, Berkeley Division, Academic Senate, University of California. *The Implementation of the Karabel Report on Freshman Admissions at Berkeley: 1990-1993*. Berkeley, CA: The Senate, 1993.

[6] Conrad, Cecilia A. and Sharpe, Rhonda V. The Impact of the California Civil Rights Initiative (CCRI) on University and Professional School Admissions and the Implications for the California Economy. *The Review of Black Political Economy*, Summer 1996.

[7] Geiser, Saul. Rede ning UC s Eligibility Pool to Include A Percentage of Students From Each High School: Summary of Simulation Results, UC Office of the President, March 1999.

[8] Holley, Danielle and Spencer, Delia. The Texas Ten Percent Plan, *Harvard Civil Rights-Civil Liberties Law Review* 34, 1999.

[9] Jencks, Christopher. Racial Bias in Testing, in *The Black-White Test Score Gap*, C. Jencks and M. Phillips, eds. Washington, DC: Brookings Press, 1998.

[10] Jencks, Christopher and Phillips, Meredith, eds. *The Black-White Test Score Gap*. Washington, DC: Brookings Press, 1998.

[11] Kane, Thomas. Racial and Ethnic Preferences in College Admissions, in *The Black-White Test Score Gap*, C. Jencks and M. Phillips, eds. Washington, DC: Brookings Press, 1998.

[12] Lott, John R. Jr. Does a Helping Hand Put Others at Risk?: Affirmative Action, Police Departments, and Crime. Mimeo, University of Chicago Law School, 1998.

[13] Loury, Glenn C. How to Mend Affirmative Action. *Public Interest*, Spring 1987.

[14] Scanlan, Laura C. Hopwood *v.* Texas: A Backward Look at Affirmative Action in Education. *New York University Law Review* 71, December 1996.

[15] Steele, Claude M. and Aronson, Joshua. Stereotype Threat and the Test Performance of Academically Successful African Americans, in *The Black-White Test Score Gap*, C. Jencks and M. Phillips, eds. Washington, DC: Brookings Press, 1998.

[16] Vars, Frederick E. and Bowen, William G. Scholastic Aptitude Test Scores, Race and Academic Performance in Selective Colleges and Universities, in *The Black-White Test Score Gap*, C. Jencks and M. Phillips, eds. Washington, DC: Brookings Press, 1998.

[17] Wilgoren, Jodi. Texas Top 10% Law Appears to Preserve College Racial Mix, in *The New York Times*, November 24, 1999.