UC Berkeley Admissions in 2015 and 2016:
An Analysis of the Role of Letters of Recommendation and Augmented Review

July 13, 2016
Jesse Rothstein,
Professor of Public Policy & Economics

Executive Summary

UC Berkeley made many changes to its undergraduate admission processes in 2016, including the addition of letters of recommendation (LORs) for many applicants and the elimination of the separate "Augmented Review" (AR) pool. These changes were intended in part to better identify hidden gems in the applicant pool – because so many more students could be asked for LORs than were ever considered under AR, if LORs were anywhere near as effective at identifying underrepresented students who could succeed at Cal the net impact would be to increase diversity. On the other hand, the changes might have raised barriers to admission for disadvantaged students, who might not have access to teachers or counselors willing and able to write strong letters and whose hidden strengths might not have been recognized without AR.

Vice Chancellor Koshland asked me in June to conduct an independent analysis of Berkeley undergraduate admissions, focusing on the LOR and AR changes. I have done academic research on college admissions at UC and elsewhere, but I have never been involved with Berkeley admissions processes. To prepare this report, I consulted extensively with the admissions office and with the Senate Committee on Admissions, Enrollment, and Preparatory Education (AEPE) in an effort to understand the admissions process. This report reflects my independent analysis and conclusions.

I reach two main conclusions. First, the number of applicants from underrepresented groups (low-income, first-generation, from low-API high schools, and/or underrepresented minorities) who were admitted rose in 2016. But because the number admitted who were not from these groups rose by a larger proportion, the *share* of admitted students from underrepresented groups fell somewhat.

Second, neither the addition of LORs nor the removal of AR contributed meaningfully to this decline. If anything, asking for LORs *raised* the relative admissions rates of applicants from underrepresented groups. I am unable to precisely identify the impact of the elimination of AR – while some estimates indicate that this slightly reduced admissions for those who would have been considered via AR in 2015, others indicate zero or even a positive effect. All of the estimates agree that the impact was small in any case.

It is beyond the scope of this report to identify what *did* cause the shift in 2016. My preliminary investigation, however, suggests that the decline in the share of admits from underrepresented groups is in large part a statistical artifact due to the expanded use of the waitlist in 2016. There were also reductions in the admissions chances of the underrepresented applicants with the strongest numeric records that cannot be attributed to the waitlist. Future investigation should focus on understanding what in the scoring process harmed these applicants.

Introduction

Berkeley made a number of changes to its undergraduate admissions processes in 2016:

- It requested letters of recommendation (LOR) for many applicants.
- It eliminated the augmented review (AR) pool.
- It shifted from a point system for scoring applications to a three-category (Yes/No/Maybe) rating system.
- Readers began scoring applicants on a list of holistic / non-cognitive factors.
- Every application was read twice, where in the past many were read only once
- A third read, by members of the faculty, was added for many applications.
- The wait list was used much more extensively than in the past, and many applicants who in 2015 would have been admitted or rejected outright were instead offered positions on the waitlist in 2016.

Table 1 presents simple summaries of admissions outcomes in 2015 and 2016, in the upper panel pooling all in-state and out-of-state applicants and in the lower panel restricting attention to California residents not being recruited as athletes. Of particular concern is the decline in the share of admitted students from disadvantaged backgrounds or groups that are underrepresented at Berkeley. In Table 1 and throughout this report, I consider four groups of such students: low-income students (with family incomes below $40,000); first-generation college students (those whose parents did not graduate from college); students from disadvantaged schools (with API indexes of 5 and below); and under-represented minority students (UREMs). I refer to them collectively as “underrepresented.”

I was asked to conduct an analysis of admissions data for the 2015 and 2016 cycles, focusing on the impact of the addition of letters of recommendation and the elimination of augmented review on the admission of underrepresented students. In all of the analyses below, I restrict attention to California resident applicants who were not classified in the admissions process as recruited athletes. The lower portion of Table 1 shows statistics for these applicants.

The overall context

I begin with several preliminaries necessary to understanding the impacts of letters of recommendation and augmented review on admissions outcomes.

The composition of the applicant pool

The first place to look for an explanation for changing outcomes is changes in the applicant pool itself. If the applicants from underrepresented groups were weaker, on average, in 2016 than in 2015, this could account for the overall observed changes, even without a change in policy, and would confound my LOR and AR analyses. Understanding the distribution of applicant strength is also helpful as a way of gauging which types of students were affected by the LOR and AR programs.

Table 1. Summary of admissions outcomes				
	Applicants		**Admits**	
	2015	**2016**	**2015**	**2016**
All applicants				
Number	78,924	82,578	13,266	14,423
Admission rate			16.8%	17.5%
Share from underrepresented groups				
Low income (< $40K)	24%	23%	19%	18%
First generation	17%	16%	13%	12%
Low API	13%	12%	12%	11%
UREM	23%	23%	18%	18%
Any of these four	40%	38%	33%	31%
California residents, excluding recruited athletes				
Number	45,570	45,626	8,570	9,610
Admission rate			18.8%	21.1%
Share from underrepresented groups				
Low income (< $40K)	31%	30%	24%	22%
First generation	24%	23%	18%	15%
Low API	23%	22%	19%	16%
UREM	33%	34%	23%	23%
Any of these four	53%	53%	42%	39%

I construct a measure of applicant strength by aggregating a number of available student characteristics.[1] My "admissions score" represents the predicted probability that a student with a given set of characteristics would have been admitted in the first round (i.e., not off the wait list, and not through Augmented Review), had he/she applied as a California resident to the College of Letters and Sciences in 2015.[2]

[1] The variables are those used in the model used by the admissions office to predict application read scores, and include measures of traditional academic strength, measures characterizing the high school, (including, for example, the extent to which the student took advantage of the school's advanced course offerings), and three of the four disadvantage measures considered here. Applicants' race and ethnicity, which cannot be considered in admissions, is not included.
[2] I use only L&S applicants because the other colleges may weight characteristics differently than does Letters and Sciences. Nevertheless, my L&S-based score is nearly perfectly correlated with a score constructed based on admissions in the College of Engineering.

Two aspects of this score must be emphasized: First, it captures only the quantitative characteristics that are coded in the admissions office's database; readers see more information, and may identify applicants as stronger or weaker than is indicated by my score. Second, the characteristics are weighted to best predict admission in 2015. The weight put on different characteristics – say, on high SAT scores vs. taking all of the AP courses offered at your high school – might vary from year to year, and indeed seems to have changed somewhat in 2016 (as discussed below).

But even with these caveats, the admissions score nevertheless presents a useful summary. To take one example, 38% of applicants in 2016 have scores under 1%. While a very few of these students might have characteristics not in the database that merit admission, this is quite rare; the vast majority of students in this group would not be admitted under the regular 2015 processes. Indeed, only 1.8% of them were admitted in 2015, and 2.5% in 2016.

Figure 1 shows the distribution of admissions scores across California resident, non-athlete applicants to all colleges in 2015. This is heavily left-skewed: Most applicants have very low chances of admission, though there are a few who are so strong on the dimensions captured by my index that it is rare for other factors to prevent them from being admitted.

Figure 1. Distribution of admissions scores in 2015 for California resident applicants

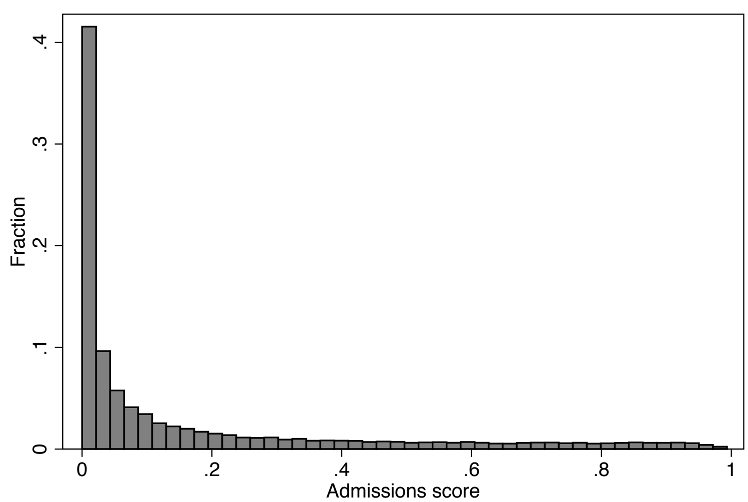

Because the overall distribution of admissions scores is so dominated by applicants with extremely low chances of admission, I find it helpful to focus on those who are more likely to be admitted. Figure 2 shows the distribution of admissions scores for those who were actually admitted in 2015 (including AR admits and those admitted off the waitlist), while Figure 3 repeats this for the four underrepresented groups and Figure 4 repeats it for applicants not from these groups.

Figure 2. Distribution of admissions scores for admitted California resident applicants in 2015

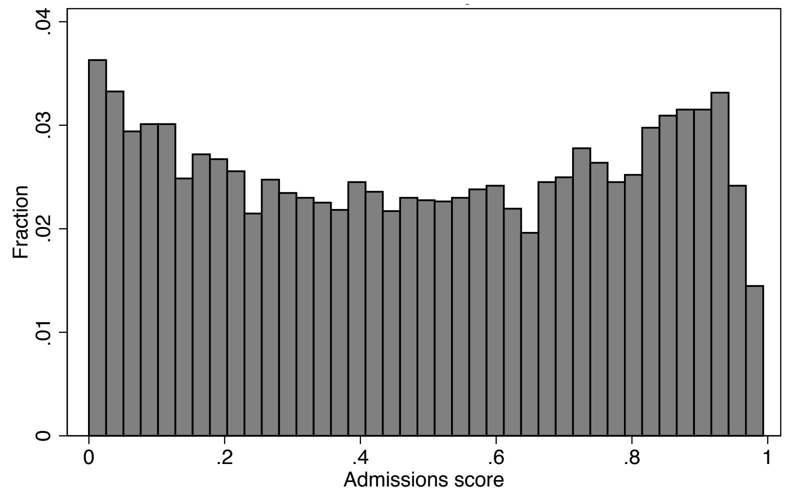

It is apparent in Figure 3 that the admissions score distribution is quite skewed to the left for students from the underrepresented groups. This is true even though the prediction model used to generate the admissions score *includes* indicators for low income, first generation, and low API (but not UREM) students. Evidently, many of the students who are admitted from these groups are picked out from large pools with similar observable credentials who are not admitted. This is much less true for students not from these groups, for whom the distribution is shown in Figure 4: Here, admitted students are much more likely to have admissions scores above 0.6.

Figure 3.Admissions score distributions for admitted California resident applicants from four underrepresented groups in 2015

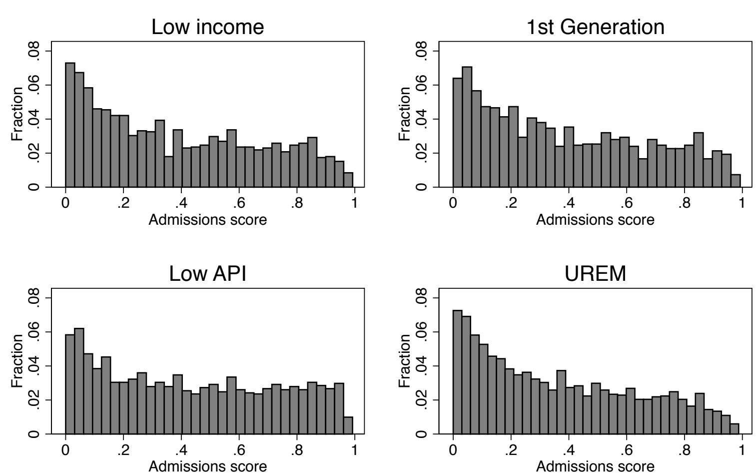

Figure 4. Admission score distribution for California residents not from underrepresented groups, 2015

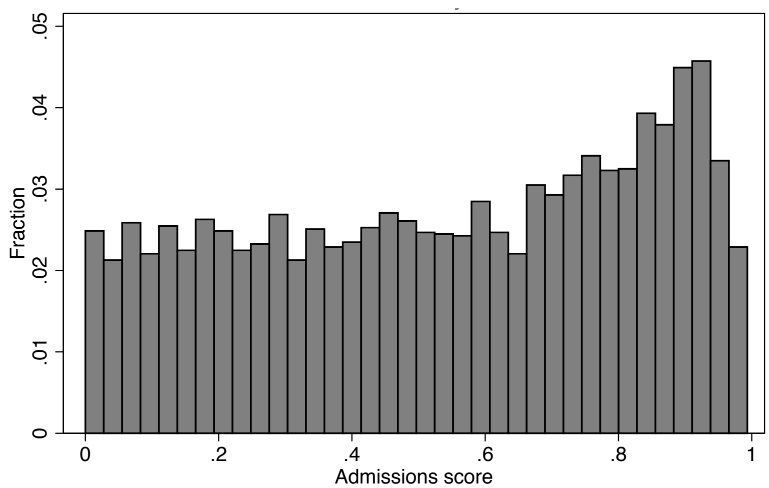

Table 2 shows summaries of the distribution of admissions scores for California resident applicants and admitted students in 2015 and 2016, both for all applicants and for applicants from the underrepresented groups. It indicates that the distribution of admissions scores changed somewhat between years, with more students with very low and very high admissions scores in 2016 than in 2015. These roughly offset each other, however, and average admissions scores, both overall and for applicants from underrepresented groups, were quite similar in 2016 as in 2015. Overall, changes in the distribution of observable characteristics among applicants, on its own, would not likely have produced substantial changes in application outcomes.

Table 2. Distribution of admissions scores in 2015 and 2016

California residents, excluding recruited athletes

	All students				**Underrepresented groups**			
	Applicants		**Admits**		**Applicants**		**Admits**	
	2015	**2016**	**2015**	**2016**	**2015**	**2016**	**2015**	**2016**
Mean	0.17	0.17	0.49	0.45	0.13	0.13	0.42	0.40
Fraction below 1%	33%	38%	1%	3%	39%	45%	2%	5%
Fraction below 5%	53%	57%	7%	12%	61%	64%	11%	15%
5th percentile	0.00	0.00	0.04	0.02	0.00	0.00	0.02	0.01
10th percentile	0.00	0.00	0.08	0.04	0.00	0.00	0.05	0.03
25th percentile	0.00	0.00	0.22	0.14	0.00	0.00	0.15	0.10
50th percentile	0.04	0.03	0.49	0.43	0.02	0.02	0.38	0.33
75th percentile	0.22	0.22	0.76	0.75	0.14	0.13	0.67	0.67
90th percentile	0.61	0.65	0.89	0.90	0.45	0.49	0.85	0.87
95th percentile	0.79	0.83	0.93	0.94	0.67	0.72	0.91	0.92
99th percentile	0.94	0.95	0.97	0.98	0.90	0.92	0.96	0.98

Figure 5 shows the share of students at each admissions score level who were considered under AR in 2015 (solid line) or asked for LORs in 2016 (dashed line). The feature that jumps out the most is that the LOR program was massively larger than AR had been – even at the lowest admissions scores, where AR students are concentrated, the share of 2016 students from whom LORs were requested greatly exceeds the share considered under AR in 2015. The second thing to notice is that, while AR students were concentrated around the 0.3 mark – that is, these students were, based on observables alone, much weaker than the average applicant – LORs were used most among students with much higher admissions scores, 50-80%. In this range, nearly all 2016 applicants were asked for LORs, but 10% or less were considered under AR in 2015.

Figure 5. Share of California resident applicants considered under AR or asked for letters

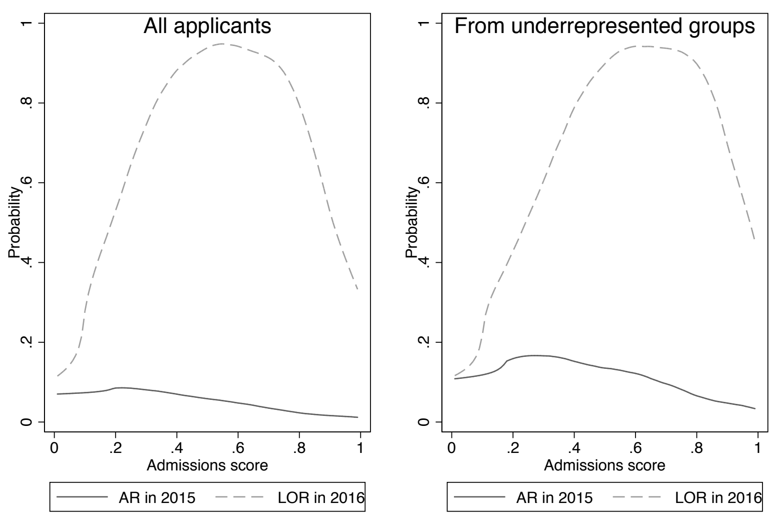

The waitlist
Beyond the AR/LOR shift, another important change in 2016 was a greatly expanded use of the waitlist. The share of applicants offered positions on the waitlist nearly doubled (from 4.5% to 8.6%) in 2016. Many students declined these offers – fully 4% of 2016 applicants declined offered positions on the waitlist, as compared to 1.5% in 2015. This greatly complicates comparisons of 2015 and 2016 outcomes, as some of the students who turned down positions on the waitlist in 2016 would have been admitted outright in 2015.

Many Berkeley applicants are choosing between Berkeley and other excellent universities, and many who are accepted wind up going elsewhere. In 2015, less than half of admitted students came to Berkeley, and this share was smaller for stronger applicants. In many cases, students will have already decided to enroll elsewhere by the time Berkeley's initial admissions offers are made. Consider, for example, a student admitted elsewhere under an early decision program. In principle, this student might withdraw her Berkeley application, but this has not been easy to do, and in any event some students might not bother with nothing at stake.

If these students are chosen for initial admission, they count toward our statistics on admitted students. But if they are offered positions on the waitlist, they are likely to turn down this option, and thus will not count as admissions even if they would eventually have been admitted off the waitlist. Thus, a shift of some admissions from the first to the second round, as occurred in 2016, will reduce the share of these uninterested students in the admit pool (and, as a side effect, raise the enrollment rate among those admitted).

Table 3 shows the distribution of admissions outcomes, aggregating non-athlete California residents across each of the separate applicant pools (for different colleges and divisions) but separating the different stages. This illustrates the potential distortion caused by the waitlist: Note that the share of underrepresented applicants who were admitted rose by only 0.7 percentage points (and the share admitted in the first round *fell* by 0.2 p.p.), while the share who enrolled rose by 1.9 p.p. This is because the "yield" rate for admitted students rose by 4 percentage points, from 47% to 51% overall, and by 10 p.p., from 59% to 69%, for admitted students from the underrepresented groups.

Table 3. Stages of the admissions process
California residents, excluding recruited athletes

	Share of all applicants (%)			Share of applicants from 4 groups (%)		
	2015	2016	Change	2015	2016	Change
Initial admit	16.1	18.1	2.0	13.3	13.0	-0.2
Initial deny	79.4	73.3	-6.1	83.8	78.9	-4.8
Offer waitlist	4.5	8.6	4.1	3.0	8.0	5.1
Decline waitlist offer	1.5	4.0	2.5	1.2	4.3	3.1
Accept waitlist offer	3.0	4.6	1.6	1.8	3.8	1.9
Admitted from WL	2.7	3.0	0.3	1.6	2.6	1.0
Not admitted from WL	0.3	1.6	1.3	0.2	1.2	1.0
Ultimate outcomes						
Admitted	18.8	21.1	2.3	14.9	15.6	0.7
Enrolled (SIR)	8.8	10.7	1.9	8.8	10.7	1.9
Did not enroll	10.0	10.4	0.3	6.1	4.9	-1.2
Denied	79.7	74.9	-4.8	83.9	80.1	-3.9
Withdrew after WL offer	1.5	4.0	2.5	1.2	4.3	3.1

Table 3 reinforces my concern that issues of self-selection are quantitatively important. Many students who under 2015 processes would have been admitted but gone elsewhere were instead in 2016 offered positions on the waitlist only to decline the offers – perhaps as many as 1% of 2016 applicants. These students would have been counted as admits in 2015 but not in 2016. Importantly, this affects the statistics for underrepresented students, as these students were disproportionately likely to decline positions on the waitlist: 47% of all students offered positions on the waitlist declined them, but this share was 53% for students from the underrepresented groups.

Unfortunately, there is no admissions measure that is perfectly comparable across years – in particular, the composition of both the pool of initial admits and the pool of eventual admits is affected by the increased use of the waitlist, even with no other changes. In this report, I present results for four different measures, each imperfect:

- Initial admissions (including both Fall and Spring admits)
- Admitted in the initial round or offered a position on the waitlist
- Ever admitted, either in the initial round or off the waitlist
- Admitted and enrolled (as proxied by filing an SIR, either for Fall or Spring enrollment)

The last of these, of course, reflects student as well as campus decisions (as does the third, which reflects student decisions to accept a spot on the wait list). Nevertheless, in my view it is the closest to comparable across years. If students' propensity to accept Berkeley admissions offers, if made, did not change across years, and if a student who would accept an initial offer is not put off by being admitted off the waitlist, changes in the pool of enrolled students can be attributed to changes in admissions criteria.

Augmented Review and Letters of Recommendation

This report focuses on the Augmented Review and Letters of Recommendation components of the admissions process. Table 4 shows the number of students considered under AR in 2015, the number asked for LORs in 2016, and the outcomes of each group of applications. As already indicated by Figure 5, this makes clear that the LOR program was much, much larger than the AR pool, which I understand was kept small due to the enormous staff time required to review AR applications.

Table 4 also shows that 15% of students who were asked for letters did not request any. This might reflect what many were concerned about, that students would not have access to teachers willing to write letters. But the above self-selection discussion points to another potential explanation: Students admitted Early Decision elsewhere, and others not very interested in Berkeley, might simply not have bothered to request letters. For this reason, I do not emphasize requests of or receipt of letters as outcomes, and focus on the impact of the letters *request* on the student's likelihood of being admitted or of enrolling.

Table 4. Outcomes for AR and LOR students
California residents, excluding recruited athletes

	All applicants		Underrepresented groups	
	AR in 2015	**LOR in 2016**	**AR in 2015**	**LOR in 2016**
Number affected	3,046	14,406	2,793	6,337
Share of all applicants	7%	32%	12%	26%
LOR outcomes				
Any requested		88%		85%
Two requested		77%		72%
Any received		87%		84%
# received = # requested		83%		79%
Two received		73%		67%
Admissions outcomes (shares)				
Initial admit	27%	40%	27%	35%
Admit or WL offer	33%	59%	33%	54%
Ever admit	30%	46%	30%	40%
Admit and matriculate	17%	23%	16%	19%
Admissions outcomes (numbers)				
Initial admit	816	5,833	759	2,200
Admit or WL offer	995	8,438	923	3,414
Ever admit	907	6,672	841	2,558
Admit and matriculate	503	3,364	459	1,211

Assessing the letters of recommendation component of the change

I begin my analysis by attempting to assess the impact of the letter of recommendation component of the 2016 admissions process. Analyses by the admissions office have contrasted those who provided letters to those who were asked for letters but did not provide them. These are useful in understanding which students may have trouble obtaining letters (though as noted above, a student who does not obtain letters might just have decided to go elsewhere). I take a different approach: The factor that is under Berkeley's control is whether applicants are *asked* for letters, so I attempt to uncover the impact of this on admissions outcomes, without trying to distinguish effects coming from difficulty in obtaining letters from those coming from (for example) the submission of weak letters.

As it happens, the way that the LOR process was implemented allows for a compelling analysis of the LOR request, based on a comparison of students asked for LORs with nearly identical students who just missed being asked. There were two ways that students were selected for LORs:

- The admissions office estimated a statistical model to predict read scores in 2015.[3] Those 2016 applicants whose read scores were predicted to be 2.5, 2.75, or 3 were automatically asked for LORs.
- Applicants who were scored by the first reader as "Possible" (on the 2016 Yes/Possible/No scale), when this was done before the deadline for requesting LORs, were asked for LORs. About two-thirds of initial reads scored "possible" were completed by the deadline.

Approximately 80% of those asked for LORs were identified by the first method, and 40% by the second method. (20% were identified by both methods.)

Although 2015 read scores used discrete categories (with each reader assigning a score of 1, 2, 2.5, 3, or 4, and with lower numbers given to stronger applicants), the statistical model used for the first method generated continuous predictions – that is, an applicant might have been predicted to get a read score of 2.47. Students with predicted read scores between 2.38 and 3.26 were all asked for letters, while students with read scores just outside this range were asked only if they were captured by the second method.

Figure 6 shows the share of applicants at each predicted read score who were asked to submit letters. For applicants with predicted scores between 2.38 and 3.26, this was 100%. But only about 60% of students who were just a bit stronger than this range (predicted scores of 2.37) were asked for letters due to the first reader's score, and only about 3% of students who were just a bit weaker than this range (predicted scores of 3.27) were asked.

These sharp breaks permit a "regression discontinuity" analysis of the effect of the LOR request on admissions outcomes. Students with predicted scores of 2.37 are essentially identical, on average, to those with predicted scores of 2.38, and would almost certainly have had very similar admissions outcomes had LORs not been requested for so many more of the latter. Thus, any difference in their outcomes can be attributed to the LOR request.[4]

---

[3] The predicted read score is similar in spirit to, and relies on the same variables as, my admissions score discussed above. They differ because lower read scores are better, and because the predicted read score weights characteristics to predict the 2015 read score, while the admissions score weights the characteristics to predict first-round 2015 admissions. The correlation between the two scores is around -0.85.

[4] Another strategy might be to compare admissions outcomes of those with predicted read scores between 2.38 and 3.26 in 2015 and 2016, relative to those outside this range. Unfortunately, I do not have access to predicted read scores for 2015 applicants, and have been unsuccessful in re-creating them. I expect to be able to eventually, but as I write this the key staff person (Greg Dubrow, Director of Research and Policy Analysis in the Office of Undergraduate Admissions) is on vacation. I thus defer this to future study.

Figure 6. Fraction of California resident applicants asked for letters of recommendation, by predicted read score

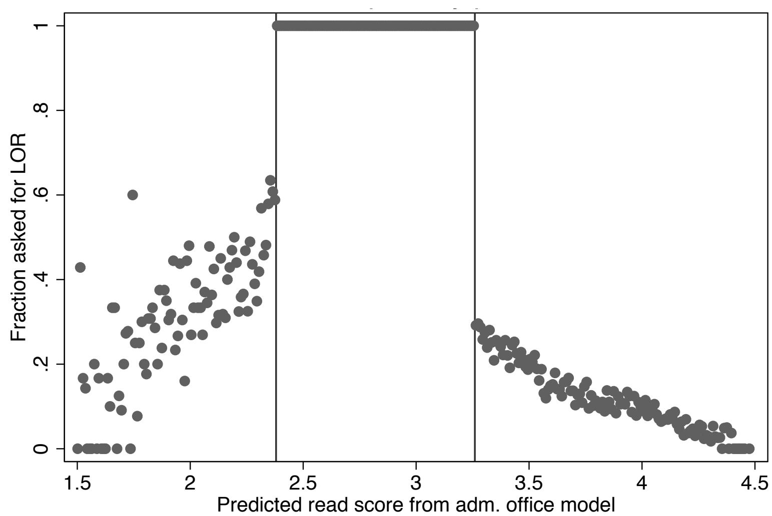

Figure 7 shows 2016 admissions outcomes as functions of the predicted read score. Not surprisingly, students with lower predicted read scores are more likely to be admitted. But notice the area around the vertical lines at 2.38 and 3.26, where I allow for the average admissions outcomes to change discontinuously. Applicants to the left of the first line, only 60% of whom are asked for LORs, are somewhat more likely than students to the right of the line, all of whom were asked for LORs, to be admitted, to be invited to the wait list, and to enroll. Because there is no reason to expect differences in outcomes between these students except for the difference in their LOR treatment, this is clear evidence that for students with predicted read scores around 2.38 – stronger than 95% of applicants and 81% of admits in 2016 – being asked for an LOR reduced the probability of admission.

Figure 7. Admissions probabilities by predicted read score, 2016, all California resident applicants

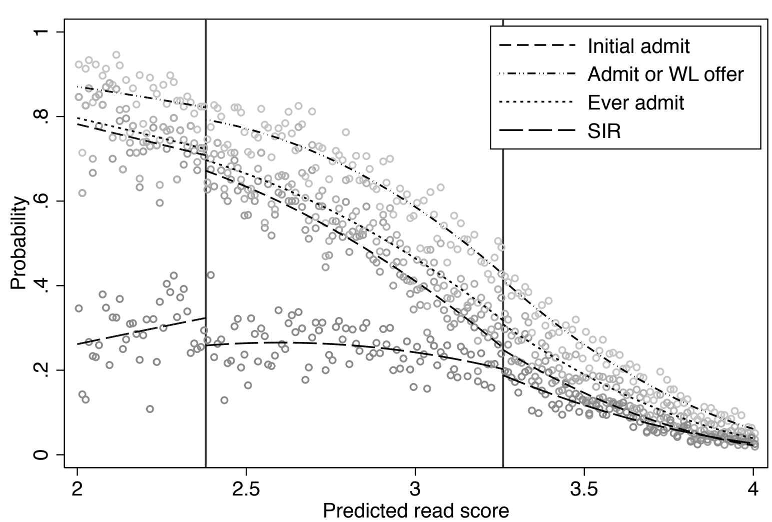

Now turn to the second line. There may be smaller discontinuities here, generally pointing to *higher* admissions probabilities for the 3.25s who were definitely asked for letters than for the

3.27s who had only a 30% chance of being asked. The discontinuities are smaller here and may be entirely attributable to statistical noise. However, there were many more applicants with predicted scores in this range than around 2.38, so even a small effect of LORs on 3.25 students would be quantitatively important.

It is not possible to say whether the LOR impacts seen in Figure 7 reflect better admissions decisions or worse, as it is not possible to identify the specific students who were admitted if not asked for letters but not admitted otherwise (or vice versa). But it is worth noting that this is exactly what we would expect if the LORs provided useful information – some students who would have gotten the benefit of the doubt due to their strong numeric credentials without LORs were revealed by the LORs to be weaker than they appeared, while others who would not have gotten the benefit of the doubt were revealed by their LORs to be worth admitting. Of course, the results are also consistent with the possibility that some strong students were denied admission because they were unable to provide LORs.

Figure 8 repeats the analysis for applicants from the underrepresented groups. These data are noisier, due to the smaller number of observations here. (Dots around the 2.38 threshold represent an average of 60 applicants in Figure 7, and only 10-15 in Figure 8.) There is no sign here that the stronger students were hurt by the LOR request, on average: Those on the left of the 2.38 line are admitted at essentially the same rate as those on the right.

Figure 8. Admissions outcomes by predicted read score, 2016, California resident applicants from four underrepresented groups

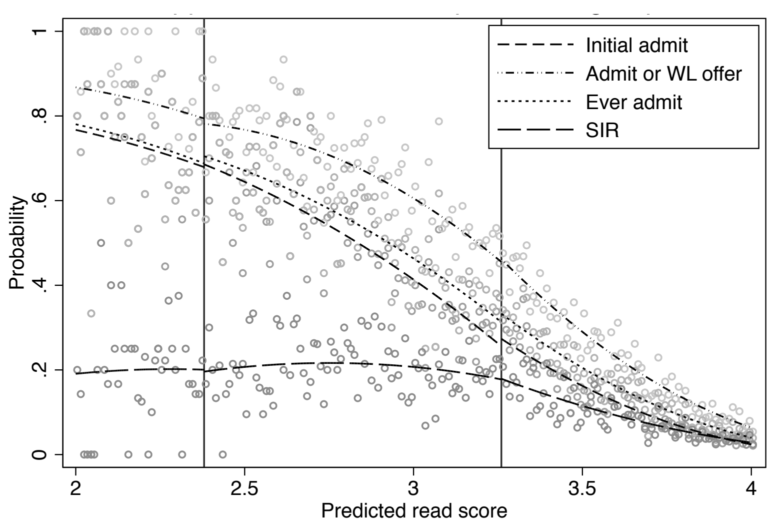

Table 3 presents quantitative estimates of the effect of an LOR request on admissions outcomes, separately for those near the 2.38 threshold and for those near the 3.26 threshold.[5]

[5] In technical terms, these are instrumental variables estimates from a fuzzy regression discontinuity design. They reflect the local average effect of the LOR request on students near the relevant threshold.

In the full applicant pool, strong students from whom LORs were requested were 8-10 percentage points less likely to be admitted than they would have been had they not been selected for LORs. For weaker students, the effect was to increase admissions chances by 2-4 percentage points. (Note that there are about four times as many students near the 3.26 threshold as near the 2.38 threshold, so the implied number of students admitted due to letters near the former is comparable to the number denied due to letters near the latter.) For underrepresented applicants, the 2.38 threshold effect is smaller, suggesting that LOR requests were not harmful to strong students from this group.

Table 3. Regression discontinuity estimates of the effect of LOR requests on the probability of admission

California residents, excluding recruited athletes. Standard errors in parentheses

	Initial admit	**Admit or WL offer**	**Ever admit**	**SIR**
All applicants				
Stronger applicants (low predicted read scores)	-10.4%	-6.3%	-7.6%	-13.6%
	(5.5%)	(4.7%)	(5.4%)	(5.2%)
Weaker applicants (high predicted read scores)	+2.0%	+3.6%	+4.2%	+3.6%
	(2.3%)	(2.5%)	(2.4%)	(2.0%)
Applicants from four underrepresented groups				
Stronger applicants (low predicted read scores)	-6.5%	-6.5%	-4.1%	-2.3%
	(9.4%)	(8.3%)	(9.3%)	(7.9%)
Weaker applicants (high predicted read scores)	-1.8%	+0.7%	+1.6%	+3.2%
	(3.8%)	(4.1%)	(3.9%)	(3.2%)
Net impact on number of admitted students				
All applicants				
Number	-304	54	40	-300
Proportion	-4%	0%	0%	-6%
From four underrepresented groups				
Number	-181	-60	19	123
Proportion	-5%	-1%	1%	5%

The final rows of the table attempt to estimate the net impact of LORs on admissions at each stage – positive numbers indicate a positive net effect, and negative numbers a negative net effect.[6] Focusing on the last column, we see that the LOR requirement raised the number of underrepresented students who enrolled by 123, while reducing the number of enrollees from other groups by 423. These are relatively small numbers, but show no sign of negative effects of LORs on diversity and indeed imply that LORs *raised* the underrepresented share of enrolled students by several percentage points.

As noted above, we cannot tell whether the LOR aspect of the 2016 procedures led to better or worse decisions. But there is no indication that it reduced admissions chances for underrepresented or weaker students, who seem most likely to have faced challenges in obtaining suitable letters.

Assessing the Augmented Review component of the change

The second major question I address is whether the elimination of Augmented Review made it harder for the types of students formerly identified for AR to be admitted, or whether other changes made to admissions processes were able to compensate for the lack of a separate AR pool.

Unfortunately, there is no regression discontinuity research design available for assessing the impact of Augmented Review. Moreover, there is no way to identify in the 2016 data exactly which applicants would have been referred to AR had it been in place, so we cannot compare outcomes for this pool over time.

As an alternative, I identify candidates who, based on their characteristics, would have been likely to be referred to AR in 2015, and examine how their outcomes changed over time relative to others who, based on their characteristics, would have been unlikely to have been referred to AR. Specifically, I create yet another score from the same variables considered to date, this one representing the likelihood that a student with these characteristics would have been referred to and confirmed for AR in 2015.[7]

---

[6] These calculations require rather heroic assumptions. I assume that the estimated effects found at the two discontinuities extend identically outside them, and I linearly interpolate effects between the discontinuities. I make no allowance for sampling error in this extrapolation.

[7] Specifically, I fit a logistic regression, using 2015 data, where the outcome is an indicator for having been referred to and confirmed for AR and predictors are a quartic in the predicted read score, indicators for the three disadvantage groups and interactions among them; and indicators and separate quadratics in the predicted read score for those with 1, 2, or 3 disadvantage factors. (A more flexible model that includes all of the underlying variables in the model used to generate the predicted read score does not generate meaningfully better predictions.)

50% of applicants have AR probabilities below 2%, but 10% have probabilities above 20%. Figure 9 shows how the AR probability score relates to the read score, separately for students with different numbers of disadvantage factors. (The multiple lines in each series represent different combinations of which disadvantage factors the student has.) Across all demographic groups, students with predicted read scores near 3.25 are more likely to be confirmed for AR than those with higher or lower predicted read scores. For any given predicted read score, AR probability scores are higher for those with more enumerated disadvantage factors. For students with three disadvantage factors and predicted read scores between 2.64 and 3.91, AR probability scores are above 0.2, and sometimes substantially so. Students who have only two disadvantage factors must have predicted read scores in a narrower range, between about 3.1 and 3.7, to achieve AR probability scores this high, while students with zero or one disadvantage factors never have AR probabilities above 0.12.

Figure 9. Estimated probability of Augmented Review by predicted read score and number of underrepresentation factors, California residents in 2015

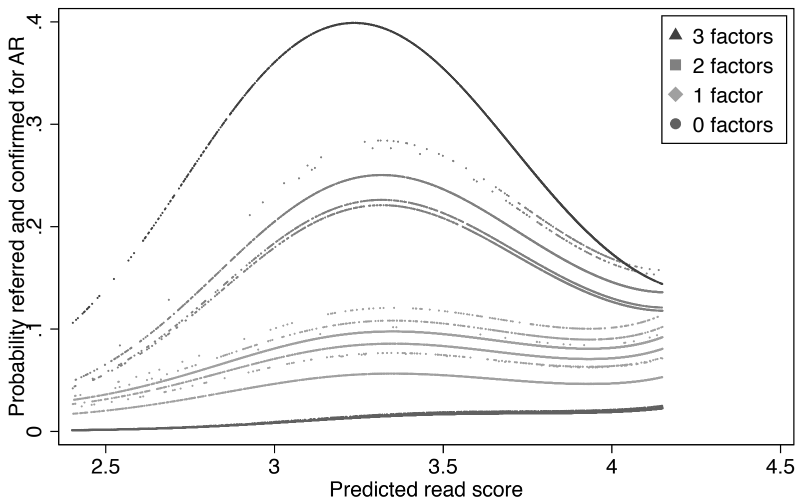

Unfortunately, while this prediction model is fairly successful, it does not achieve a sharp distinction between AR and non-AR students – even the students with the absolute highest AR probabilities have only a 40% chance of being confirmed for AR. In light of this, I consider two definitions of students most likely to be considered in the AR pool:

- Students with AR probability scores above 0.2 (10% of applicants and 38% of those confirmed AR in 2015)
- Students who are low income, first generation, *and* from low API schools, with AR probability scores above 0.2 (7% of applicants, and 28% of those confirmed for AR in 2015).

Table 4 shows the admissions outcomes for students in each of these groups in 2015 and 2016, as well as for their complements (students with lower AR probability scores). Relative changes at all margins except initial outcomes are positive or close to zero. (Across each definition, the relative changes are most in favor of the high-risk group when the outcome is admission or the offer of a waitlist spot – it seems that many students who would have been in the AR pool in 2015 were offered waitlist spots in 2016 but not admitted, either because they were not

Table 4. Probability of admission, by AR probability, 2015 and 2016

California residents, excluding recruited athletes

	Initial admit	Admit or WL offer	Ever admit	SIR
Definition 1: Predicted AR probability >20%				
High probability AR				
2015	17.5%	21.5%	19.4%	9.9%
2016	18.4%	32.5%	22.5%	12.3%
Change	0.9	11.0	3.1	2.5
Change for low probability AR group	2.7	6.7	2.9	2.3
Difference in changes	-1.8	4.3	0.2	0.2
Definition 2: 3 disadvantage factors and predicted AR probability >20%				
High probability AR				
2015	16.3%	19.2%	17.5%	8.9%
2016	16.0%	28.5%	19.7%	11.0%
Change	-0.3	9.3	2.2	2.0
Change for low probability AR group	2.7	6.9	3.0	2.3
Difference in changes	-3.0	2.4	-0.8	-0.3
Impact of AR elimination on number admitted				
Definition 1	-248	603	30	25
Definition 2	-396	318	-105	-38

selected from the waitlist or because they declined the offer.) Neither of the definitions indicates meaningful effects of AR on the number of students who enrolled, and in general, it is hard to discern changes of meaningful magnitude in the admissions outcomes of AR-type students between years, suggesting that changes in other aspects of the admissions process enabled these students to get the extra consideration in 2016 that they got through AR in 2015.

Toward an understanding of the overall impact of admission process changes

The results thus far suggest that LORs, if anything, increased diversity of the entering class in 2016, and that the elimination of AR had a trivial effect. But there were a number of other changes made in 2016, and overall the impact was somewhat less than satisfactory – the share of students from underrepresented groups among admitted students fell, though the share among students who enrolled was stable. In this final section, I present some analyses of overall outcomes that point to possible contributing factors.

Figure 10 shows estimates of the share of students at each application score who were successful in 2015 and 2016, for each of the definitions of success defined earlier. Here, I adjust the 2015 applicant pool to match the distribution across colleges seen in 2016, to remove the influence of shifts across admissions processes that are more or less competitive. We see that weaker applicants (as measured by admissions scores between 0.1 and 0.4) were more likely to be admitted in 2016 than in 2015, but stronger applicants (scores above 0.7) were somewhat less likely to be admitted. The latter change disappears when we include waitlisted students with initial admits, but it reappears and is even larger when we examine the share of applicants who were ever admitted (counting as failures those who were rejected outright as well as those who were offered waitlist spots but either declined them or were not admitted off the waitlist). By contrast, the strongest applicants were *more* likely in 2016 than in 2015 to matriculate.

Figure 10. Admissions probabilities by admissions score, California residents in 2015 and 2016

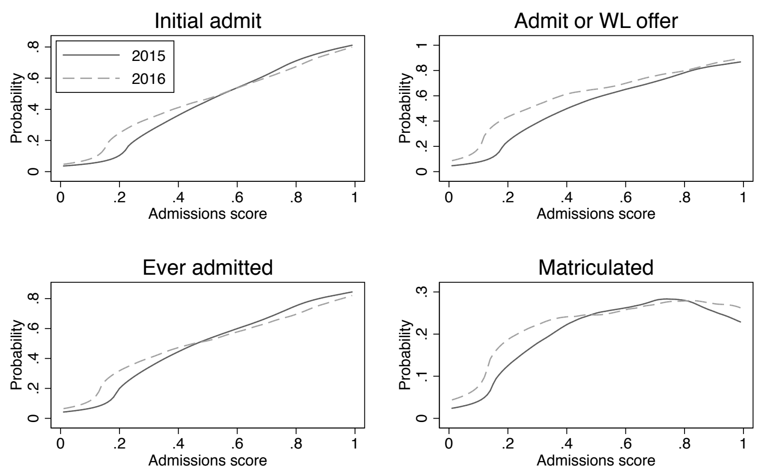

The contrast between the 3rd and 4th panels is informative – it suggests that some very strong students who were admitted in 2015 but matriculated elsewhere were reclassified as non-admits in 2016, either because the admissions decision took account in some way of the likelihood of matriculation or because these students dropped out at the waitlist stage. In any event, we see that both the very strongest and weaker applicants were more likely to matriculate in 2016 than in 2015, while there was little change for those in the middle range (between 0.4 and 0.8).

Figure 11 repeats this exercise, this time only for applicants from the four underrepresented groups. As in the overall pool, we see increased admissions chances in 2016 for applicants with admissions scores around 0.2. But here we see fairly dramatic declines in admissions of applicants with scores above 0.5 that translate into reduced matriculation as well.

Figure 11. Admissions probabilities by admissions score, California resident applicants from underrepresented groups

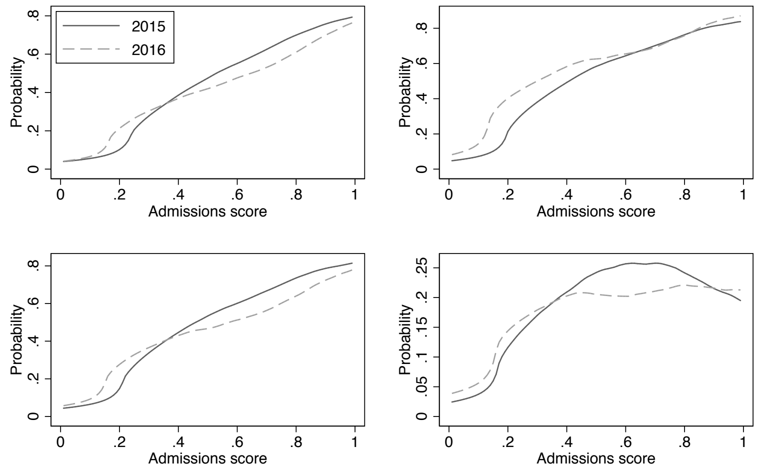

Evidently, something in the admissions processes used in 2016 reduced the admissions chances of the students from underrepresented groups who were, by 2015 standards, the strongest in their observed characteristics. One candidate explanation is the use of non-cognitive scores, which might have been subtly biased against students from underrepresented groups; another is that readers might have put less weight on the factors measuring students relative to their schools in evaluating 2016 applications. Unfortunately, in the limited time I had to prepare this report, I was not able to get to the bottom of this change. It bears further study. It is worth noting, however, that Figures 10 and 11 constitute strong evidence against the view that the elimination of AR played a major role – recall that AR students are concentrated around admissions scores near 0.2, where admissions chances went up the most in 2016.

Conclusion

Berkeley admissions outcomes for underrepresented students were, by some measures, disappointing in 2016: Although more were admitted overall, and their share of enrolled students was steady, they made up a lower share of admissions offers and particularly of first-round admissions offers. It was natural to wonder whether the elimination of Augmented Review and the addition of Letters of Recommendation contributed to this change.

My analysis offers no support for these possibilities. Letters of recommendation seem to have hurt the admissions chances of otherwise-strong applicants not from underrepresented groups, with smaller or no effects on applicants from those groups, and thus to have raised the share of underrepresented students among admissions. It is difficult to identify any clear effect of AR either way, but in any event it was small. The explanation for the change in outcomes in 2016 must lie elsewhere, in one of the other changes made to admissions processes.