How the Justice Department Built Its Title VI Case Against UC San Diego’s Medical School, What the Hardship Essay Evidence Actually Shows, and What Happens to Physician Pipeline Numbers If It Wins

JUL 31, 2026

Share

🎧 Part I Podcast free on Spotify.

Thoughts on Healthcare Markets & Technology

Part I: How the Justice Department Built Its Title VI Case Against UC San Diego’s Med School, What the Hardship Essay Evidence Actually Shows, & What Happens to Physician Pipeline Numbers If It Wins

DOJ sent UC San Diego’s med school an 8-page findings letter in July 2026 concluding intentional racial discrimination in admissions. The case rests on one email...

Listen now

24 minutes ago · Thoughts on Healthcare

🎧 Part II Podcast episode for paid subscribers only. Also available on Spotify.

Thoughts on Healthcare Markets & Technology

Part II: How the Justice Department Built Its Title VI Case Against UC San Diego’s Med School, What the Hardship Essay Evidence Actually Shows, & What Happens to Physician Pipeline Numbers If It Wins

DOJ sent UC San Diego's med school an 8-page findings letter in July 2026 concluding intentional racial discrimination in admissions. The case rests on one email...

Listen now

23 minutes ago · Thoughts on Healthcare

To listen to paid episodes in Apple or Spotify, li
settings on those platforms (instructions inside th
Podcast).

Abstract

- Eight pages, seventy five documents, and
- What the data shows, and what it very m
- The legal standard nobody wants to say
- Proxy math, or why hardship was never
- What a settlement looks like and who ge
- Pipeline arithmetic, run honestly
- Downstream for everyone who is not a d
- On July 20, 2026, DOJ Civil Rights sent UC San Diego an 8 page findings letter concluding that its School of Medicine intentionally discriminated by race in admissions, in violation of Title VI and SFFA.
- The investigation opened March 25, 2026. The school produced 75 documents plus applicant level data for entering classes 2020 to 2025. It produced nothing in response to a request to explain score gaps between racial groups.
- Core allegation: a "hardship" essay prompt was used to infer which applicants were URiM, then those applicants were sorted into hardship subgroups within MCAT/GPA bands A, B, and C, which raised their odds of an interview invite. Reviewers also had race visible in the 2023-24 and 2024-25 cycles.
- Statistical showing: Black admit rate more than 2x white in the 2023, 2024, and 2025 entering classes (2.45% vs 6.18%, 2.10% vs 5.25%, 1.96% vs 5.14%), plus median credential gaps reaching 17 MCAT percentile points in 2025.
- The document evidence is strong. The statistical evidence, as published, is thin: no model, no controls, no coefficients, medians conditioned on admission.
- Fourth medical school finding after UCLA, Yale, and UC Davis. 15 more investigations opened June 4, 2026.

- Realistic downstream: national MD matriculation already fell 11.6% for Black and 10.8% for Hispanic students in the first post-SFFA cycle while applications from both groups rose. Steady state Black physician share converges toward roughly 8%, against 13.7% of population.
- The binding constraints on who actually practices in shortage areas are GME slots, loan repayment, and training location, none of which this fight touches.

Eight pages, seventy five documents, and a very bad email

Start with the physical object, because it tells you a lot. The Justice Department's findings letter to UC San Diego runs eight pages. Eight. The compliance review opened on March 25, 2026, gave the school until June 12 to produce, and closed out with a finding on July 20. That is under four months from open to formal accusation of intentional racial discrimination by a federal civil rights enforcer, which is roughly the amount of time it takes most academic medical centers to schedule a meeting about scheduling a meeting.

What the school handed over in that window was 75 documents, including correspondence, plus applicant level data covering entering classes 2020 through 2025. Seventy five documents. For a six year admissions record at a major public medical school. The Department also asked for an explanation of the deviations in test scores and GPAs across racial groups and got nothing back, a fact the letter mentions twice and then uses as a rhetorical crowbar in the findings section. There is a lesson in there for every general counsel reading this, and it has nothing to do with civil rights law: a thin production does not look like a school with nothing to hide. It looks like a school that ran out of time or nerve.

The substance is worse than the volume. The letter reproduces an admissions committee slide from June 8, 2023, three weeks before the Supreme Court decided Students for Fair Admissions. Under a header reading "Proposals/things to chew on," the first bullet is "Continue as is!" The exclamation point is theirs. The reasoning that follows is that inviting more URiM applicants to submit secondary applications leads to more interviews, and that this does not equal preferential treatment at the final

decision. That is a real argument, sort of, in the way that arguing about where exactly in the funnel the thumb went on the scale is a real argument. It is also the kind of thing that reads very differently on a slide than it does in a deposition.

Then comes the email. On June 15, 2023, the director of admissions wrote that roughly 550 URiM students fall out of Group B and into Group C under the sorting bands, and recommended purposely selecting some students from Group C. He built a chart illustrating a "bump up group C" that could pull as many as 62 additional applicants with substantially lower MCAT and GPA scores into the interview pool, with the stated goal of enrolling more URiM students. A July 31, 2023 follow up to the admissions team elaborated on manipulating batches to prioritize applicants flagged for hardship, disadvantage, or distance traveled, using parental education and income, over applicants with more advantages.

Understand what makes that fatal. Nothing in it is illegal on its face. Prioritizing first generation applicants and low income applicants is lawful everywhere, including California, including after SFFA. The problem is the sentence structure. The email starts from a racial count, moves to a mechanism, and ends at a racial goal. The proxy is not doing independent work in that paragraph. It is doing errand work for the racial variable, and it says so.

The mechanics themselves are almost comically legible. Pre-SFFA, the school sorted into three bands. Group A required an MCAT at or above the 93rd percentile and a cumulative GPA at or above 3.73. Group B required an MCAT at or above the 83rd percentile with a GPA at or above 3.50, or, and here is the part that never should have survived a legal review, an MCAT at or above the 50th percentile with a GPA at or above 3.10 only if the applicant was URiM or from a disadvantaged background. That second prong is a written racial classification in an admissions rubric, in a state that banned race conscious admissions by constitutional amendment in 1996. Group C caught California residents at the 50th percentile and 3.10 who did not land in B.

Post-SFFA, the bands got cleaned up into pure numbers: Group A at MCAT 517 and 3.73, Group B at 512 and 3.50, Group C at 505 and 3.10, no race, no residency prong. Then the director of admissions reviewed applicant submitted material about overcoming a hardship and split each band into two, producing six categories. Batches of 30 went to reviewers. Only a fixed number per batch could receive a rating that guaranteed an interview: 8 out of a Group A batch, 6 out of Group B, 3 out of Group C. And in the 2023-24 and 2024-25 cycles, reviewers could see the applicant's race while rating.

That last detail is the one that turns a defensible holistic review into a hard case. You can argue all day that hardship is a legitimate, race neutral, mission aligned attribute, because it is. You cannot easily argue it while the person scoring the file is looking at a race field.

What the data shows, and what it very much does not

The statistical section is where analytically minded readers should slow down, because it is simultaneously the most quotable part of the letter and the weakest.

The headline is that Black applicants were admitted at more than double the rate of white applicants for three consecutive entering classes. For 2023, 2.45 percent of white applicants versus 6.18 percent of Black applicants. For 2024, 2.10 versus 5.25. For 2025, 1.96 versus 5.14. Then come the median tables for admitted students. In the 2023 entering class, admitted Asian students had a median 3.92 GPA and 517 MCAT at the 94th percentile, white students 3.86 and 518 at the 95th, Black students 3.79 and 514 at the 88th, Hispanic students 3.71 and 511 at the 80th. By the 2025 entering class the spread widened: Asian admits at 3.96 and 519, 96th percentile, white admits at 3.93 and 516, 92nd percentile, Black admits at 3.81 and 510, and Hispanic admits at 3.86 and 510, both at the 79th percentile. Seventeen percentile points between Asian and Black admits at the median, in a cohort where everybody cleared a 505 floor.

Those are real gaps and they are not small. But look at what the letter does not contain. There is no logistic regression. No odds ratios. No marginal effects. No controls for state residency, undergraduate institution, science GPA versus cumulative, MCAT section profile, clinical hours, research output, reapplicant status, or the school's own mission fit scores. No confidence intervals. No sample sizes by cell. The one conditional claim in the entire section, that a Black or Hispanic applicant had a substantially higher likelihood of admission than a white or Asian applicant with the same academic credentials, appears in a single sentence with no supporting output whatsoever.

Compare that to the record in SFFA v. Harvard, where two very well paid economists spent years building competing models with published coefficients, dueling specifications, and arguments about whether to include legacy status and athlete status in the baseline. Whatever anyone thinks of that fight, it was a fight on the merits with numbers you could interrogate. This is an assertion.

Two of the analytical moves in the letter are the kind of thing that would get flagged in any decent internal review. First, comparing raw admit rates across groups without

conditioning on the distribution of the applicant pool tells you almost nothing about the decision rule. Selective medical school admit rates in the low single digits mean the pool is enormous relative to a class of roughly 130, and the composition of that pool by credential band drives raw rates before anyone reads a file. Second, and more seriously, comparing median credentials among admitted students is selection on the dependent variable. Those medians are conditional on having been admitted. They are consistent with a racial preference, and they are also consistent with a dozen other selection rules operating on a pool with different credential distributions by group. You cannot distinguish between those stories from the admitted set alone. You need the applicant set and a model.

There is also a delicious irony sitting in the middle of all this. Title VI private suits require proof of intentional discrimination, per Alexander v. Sandoval in 2001, and this administration has spent considerable effort formally deprioritizing disparate impact theory across federal enforcement. Yet the statistical section here is pure disparate impact reasoning, outcome gaps offered as circumstantial proof of intent. The theory that outcome disparities imply discriminatory intent is, in fact, the theory being used. It just points the other direction now.

None of this means the case is weak. It means the case is not a statistics case. It is an intent case, and intent cases are won on documents. In employment discrimination terms this is close to direct evidence, and direct evidence beats regression analysis every time in front of a jury. The numbers are corroboration. The email is the case.

The legal standard nobody wants to say out loud

Here is the part that gets muddled in most coverage, including some written by people who should know better.

SFFA did not ban considering hardship. It did not ban essays about race. Chief Justice Roberts wrote explicitly that nothing prohibits a university from considering an applicant's discussion of how race affected their life, so long as the consideration is tied to a quality of character or a unique ability that applicant would bring, and the student is treated as an individual rather than on the basis of race. That is a real carve out and it is not narrow in theory.

The very next move in the opinion is the trap door. The Court said universities may not simply establish through application essays or other means the regime held unlawful, and cited an 1867 case for the proposition that what cannot be done directly cannot be done indirectly, that the prohibition is aimed at the thing and not the name.

So the legal line is not about the essay prompt. It is about what the institution does with the inference.

Reading an applicant's account of a rough upbringing and valuing the grit it demonstrates: lawful. Reading the same account, inferring group membership from it, and applying a boost keyed to that group membership in order to move a demographic count: unlawful. The identical prompt, the identical essay, the identical file, and the entire legal outcome turns on the institution's internal purpose. Which means the dispositive evidence is almost always going to live in email, meeting minutes, and dashboards, not in the admissions rubric.

That is a strange and slightly unstable place for the law to land, and it has an obvious perverse consequence. The behavior that gets punished is not preference. It is documented preference. Every general counsel in academic medicine understood this by the second week of July 2023, which is why the sophisticated schools stopped writing things down and the unsophisticated ones kept building charts titled "bump up group C."

A few other legal points worth nailing down. There is no professional school exception waiting to be discovered. Grutter itself was a law school case, and Title VI reaches any program receiving federal funds, which every medical school does through student aid, NIH, HRSA, and Medicare GME adjacency. The argument that graduate and professional admissions get more latitude because the applicant pool is smaller and the mission is more specific is a reasonable policy argument and a dead legal one. California adds a second layer, because Proposition 209 has barred race conscious admissions at public institutions since 1996, meaning UC San Diego is exposed under both a federal statute and a state constitutional provision, and the pre-SFFA Group B rule with the explicit URiM prong is a problem under state law regardless of what the Supreme Court said in 2023.

On remedy, Title VI enforcement requires the agency to seek voluntary compliance first, then either terminate funding or refer for suit. Fund termination is a nuclear option that is politically loud, procedurally slow, and almost never used to completion. The realistic leverage is litigation risk, the cost of discovery, and the reputational drag of a public findings letter, which is exactly why the letters are public.

Proxy math, or why hardship was never going to work

Set the law aside for a minute and just do the arithmetic, because the substitution strategy was doomed on the numbers before it was doomed in court.

Proxies work when the correlation between the proxy and the target is high. The correlation between socioeconomic disadvantage and race in the general population is meaningful. In the medical school applicant pool it is much weaker, because that pool has already been filtered through college completion, a science curriculum, an MCAT sitting, and a few thousand dollars of application costs. What survives that filter is a population where low income and first generation applicants exist in every racial group and where, in absolute counts, the majority of disadvantaged applicants are white and Asian, because the majority of applicants are white and Asian.

So a modest hardship boost distributes mostly to non-URiM applicants and barely moves the demographic count. To move the count you have to make the boost large. A large boost creates visible credential gaps in the admitted class, because you are reaching deep into lower bands. The gaps show up in exactly the median tables the Justice Department published. This is the whole trap in one sentence: a proxy weak enough to be innocent is too weak to work, and a proxy strong enough to work is strong enough to be detected.

UC San Diego's own numbers make the point better than any model could. Roughly 550 URiM applicants fell out of Group B into Group C under the post-SFFA bands. The contemplated remedy pulled up to 62 additional people. That ratio tells you how much force you have to apply to a proxy to get a demographic result, and how visible the application of that force is.

The national data confirms the mechanism. In the first full admissions cycle after SFFA, per AAMC, MD matriculants declined 11.6 percent for Black students, 10.8 percent for Hispanic students, 22.1 percent for American Indian and Alaska Native students, and 4.3 percent for Native Hawaiian and Pacific Islander students. Applications went the other way, up 2.8 percent among Black applicants and up 2.2 percent among Hispanic applicants. More people applied and fewer got in. Whatever happened, it did not happen in the pipeline. It happened at the selection step, and it happened fast, in a single cycle, which is more consistent with schools stripping out anything legally exposed than with any change in applicant quality.

For the 2025-26 entering class, 8.4 percent of matriculants identified as Black or African American and 11.5 percent as Hispanic or Latino, against 0.9 percent American Indian or Alaska Native and 0.4 percent Native Hawaiian or Pacific Islander. Read those numbers with a caveat, though, because AAMC changed its race and ethnicity collection methodology that same year, adding a Middle Eastern or

North African category. Year over year comparisons across that boundary are noisier than most people quoting them acknowledge. Anyone building a trend line through 2025 without flagging the methodology break is selling something.

What a settlement looks like and who gets paid

The Department said it wants a voluntary resolution agreement and will sue if talks fail. Based on the shape of the findings and the pattern across the other four schools, the terms write themselves.

Expect a hard wall between demographic data and anyone evaluating a file, enforced at the system level rather than by policy memo. Expect elimination of any subgroup, tag, or flag that correlates with race and was created after June 2023, which will include the hardship subgroups by name. Expect a prohibition on numeric targets, aspirational or otherwise, and on any internal reporting that measures admissions success by URiM count. Expect multi year applicant level data reporting to the Department, a document retention obligation, certification requirements from senior officers, and probably an outside monitor with a term of three to five years. Expect the essay prompt itself to survive in some sanitized form, because banning applicants from describing their lives is not a thing the Department has asked for and would be its own legal problem.

The pattern matters more than the single case. UCLA and Yale got findings in May 2026, UC Davis in June, San Diego in July. Fifteen additional investigations opened on June 4. Three of the six University of California medical schools have now been formally accused. This is not a one campus enforcement action, it is a sector sweep, and the sequencing suggests the Department found the same fact pattern often enough that it stopped treating each one as a bespoke investigation.

The compliance economics are almost funny. The actual deliverable being purchased across academic medicine right now is not a fairer admissions process. It is document hygiene training for admissions deans. The consultants, outside counsel, and audit shops are going to do extremely well. Whether the underlying decisions change is a separate and much harder question, and enforcement built on internal documents systematically selects for institutions with sloppy email practices rather than institutions with the strongest preferences. That is a real detection bias and nobody involved has an incentive to mention it.

Meanwhile the accreditation scaffolding that used to push the other direction is gone. LCME eliminated Element 3.3, its diversity programs and partnerships standard, in

May 2025, and stripped the remaining diversity language from its standards afterward. ACGME suspended enforcement of its diversity common program requirement the same month and later removed the requirements entirely and closed the department that ran them. So schools spent a decade being told by their accreditor to demonstrate diversity efforts, and now face a federal enforcer treating the residue of those efforts as evidence of intent. Institutions did not build these programs in a vacuum. They built them because the accreditor asked. That context will not be a defense, but it is worth remembering when the coverage frames this as freelance defiance.

Pipeline arithmetic, run honestly

Now the part that actually matters for anyone thinking about workforce, access, or the economics of care delivery over a twenty year horizon.

The United States has roughly a million active physicians. About 6 percent identify as Black and about 6 to 7 percent as Hispanic, against population shares of roughly 13.7 percent and 19 percent. The stock turns over slowly. Something on the order of 24,000 MD matriculants and 9,000 DO matriculants enter each year against a working career of three to four decades, which means the physician workforce is a very long moving average of matriculation shares. In steady state, the composition of the stock converges to the composition of the flow.

That gives you a clean way to think about the ceiling. At an 8.4 percent Black matriculation share, the long run Black physician share converges toward roughly 8 percent. Not 13.7. Even holding the pre-SFFA share, the convergence takes most of a career to show up. A single cycle's 11.6 percent decline removes something on the order of 200 to 250 Black students from a cohort and a similar number of Hispanic students. Sustained over a decade, that is a few thousand physicians against a base of a million. In aggregate terms it is close to a rounding error.

Which is exactly why the aggregate framing is the wrong one. Physician practice location is not randomly distributed. The foundational California work published in the New England Journal in 1996 found that Black physicians practiced in areas where the share of Black residents was nearly five times higher on average than in the areas where other physicians practiced, and that communities with high Black and Hispanic populations were four times as likely to face physician shortages regardless of community income. Later work has been consistent: Black, Hispanic, and Native American primary care physicians are more likely than white counterparts to practice in federally designated shortage areas and to carry higher Medicaid and uninsured panels. County level analyses have found associations between Black primary care

physician representation and higher life expectancy and lower mortality among Black residents.

Be honest about the strength of that evidence, though, because this audience will check. The individual patient outcome literature on racial concordance is genuinely mixed. One systematic review of 27 studies found concordance associated with positive outcomes in about a third, no association in about 30 percent, and mixed findings in the rest, with language, education concordance, and continuity of relationship often outperforming race as predictors. The stronger and more consistent findings sit in communication, trust, preventive service uptake, and adherence, plus the geographic and payer mix effects. There are no randomized trials here and there never will be. Anyone claiming a crisp effect size on mortality from concordance alone is overreaching, and anyone claiming the workforce composition has no access consequences is ignoring fifty years of practice location data.

On the merit side, the AAMC's own validity research says MCAT scores predict preclerkship and clerkship performance, Step 1 and Step 2 CK results, and on time progression, and that prediction operates similarly across groups at a given score. That last point is the one that gets misread constantly in both directions. It does not mean scores are destiny. Published single institution work has found the MCAT component scores explaining something like 17.7 percent of the variance in Step 1 and 12 percent in Step 2 CK. It also does not mean scores are noise. A 510 versus a 518 is a real difference in preparation. It is also the difference between the 79th and the 95th percentile of a group of people who have already completed a science degree and sat for one of the harder standardized exams in American education. Calling the 79th percentile unqualified is a rhetorical choice, not an empirical one. And since Step 1 went pass/fail in 2022, a chunk of the downstream sorting that score gaps used to drive has migrated to Step 2 CK anyway.

Downstream for everyone who is not a dean

If this all resolves the way it looks like it will resolve, and the challenged practices are prohibited across the sector by settlement rather than by a litigated opinion, the realistic effects sort into three buckets.

The first is a modest, durable reduction in URiM matriculation, on the order of what the 2024 cycle already showed, partly locked in by over compliance. Texas showed the pattern after SB17: institutions facing legal risk do not calibrate to the line, they retreat well behind it, because no dean has ever been fired for being too cautious. Expect essay prompts to get blander, expect mission fit language to get scrubbed,

expect socioeconomic factors to survive in a form so thoroughly documented as race blind that they lose most of their bite.

The second is a redistribution of who bears the cost, and it is not the applicants who get most of the attention. It is patients in shortage counties, FQHC and National Health Service Corps staffing, and network adequacy in Medicaid managed care and Medicare Advantage. Brown researchers found one in five Black and Hispanic Medicare Advantage enrollees had no in network Black or Hispanic primary care physician, with roughly four in ten counties lacking any Black physicians in MA networks. That gap does not improve when the training pipeline narrows. And there is a policy contradiction sitting right there in plain sight: CMS is moving toward stratified accountability in Medicare Advantage star ratings through the Health Equity Index, which keys off dual eligibility, low income subsidy status, and disability rather than race, while federal enforcement narrows the workforce that disproportionately serves those same populations. Plans get graded on outcomes for the hardest to serve members while the supply of clinicians who historically go serve them gets squeezed. Somebody is going to have to reconcile those two, and it will probably be a plan actuary rather than a lawyer.

The third is the reframe that ought to matter most and gets the least attention. If the actual policy goal is more physicians practicing in underserved places, medical school admissions is a slow and low leverage instrument. The binding constraints are elsewhere. Residency slots have been capped by Medicare since the Balanced Budget Act of 1997, with a thousand added in the 2021 appropriations act and a couple hundred more in 2023, against an AAMC projection of a shortage reaching as high as 86,000 physicians by 2036. Physicians overwhelmingly practice near where they train, which makes GME geography a stronger lever than admissions composition. Loan repayment through the National Health Service Corps buys practice years in shortage areas directly and measurably. Post baccalaureate and pipeline programs, the ones California built after 209 and credited for the slow recovery in medical school diversity through 2019, operate upstream of the legal fight entirely and are lawful as long as they are open to everyone.

So the entire national argument is being conducted over roughly two to three percentage points of an entering class, in a system where the number of training slots, the location of those slots, and the debt load of the people filling them determine far more about who ends up practicing where. That is not an argument that the legal question is unimportant. It is an argument that if the outcome anyone claims to want is a physician workforce that shows up in the counties that need one, both sides are fighting over the wrong variable with unusual intensity.

Which is a familiar shape in health policy. The fight is loud, the mechanism is legally interesting, the enforcement is real, and the thing that would actually move the number sits in a Medicare payment formula that nobody has opened since 1997

.

← Previous

Discussion about this post

Comments Restacks

Write a comment...