deaths, many of which are in the western United States. Arizona, Colorado, and Washington State had mortality nearly twice the national average or higher in 2020–2022. An assessment of hotspot counties — defined as those that had at least 20 overdose deaths and mortality higher than the national average during this period — revealed 19 counties of high concern. Maricopa County, Arizona, and Los Angeles County, California, had the highest number of fatal overdoses, with 117 and 111 deaths, respectively. School-based interventions could be prioritized in these counties.

An audio interview with Joseph Friedman is available at NEJM.org

The U.S. overdose crisis is affecting Americans at younger ages — even as adolescent substance use is decreasing. In response, parents, physicians, and educators must equip adolescents with the knowledge and tools they need to keep themselves safe, and adolescents must be supported with stronger mental health and addiction services.

Disclosure forms provided by the authors are available at NEJM.org.

From the Center for Social Medicine and Humanities, University of California, Los Angeles, Los Angeles (J.F.); and the Division of Adolescent and Young Adult Medicine, Mass General for Children, and the Department of Pediatrics, Harvard Medical School — both in Boston (S.E.H.).

1. Goldstick JE, Cunningham RM, Carter PM. Current causes of death in children and adolescents in the United States. N Engl J Med 2022;386:1955-6.

2. Friedman J, Godvin M, Shover CL, Gone JP, Hansen H, Schriger DL. Trends in drug overdose deaths among US adolescents, January 2010 to June 2021. JAMA 2022;327:1398-400.

3. Tanz LJ, Dinwiddie AT, Mattson CL, O'Donnell J, Davis NL. Drug overdose deaths among persons aged 10–19 Years — United States, July 2019–December 2021. MMWR Morb Mortal Wkly Rep 2022;71:1576-82.

4. Palamar JJ, Ciccarone D, Rutherford C, Keyes KM, Carr TH, Cottler LB. Trends in seizures of powders and pills containing illicit fentanyl in the United States, 2018 through 2021. Drug Alcohol Depend 2022;234:109398.

5. King C, Beetham T, Smith N, et al. Treatments used among adolescent residential addiction treatment facilities in the US, 2022. JAMA 2023;329:1983-5.

DOI: 10.1056/NEJMp2312084

Accuracy and Equity in Clinical Risk Prediction

Emma Pierson, Ph.D.

When I was 19 years old, I learned that my mother carried a *BRCA1* mutation, which confers a high risk of breast and ovarian cancer. I've spent much of my life since then thinking about risk. As I waited to find out whether I had inherited the mutation, I tried to come to terms with what it meant to have an exactly 50% risk of having it. "Sometimes I feel that I know I don't have it," I wrote in my diary. But a second later, I dismissed that idea as irrational. When I learned that I carried the mutation, I began, at age 25, the annual MRIs recommended for high-risk women. I began thinking about false-positive and false-negative risks from screening — the imperfections of mammography that necessitate MRIs for young women with my mutation, the unreliable results of ovarian cancer screening tests that render them mostly useless. Now, as a 32-year-old trying to plan with my girlfriend when we will have children, I wish desperately that doctors could tell me precisely how my risk of ovarian cancer will climb year by year: How long can we safely wait?

Health risk estimates have shaped my most important life decisions. This has convinced me that we have a profound obligation to patients to predict their risks as accurately as possible. As a professor of computer science and population health, I build algorithms that predict health risks, with a particular focus on ensuring that they perform equitably across groups.

Over the past few years, I've seen a welcome and overdue surge of interest in algorithmic equity. But I've also watched, disquieted, as my field has sometimes — with the laudable intention of ensuring equity — deviated from our basic mandate to predict patients' risk as accurately as possible. Equity and accuracy need not conflict — improving the accuracy of risk prediction can often improve equity as well — but we have made choices that reduce accuracy in the name of equity, ultimately achieving neither.

One key example is the explosion of "debiasing" methods — turnkey technical solutions that can be run on any clinical algorithm, even in settings in which they make little sense, as recent reviews have documented.[1] Frequently, these methods seek to achieve equity by equalizing some measure across two groups. For example, they might try to equalize the fraction of people in each group — men and women, or

Black patients and White patients — identified as high-risk. But often this goal makes no sense. If you are trying to predict the risk of breast cancer, you do not want to equalize the fraction of men and women who are predicted to have high risk — their true risks differ. Although that's a fairly obvious point, many analogous debiasing methods try to equalize other measures of performance, such as false-positive or false-negative rates, across groups. This approach seems intuitively more reasonable than equalizing the fraction of each group predicted to be high-risk, but it turns out to present similar issues. A systematic study of such approaches found that they caused a "nearly-universal degradation of multiple performance metrics."[1]

I believe we should also be cautious about issuing categorical calls to remove race as an input in every health risk prediction. Certainly there are sometimes excellent reasons to remove race from prediction algorithms: past uses have relied on nonrigorous data, stemmed from racist beliefs, increased health disparities, or all of the above.[2,3] But in other cases, removing race can harm the accuracy of risk prediction for the very groups we seek to help. Algorithms that do not incorporate race information have been shown to underpredict diabetes risk for Asian patients and colorectal cancer risk for Black patients. I would argue that we do these patients no favors — and, indeed, risk compounding health inequities — if we knowingly understate their health risks in the name of equity.

Advocates of removing race from algorithms often argue that it is only an imperfect proxy for variables more directly related to risk — for example, genetic ancestry and social determinants of health — and so we must collect better data. I agree entirely. If we want to predict whether a patient has an elevated risk of carrying a sickle-cell mutation, for example, genetic ancestry is more directly relevant than socially constructed race. My work has therefore quantified the limitations of existing race data[4] and called for better data collection.[5]

At the same time, having appealed directly to hospital systems, health insurers, and clinical researchers for such data, I know that obtaining them will be a difficult and slow process. I believe the pragmatic approach is to fight a two-front battle: while we maintain our advocacy and pressure to collect better data and wait for health care systems to do so, we should also do the best we can for patients who need to make life-or-death decisions today. That means making the most accurate predictions we can from the data we have, not the data we wish we had.

Making less-accurate risk predictions by excluding factors that have proven to be predictive albeit highly imperfect — features that are, at best, proxies for the true characteristics of interest — can sometimes carry great human costs. When my mother was diagnosed with breast cancer, her doctors tested her for cancer-causing mutations because they predicted she was at high risk for them. This risk prediction was influenced by the fact that she was Ashkenazi Jewish. The predictive power of this feature is well documented, since Ashkenazi Jews are a small and unusually well-studied group. But use of this feature in predictions is also fraught: its predictive power derives in part from centuries of anti-Semitic murders and expulsions, and it is only an imperfect proxy for the information we would ideally have — my mother's genetic ancestry. But I still believe the doctors made the right call. My mother was 10 times as likely as women who aren't Ashkenazi Jewish to carry these mutations, which increased her risk of ovarian cancer by a factor of 30; ovarian cancer has a 10-year survival of 36%. Had the doctors not tested her, she might not have lived to see us grow up.

Predicting risk as accurately as possible for all patients does not mean ignoring equity. On the contrary, it requires engaging with equity more deeply, rejecting quick technical fixes and overbroad solutions, and working with domain experts to understand the context-specific factors producing inequity. In assessing the equity of a cancer risk prediction algorithm, there are many questions you should ask. Are you training the algorithm to predict the appropriate target, or might your measure of cancer be skewed by underdiagnosis or other biases? Does the patient population you trained your algorithm on reflect the diversity of the patients it will serve, or are biases in care access affecting your data set? There are many other important questions, but all of them require engaging with the factors contributing to inequity in a specific setting. Applying quick technical fixes that weren't developed for that setting, without understanding what they do or whether it's relevant, is often worse than doing nothing at all. Many generic debiasing methods, for example, were motivated by and developed using data from criminal risk predictions — a wildly different setting from cancer risk prediction.

A few weeks ago, I went to see

my oncologist. She is the sort of clinician I'm happy to trust with my life, and she answered my questions about ovarian cancer risk with compassion and precision. By the time I'm 40, she explained, my cumulative risk of ovarian cancer will have risen to 3%. That number sounds so tiny, and yet it's the basis on which doctors warn me to have my ovaries removed by the time I'm 40. I would be furious if, in the name of treating me "equitably," my doctors told me that my risk was 3% when it was really 6%, or 1%, because my most critical life decisions hinge on those numbers. So I want for other patients what I want for myself: give us your best estimate of our risk, engaging deeply with the context-specific inequities that distort risk predictions, so we can decide what to do.

Disclosure forms provided by the author are available at NEJM.org.

From the Department of Computer Science, Cornell Tech, and the Department of Population Health Sciences, Weill Cornell Medical College — both in New York.

This article was published on January 6, 2024, at NEJM.org.

1. Pfohl SR, Foryciarz A, Shah NH. An empirical characterization of fair machine learning for clinical risk prediction. J Biomed Inform 2021;113:103621.

2. Roberts DE. Abolish race correction. Lancet 2021;397:17-8.

3. Vyas DA, Eisenstein LG, Jones DS. Hidden in plain sight — reconsidering the use of race correction in clinical algorithms. N Engl J Med 2020;383:874-82.

4. Movva R, Shanmugam D, Hou K, et al. Coarse race data conceals disparities in clinical risk score performance. In: Proceedings and Abstracts of the 2023 Machine Learning for Health Care Conference, August 11–12, 2023. New York: Machine Learning for Health Care, 2023.

5. Chen IY, Pierson E, Rose S, Joshi S, Ferryman K, Ghassemi M. Ethical machine learning in healthcare. Annu Rev Biomed Data Sci 2021;4:123-44.

DOI: 10.1056/NEJMp2311050

Centering Women of Color to Promote Excellence in Academic Medicine

Christina Mangurian, M.D., Nancy D. Spector, M.D., and Ruth S. Shim, M.D., M.P.H.

The recent U.S. Supreme Court decision striking down the use of affirmative action in university admissions threatens decades of progress in the areas of diversity, equity, and inclusion in academic medicine. Although women accounted for the majority of medical school enrollees in 2022, they represented only 28% of full professors, 23% of department chairs, and 27% of deans that same year,[1] and gender-based disparities in compensation persist at the highest levels of academic medicine.[2] Similarly, 11% of full professors and 13% of department chairs and deans in 2022 were members of racial or ethnic groups that are underrepresented in medicine.[1] Only about 3% of full professors and department chairs were women from underrepresented groups, including Black, Latina, and Indigenous women.[1]

Beyond the Supreme Court decision, leaders in academic medicine have long recognized the importance of diversity, but they have had difficulty with focus. This lack of focus has led to diffusion of efforts and to "condensation," the process by which any number of loosely related elements are consolidated under a single concept.[3] Broad definitions of "diversity" can result in academic institutions making little progress on racial equity specifically.[3] Existing inequities made more evident by the Covid-19 pandemic and the murder of George Floyd led to a renewed focus on recruitment of Black faculty members, cohort hires, and temporary policy reforms to support caregivers, many of whom are women. Although recent efforts represent progress, they continue to be piecemeal and have failed to support retention of faculty members from groups that are underrepresented in medicine by creating a truly inclusive and equitable climate in which all faculty members can thrive.

We propose centering efforts on retaining and advancing women of color (and, in particular, Black women) at multiple levels (including among students, trainees, staff, faculty, and institutional leaders) in academic medicine. By "centering" women of color, we mean that leaders should focus attention, decision making, and policy interventions specifically on dismantling the structural racism and sexism that exist in academic medical institutions. Using an intersectional lens to examine how racism and sexism interact makes it clear that, throughout many industries, the experience of women of color diverges the most from that of White men. Women of color face multiple forms of discrimination and have less access to career-enhancing work than do members of other groups. Women of color also tend to face more external pressures, including responsibility for domestic work and caregiving — not just for children, but for parents and extend-