

Ethical and legal race-responsive vaccine allocation

Bastian Steuer¹  | Nir Eyal^{2,3,4} 

¹Department of Political Science, Ashoka University, Sonapat, HR, India

²Center for Population-Level Bioethics, Institute for Health, Health Care Policy and Aging Research, Rutgers University, New Brunswick, New Jersey, USA

³Department of Health Behavior, Society and Policy, Rutgers School of Public Health, Piscataway, New Jersey, USA

⁴Department of Philosophy, Rutgers University, New Brunswick, New Jersey, USA

Correspondence

Nir Eyal, Center for Population-Level Bioethics, Institute for Health, Health Care Policy and Aging Research, Rutgers University, 112 Paterson St, New Brunswick, NJ 08901, USA.

Email: neyal@cplb.rutgers.edu

Funding information

Longview Philanthropy

Abstract

In many countries, the COVID-19 pandemic varied starkly between different racial and ethnic groups. Before vaccines were approved, some considered assigning priority access to worse-hit racial groups. That debate can inform rationing in future pandemics and in some of the many areas outside COVID-19 that admit of racial health disparities. However, concerns were raised that “race-responsive” prioritizations would be ruled unlawful for allegedly constituting wrongful discrimination. This legal argument relies on an understanding of discrimination law as demanding color-blindness. We argue that a color-blind understanding of discrimination would be hostile only to one of two rationales for prioritizing the relevant racial minorities in settings of racial health disparities. We also propose a method for incorporating appropriate race-responsive concerns that is in many ways ethically and legally superior to ones suggested thus far. That method turns artificial intelligence, thanks precisely to its artificial and “black box” nature (features that underlie recent concerns about artificial intelligence's discriminatory potential), into an instrument of social justice.

KEYWORDS

coronavirus, discrimination, health inequality, racial justice, vaccine

1 | INTRODUCTION

Growing awareness of racial (as well as wider social and demographic)¹ injustice in health care systems around the world has put efforts to ameliorate that injustice on progressive agendas. At the same time, in some countries, there has been pushback against using the tools of affirmative action for this purpose. In the United States, for example, race-based affirmative action has been severely limited, if not banned.²

At first blush, these two developments appear to be on a collision course. If public authorities are legally or politically barred from tackling racial injustices head on by means of affirmative action, this would appear to undermine efforts to redress racial inequalities. We suggest a way forward for important and effective measures against racial disparities that can be pursued even if one subscribes to a color-blind understanding of discrimination that disallows affirmative action. This article remains agnostic on whether discrimination should be understood in color-blind terms. Our proposal is race-responsive in the minimal sense that it responds to racial injustice, but it is (largely) race-neutral in its means as well as in its ultimate goal. While our article focuses on COVID-19, it is as much a contribution to COVID-19-specific policy debates as it is to the larger question of

¹We use the term racial in a broad sense to include racial as well as ethnic groups. This use makes our argument compatible with a variety of views on how to distinguish between racial and ethnic groups.

²The article was written and accepted before the U.S. Supreme Court decisions in *Students for Fair Admissions v. Harvard* and *Students for Fair Admissions v. University of North Carolina*. The decision was widely reported as overturning affirmative action and sets highly restrictive limits. All the while, the majority opinion phrased the decision as following the precedent in *Grutter v. Bollinger*, which permitted race-based affirmative action within limits. The decision's impact remains to be

seen. To be sure, our argument which charts a way for policies in light of color-blind discrimination law is even more relevant in the aftermath of *Students for Fair Admission*.

This is an open access article under the terms of the Creative Commons Attribution-NonCommercial License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited and is not used for commercial purposes.

© 2023 The Authors. *Bioethics* published by John Wiley & Sons Ltd.

how to advance racial justice within the constraints of a color-blind conception of discrimination, which is becoming stronger in the US and potentially elsewhere.³

We first recall racial disparities in COVID-19 and distinguish between two rationales for race-responsive COVID-19 resource allocation, focusing on vaccines. We suggest that one of these two rationales could bridge political divides and alleviate legal concerns. Second, we propose a race-responsive system for vaccine allocation that has important ethical (and legal) advantages compared to simple racial preferences. It remains blind to race in its goals and (almost) all of its operations. That system gives a role for artificial intelligence, surprisingly, thanks to components that have recently earned use of artificial intelligence an association with racial injustice.

As mentioned, our lead examples are COVID-19 resources, particularly vaccines. This includes boosters, vaccines for young children, future vaccines targeting specific variants, and full access to all vaccine products, for example, in the form of increased outreach and health literacy efforts. But we also see broader applications. Our argument holds lessons for additional interventions against COVID-19 and against other public health threats that exhibit a racial gradient. For example, given the general tendency of disasters and, in particular, infectious disease outbreaks, to affect minorities disproportionately,⁴ our lessons remain relevant for preparing protocols to respond to future emerging infection outbreaks, which are expected to become more frequent. Although our legal discussion focuses on the United States, our conclusions have implications elsewhere as well. Color-blindness is but one case of understanding discrimination without distinguishing between invidious discriminatory preferences and preferences that seek to reduce social inequality and hierarchy. Such a symmetric understanding of discrimination is common in many countries—for color, race, caste, religion, or other traits.

2 | FOUNDATIONAL AND DERIVATIVE REASONS FOR RACE-RESPONSIVE VACCINE ALLOCATION

In many countries, a number of local racial or ethnic minorities are both likelier to get SARS-CoV-2 infection and likelier to experience severe outcomes in the event of one. In the United States, the picture is stark for African Americans, Latinx, and Native Americans and less clear for, for example, Asian Americans.⁵ In the United Kingdom, outcomes for Black and Asian British people have been worse than for White British people.⁶ But the evidence also suggests high

variability within these groups. For example, in the United Kingdom, the risk is much higher for Bangladeshi ethnicity than for Indian ethnicity.⁷ While the causes of these disparities are complex, surely racism and structural inequalities are (in)direct driving forces.⁸ It is worth emphasizing that highlighting such racial disparities does not assume race essentialism, the view that race is a biological or genetic essence that all members of a given race share. Rather, our point is that due to largely social and economic factors (themselves at least partly the results of historic and recent racism), members of some racial and ethnic groups are at higher risk.⁹

How should racial disparities influence allocation decisions? These disparities make racial affiliation plainly relevant to these allocation questions, for the simple reason that members of some race groups face elevated risk for COVID-19, for severe COVID-19, and for secondary transmission. This *COVID-19 risk rationale* treats race just like any other trait for which elevated COVID-19 risk is observed. Prioritization for diabetics and for workers outside the home is simply data-driven public health, not invidious discrimination. To the extent that members of some racial groups are at elevated risk, failure to prioritize them on the basis of their race wastes a scarce resource and predictably leads to avoidable morbidity and mortality—both due to the personal risk to them and due to the elevated risk of spreading the infection. On some level, such failure is as plainly irrational and unjust as it would be to fail to prioritize those who face elevated risk on the basis of diabetes or working outside the home. The COVID-19 risk rationale is thus strong and straightforward. But it gives race a *derivative* importance. Race is relevant only because it is a risk factor for COVID-19. The underlying ethical principle is compatible with, for example, utilitarianism, which famously does not incorporate any concern about the distribution of benefits and burdens.

A more extensive role for race in prioritizing local racial minorities is given by a second rationale. On the basis of that second rationale, we should distribute COVID-19 vaccines so as to rectify or mitigate background disadvantage and general health inequities. Philosophically, this distinct *background injustice rationale* might be justified as a form of priority to the generally worse-off or as a priority to victims of unjust social, or even racist, structures.¹⁰ As such, the background injustice rationale gives a *foundational* status to race.

We fully agree that the unfair health inequalities in many countries, including ones impacting the racial minorities hit hardest

[uk/peoplepopulationandcommunity/birthsdeathsandmarriages/deaths/datasets/updatingethniccontrastsindeathsinvolvingthecoronaviruscovid19england](https://www.ukpeoplepopulationandcommunity/birthsdeathsandmarriages/deaths/datasets/updatingethniccontrastsindeathsinvolvingthecoronaviruscovid19england).

³Raleigh, op. cit. note 6, as well as the ONS data cited earlier.

⁴Khazanchi, R., Evans, C. T., & Marcelin, J. R. (2020). Racism, not race, drives inequity across the COVID-19 continuum. *JAMA Network Open*, 3, e2019933.

⁵See Khazanchi, R., et al., op. cit. note 8. For objections to race essentialism in medicine, see Richmond, S. P. II, & Grubbs, V. (2022). How abolition of race-based medicine is necessary to American health justice. *AMA Journal of Ethics*, 24, E226–E232. Our argument in this paper is in line with the proposal by Cerdeña, J. P., Plaisime, M. V., & Tsai, J. (2020). From race-based to race-conscious medicine: How anti-racist uprisings call us to act. *The Lancet*, 396, 1125–1128.

¹⁰See Schmidt, H. (2020). Vaccine rationing and the urgency of social justice in the Covid-19 response. *Hastings Center Report*, 50, 46–49.

³*Students for Fair Admissions v. Harvard*, op. cit. note 2.

⁴This was known long before COVID—see, for example, DeBruin, D., Liaschenko, J., & Marshall, M. F. (2012). Social justice in pandemic preparedness. *American Journal of Public Health*, 102, 586–591.

⁵See, for example, the data collected by the APM Research Lab (<https://www.apmresearchlab.org/covid/deaths-by-race>). Accessed May 3, 2022.

⁶Raleigh, V. S. (2022). Ethnic differences in Covid-19 death rates. *BMJ*, 376, o427 as well as the regularly updated statistics from the Office of National Statistics (<https://www.ons.gov>).

by COVID-19, must be addressed. We also agree that these health inequalities played a role in causing COVID-19 disparities. It is far from obvious to us, however, that the distribution of a single and potentially life-saving resource is the right occasion to address background (health) inequalities. One can be committed to the goals of prioritizing the worse-off, reducing social inequality, and alleviating historic injustice at the level of overall health priority-setting, yet believe that individual allocation decisions on this special topic should reflect medical need alone.¹¹ Consider the following illustration. Colon cancer is more prevalent in males than in females.¹² But women are generally disadvantaged. It is far from obvious that in screening, prevention, or treatment of colon cancer, we should prioritize women over men with equal colon cancer risk, simply in order to compensate women for their general background disadvantages.

A second challenge to the background injustice rationale is that how injustices are rectified and mitigated may matter. Not just any compensation will obviously be appropriate for remedying background injustices. Allowing racial minorities to skip the line at doctors' appointments for conditions that lack a racial pattern might be reasonably considered inadequate for acknowledging racial injustice—as opposed to, say, apology, reparations, and priority in the very same areas where disparities reflect earlier discrimination. Even when conditions have a racial gradient, the rationales being given matter. It is far from obvious that skipping the line for coronavirus therapies would be an adequate response to, for example, injustices in school quality or in policing.

These are not decisive arguments against the background injustice rationale. But they highlight that the rationale is controversial and that there is reasonable disagreement about it. The high stakes of distributing a life-saving resource make it more urgent to avoid ethical mistakes. More importantly, a controversial public justification seems unnecessary when another, more widely sharable, rationale exists—the COVID-19 risk rationale that we pointed out earlier.¹³

Another concern with the background injustice rationale is that it risks agitating parts of the public at a time when building consensus is crucial. It is difficult to estimate this risk empirically. Ground for optimism comes from spring 2021 surveys from the United States that showed a relatively high level of support for racial preferences.¹⁴ However, they also show that racial preferences were disfavored by Republican respondents. Strong opposition from a significant part of the public can be an important political obstacle even if the majority

supports the policy. Interestingly, the same survey found that priorities by postcodes, by contrast, found much wider support. Such priorities are more easily couched in our COVID-19 risk rationale. Moving to the more consensus-building alternative significantly increased support among those otherwise skeptical.

We are not proposing to sweep racial injustice under the carpet. One balanced messaging strategy is to be clear that disparities in COVID-19 risks result from broad background injustices while highlighting that priority is based exclusively on the elevated risk of severe harm from COVID-19. This is the COVID-19 risk rationale in a nutshell. In that spirit, Colorado's COVID-19 vaccination plan “acknowledges long-standing racism, including economic and environmental injustice” that “has created negative health outcomes.” But it then goes on to state, in line with our proposal, that “addressing healthcare inequities *specifically related to COVID-19* was considered for each vaccination phase” (our italics).¹⁵

3 | RACE-RESPONSIVE ALLOCATION AND THE COLOR-BLIND CONCEPTION OF DISCRIMINATION

Policy discussions in the United States, including at ACIP/CDC,¹⁶ and the related NASEM report,¹⁷ proposed using the CDC's Social Vulnerability Index (SVI) as one determinant for vaccine prioritization between geographical areas. (Most U.S. states ended up using a variety of disadvantage indices in their prioritization plans.¹⁸) The SVI includes, as a variable, the share of racial/ethnic minorities in each area. Race would therefore be one of many criteria used to determine who receives vaccines first in the United States.

That proposal raised concerns¹⁹; especially in its more conservative recent composition, the U.S. Supreme Court is committed to a “color-blind” conception of discrimination that treats all racial classifications alike, even ones that (successfully) seek to reduce racial inequality. The concern was that the court would therefore rule out use of the SVI as unlawfully discriminatory. This concern gains even greater traction after the Supreme Court's decision in *Students for Fair Admissions v. Harvard* overturning two affirmative action plans in higher education. One might likewise be concerned that adopting the SVI is politically difficult. Similar legal and political concerns apply to the COVID-19 Community Vulnerability Index (CCVI), which is based on the SVI and adds a number of variables specifically relevant

¹⁵Colorado Department of Public Health and Environment. (2020). *COVID-19 vaccination plan*.

¹⁶See the presentation by Dooling, K. (2021). Implementation considerations for COVID-19 Vaccines. Retrieved May 13, 2022, from <https://www.cdc.gov/vaccines/acip/meetings/downloads/slides-2021-02/28-03-01/02-COVID-Dooling.pdf>

¹⁷National Academies of Sciences, Engineering, and Medicine. (2020). *Framework for equitable allocation of COVID-19 vaccine*.

¹⁸Schmidt, H., Weintraub, R., Williams, M. A., Miller, K., Buttenheim, A., Sadecki, E., Wu, H., Doiphode, A., Nagpal, N., Gostin, L. O., & Shen, A. A. (2021). Equitable allocation of COVID-19 vaccines in the United States. *Nature Medicine*, 27, 1298–1307.

¹⁹Schmidt, H., Gostin, L. O., & Williams, M. A. (2020). Is it lawful and ethical to prioritize racial minorities for COVID-19 vaccines? *Journal of the American Medical Association*, 324, 2023–2024.

¹¹See more broadly Kamm, F. M. (2002). Health and equity. In J. L. C. Murray, J. A. Salomon, C. D. Mathers, & A. D. Lopez (Eds.), *Summary measures of population health* (pp. 692–695). World Health Organization.

¹²Arnold, M., Sierra, M. S., Laversanne, M., Soerjomataram, I., Jemal, A., & Bray, F. (2017). Global patterns and trends in colorectal cancer incidence and mortality. *Gut*, 66, 683–691.

¹³See also Wasserman, D., Persad, G., & Millum, J. (2020). Setting priorities fairly in response to Covid-19: Identifying overlapping consensus and reasonable disagreement. *Journal of Law and the Biosciences*, 7, 1–12.

¹⁴Schmidt, H., Shaikh, S. J., Sadecki, E., & Gollust, S. (2022). US adults' preferences for race-based and place-based prioritisation for COVID-19 vaccines. *Journal of Medical Ethics*, 48, 497–500.

to COVID-19, while retaining racial/ethnic minorities as one constituent variable. Commentators therefore put forth a second-best metric, the Area Deprivation Index (ADI), which tracks socioeconomic disadvantage, plausibly a large part of what makes race a predictor of severe COVID-19 outcomes.²⁰

In the next section, we argue that concerns about legal intervention are overly pessimistic if one accepts the COVID-19 risk rationale that we recommend. As we show, analogies with affirmative action cases are largely beside the point. Still, our argument reveals a possible legal obstacle. To address that obstacle, we also propose a system with the following features: (1) In line with our COVID-19 risk rationale, it adopts a neutral aim, (2) it pursues this aim through race-neutral means, but (3) it ends up being race-responsive in effect. We call this system “RE-NAM” (responsive in effect, neutral in aims and means).

We believe that this system has the potential to outperform the SVI, CCVI, or ADI in its sensitivity to the differential impact of the pandemic on certain racial minorities. It would be legally more robust than these other proposals, because it is compatible with a color-blind conception of discrimination. While our discussion is focused on the United States, there are important lessons for other countries. The moral view underpinning the legal debate is the color-blind (or symmetric) conception of discrimination and this view is popular (to varying degrees) in many societies.

4 | PRECEDENTS ABOUT AFFIRMATIVE ACTION ARE LESS RELEVANT THAN INITIALLY THOUGHT

There is no direct precedent by the U.S. Supreme Court on allocating medical resources partly by race. The precedents cited to motivate the concern about the SVI (and by implication CCVI)²¹ largely pertain to affirmative action in education contexts. Conservative Justices have been increasingly skeptical of all racial classifications in line with the color-blind conception. In *Parents Involved*, Justice Roberts wrote, “The way to stop discrimination on the basis of race is to stop discriminating on the basis of race.”²² Justice Thomas, who espouses a color-blind conception of discrimination, wrote in his dissent in *Grutter v. Bollinger* that “every time the government places citizens on racial registers and makes race relevant to the provision of burdens or benefits, it demeans us all.”²³ In the recent *Students for Fair Admissions* case, Justice Roberts reiterated that “[e]liminating racial discrimination means eliminating all of it.”²⁴ But that does not make it clear what the Court would say in the present case.

The first point to notice is that what draws most criticism is classifying *individuals* by race.²⁵ In the use of the SVI/CCVI for vaccine allocation, and potentially in other race-responsive systems, race would serve as a factor only in determining *geographic* priorities. People in the same priority phase and geographic area would be treated equally without regard to their race. Such policies may therefore be considered not to use any racial classification at all. Individual vaccine recipients would not be favored or disfavored directly because of their race; rather, they would be favored if they live in areas with many minorities.²⁶

Additionally, even the color-blind conception of discrimination is in principle compatible with racial criteria for allocating COVID-19 vaccines. Our proposed COVID-19 risk rationale helps explain why. The color-blind view on discrimination holds that it is difficult to justify any racial classification, regardless of its intent or impact, not that it is impossible. In law, this is reflected by race's status as a “suspect classification” that requires strict scrutiny, a standard according to which the policy has to be narrowly tailored in order to achieve a compelling state interest. In our own case, the state is pursuing a compelling state interest in minimizing COVID-19 morbidity and mortality impartially understood. Under our recommended rationale, responsiveness to race/ethnicity simply targets vaccines to those at highest risk of death and to thereby minimize COVID-19 mortality and morbidity.

Consider an analogy with sex discrimination. Breast cancer affects predominantly women and only exceptionally men. Prevention programs are, for this reason, focused on women rather than men. No plausible understanding of what constitutes sex discrimination, even one that is suspicious of all sex classifications, would regard this focus as discriminatory. A similar point holds for COVID-19 to the extent that race is a risk factor for medical need and public health urgency.

By contrast, in the case of university admissions, the primary rationale for racial preference is reducing background inequality or promoting diversity. There is an alternative paradigmatic basis for entitlement to the benefit in question, namely, academic merit, that critics protest has been disrupted. Even some defenders concede that affirmative action is in that respect a necessary compromise that, even when justified on balance, veers away from the ordinary norms in the relevant “distributive sphere.”²⁷ In the case of COVID-19 vaccines, the paradigmatic basis for entitlement is “medical and public health need, for example, the risk of severe COVID-19 and viral transmission. Race as a risk factor is part of this paradigmatic basis and does not disrupt any allocation.

²⁰Schmidt, H., et al., op. cit. note 19. A parallel debate focuses on triage decisions for intensive care in which proposals have been raised to incorporate ADI metrics into triage decisions with the aim of improving social and racial justice. See White, D. B., & Lo, B. (2021). Mitigating inequities and saving lives with ICU triage during the COVID-19 pandemic. *American Journal of Respiratory and Critical Care Medicine*, 203, 287–295.

²¹Schmidt, H., et al., op. cit. note 19; Persad, G. (2021). Allocating medicine fairly in an unfair pandemic. *University of Illinois Law Review*, 2021, 1085–1134.

²²*Parents Involved in Community Schools v Seattle School District No. 1*, 551 U.S. 701 (2007).

²³*Grutter v Bollinger* 539 U.S. 306, 371 (2003).

²⁴*Students for Fair Admissions v. Harvard*, 600 U.S. ____ 2023.

²⁵See also Persad, op. cit. note 21.

²⁶In that respect, proposals for individual racial preferences in the allocation of ventilators or extracorporeal membrane oxygenation (ECMO) between individual patients are clearer candidates to be ruled out as unconstitutional. See Sederstrom, N. (2020). The ‘give back’: Is there room for it? *Bioethics.net*. Retrieved September 16, 2021, from <https://www.bioethics.net/2020/07/the-give-back-is-there-room-for-it/>; Schmidt, H. (2020, April 15). The way we ration ventilators is biased. *New York Times*.

²⁷More broadly on this, see Gutmann, A. (1995). Justice across the spheres. In D. Miller & M. Walzer (Eds.), *Pluralism, justice, and equality* (pp. 99–119). Oxford University Press.

The color-blind conception of discrimination is thus friendlier to an allocation of COVID-19 vaccines by race-responsive means than it is to preferential university admissions. But one concern remains. Suppose that race is a risk factor for COVID-19 but that we could use nonracial factors that would be equally good or even better in predicting COVID-19 risks. For example, imagine that racial/ethnic disparities in COVID-19 risk are largely or fully driven by economic disparities. A proponent of the color-blind view of discrimination could protest that we should then not use race when other nonracial means are available to us. In the language of U.S. law, the criteria would not be “narrowly tailored” to achieve the public health outcomes. In this respect, we can see the SVI and the way in which it predicts COVID-19 mortality, morbidity, and incidence²⁸ as a baseline. We do not know how well the SVI performs unless we compare it to alternatives that do not make explicit reference to race. In the next section, we propose a system that plays one of two roles. Either it is better suited than the SVI or CCVI to avoid this complaint or, if it performs worse than the SVI, then it at least helps make the case that the explicit use of race in the definition of the SVI was indeed necessary in order to be narrowly tailored to reducing COVID-19 mortality and incidence. A priori, we cannot be sure that our system will in fact outperform the SVI, but we suggest a few reasons to be optimistic that it would if implemented.

5 | A RACE-RESPONSIVE BUT RACE-NEUTRAL ALLOCATION SYSTEM

Our proposed RE-NAM system incorporates the following three elements:

1. A color-blind public health goal (e.g., reducing and equalizing individual risk in line with the COVID-19 risk rationale) is adopted.
2. Epidemiological modeling, assisted by a carefully designed Artificial Intelligence (AI), is tasked to identify which factors best predict individual risk.
3. Vaccine allocation practice then follows the identified risk factors to best serve the public health goal (with corrections only when data points are missing).

Our system is race-responsive in its effect, while it is race-neutral in its aims and means. Thus, we will refer to it as the RE-NAM system. The input data for the model consist of the best available public health data of the relevant country or jurisdiction. These will differ between countries. For example, the U.K.'s NHS has access to different data than the CDC in the United States or equivalent authorities in other countries. A first step is to gather data on COVID-19 deaths and infections by demographic and other factors. These data need to be checked for potential biases. For example,

concerns about asymmetries in testing can be countered by seropositivity surveys. The input data will mirror some of the variables incorporated in the SVI and other disadvantage indices. It can also include additional data points, if available. But even if it contained similar categories as the SVI, the modeling process could reveal the appropriate weightage and interaction effects between different forms of deprivation that are pre-set by the SVI.

A similar approach is used in the United Kingdom. QCovid is a risk prediction algorithm for personal risk of COVID-19 hospital admission and mortality.²⁹ The National Health Service has used QCovid to identify individuals at high risk, who were then added to a Shielded Patients List and offered priority access to vaccines. More extensive uses are imaginable. How such an algorithm could be best used will depend on local circumstances and on the stage of vaccine rollout that the country is in. In countries like the United Kingdom, in which virtually all of society is registered with a general practitioner (a primary care physician), the algorithm could be used to generate a list for each general practitioner of who among their patients to prioritize. Alternatively, the algorithm could help identify criteria for vaccine access priority groups. Importantly, the algorithm could also be used less directly, to identify strong risk factors and allocate vaccines geographically, prioritizing, for example, postcodes identified as at high risk.

The factors used by the model would presumably include the most pertinent socio-economic factors that constitute a large part of what makes race relevant in the first place. For example, one model of COVID-19 transmission has identified the inability to reduce mobility as a key factor that explains most of the racially disproportionate outcomes.³⁰ But many other factors are possible, too, and any sophisticated model would need to heed a variety of risk factors.

Such a RE-NAM system, reliant on epidemiological modeling, should in some respects make racial priorities significantly more acceptable. The classifications that the RE-NAM system ultimately uses to track the identified risk factors raise fewer legal difficulties than express racial classifications put forth by human allocators. The proposed system's classifications are not suspect because they are not traditionally associated with social disadvantage or prejudice. Therefore, they legally require only a lower standard of justification. The state would need to show that the classifications used (e.g., inability to reduce mobility) are rationally related to a legitimate policy goal. They clearly are and while the expected output of the allocation scheme would probably prioritize racial minorities, the scheme would use only race-neutral means as its input.

²⁸Karmakar, M., Lantz, P. M., & Tipirneni, R. (2021). Association of social and demographic factors with COVID-19 incidence and death rates in the US. *JAMA Network Open*, 4, e2036462.

²⁹Cliff, A. K., Coupland, C. A. C., Keogh, R. H., Diaz-Ordaz, K., Williamson, E., Harrison, E. M., Hayward, A., Hemingway, H., Horby, P., Mehta, N., Bengler, J., Khunti, K., Spiegelhalter, D., Sheikh, A., Valabhji, J., Lyons, R. A., Robson, J., Semple, M. G., Kee, F., ... Hippisley-Cox, J. (2020). Living risk prediction algorithm (QCovid) for risk of hospital admissions and mortality from coronavirus 19 in adults: National derivation and validation cohort study. *BMJ*, 371, m3731.

³⁰Chang, S., Pierson, E., Koh, P. W., Gerardin, J., Redbird, B., Grusky, D., & Leskovec, J. (2020). Mobility network models of COVID-19 explain inequities and inform reopening. *Nature*, 589, 82–87.

But what if the system does identify race itself as a factor predicting individual risk? There are two reasons why this is less problematic under this RE-NAM modeling system. Pragmatically, if, in service of the public health goal, the AI ends up heeding racial affiliation directly, that would temporarily remain in its “black box” and opaque. Such a use of race would not require a public communication of racial preferences, one that could be described in real time by demagogues as reverse racism. Nor would it raise conservative concerns about racial divisiveness, “Balkanization,”³¹ or increased racial tensions.

It is true that sometime later, such a black box could be analyzed and the racial element could be exposed. But here is where a second point kicks in. If race is found by an impartial AI as irremediably necessary to track color-blind public health goals, then this is presumably because race is a very important factor for public health. This could either be a relatively objective and direct effect (e.g., racism creates stress that weakens immune systems) or because feasible data points do not cover (e.g., highly specific socio-economic) factors that are objectively crucial and correlated with race or ethnicity. As such, the RE-NAM system would be patently impartial and closely related to the paradigmatic basis for entitlement to vaccines, given feasibility constraints. This rationale, that modeling has shown that racial minorities are predominantly at risk for a factor irreducible to anything other than race, is likely to survive strict scrutiny. The model will have shown that other traits do not predict COVID-19 risk as accurately. This undermines the complaint that race is a coarse-grained category from a public health perspective and that nonracial criteria would be more accurate.³²

Regardless of whether race is used as an input factor, the output of RE-NAM is likely to involve a priority for racial minorities, simply because the public health data indicate that racial minorities are at higher risk. This should not create conflicts with a Court that demands color-blind *means* to generate output. One might wonder whether adopting an alternative policy that yields the same effective outcome that vaccines go disproportionately to racial minorities amounts to indirect discrimination (what is called “disparate impact” in the United States). However, the U.S. Constitution, unlike that of many other countries, does not prohibit indirect discrimination. Laws that have disparate impacts on different racial groups are constitutional under U.S. law so long as they were adopted without a discriminatory purpose.³³ Our proposed responsive but neutral allocation scheme is, as advertised, race-neutral in its means³⁴ and it follows a neutral and nondiscriminatory purpose. The RE-NAM

allocation scheme is mindful of racial disparities, but this does not make its purpose (to equalize and reduce individual health risk across the entire population, impartially to all) discriminatory. To the contrary, every person is weighted equally, and all relevant risk factors are considered. The differential priorities for racial groups are the result of the pandemic, not the result of a failure to give due respect to the majority group.

6 | ADDRESSING CONCERNS ABOUT POTENTIAL ALGORITHMIC BIAS

There is growing awareness that AI can exacerbate social injustice and in particular racial injustice.³⁵ For example, a recent study of a widely used commercial algorithm to predict health risks found that it underestimated the health of Black patients compared to its estimates for White patients.³⁶ The study's main message is that the choice of variables used for AI can drive disproportionate outcomes. Might the proposed RE-NAM system likewise exacerbate racial injustice? We believe that this can be avoided. This point goes back to our discussion of the relevant input data for the model. First, the same recent study also points out that the disparate impacts vanish once the choice of variables is correct. We see no special barrier to careful selection of the relevant variables. Such selection and choices are therefore important steps to ensure that the proposed modeling exercises promote rather than hinder racial justice.

Another concern about predictive algorithms, namely, that they mirror past behavior and thus past injustices and biases, may actually speak in favor of our RE-NAM model.³⁷ In algorithms used in criminal justice, AI models rely on past data from an already unfair and biased system. Arrest records and trial outcomes incorporate racial biases present in policing, prosecution, and the justice system. This damning characteristic is less of a concern for RE-NAM. Although putatively objective measures like COVID-19 health outcomes can be biased,³⁸ for example, when discrimination or other disparities in access make African Americans test less, public health experts have methods that partly correct for such biases. One such method is to make use of seropositivity surveys as a check of previous infections.³⁹ The

³¹Siegel, R. B. (2011). From colorblindness to antibalkanization: An emerging ground of decision in race equality cases. *The Yale Law Journal*, 120, 1278–1366.

³²Indeed, one key concern in the majority opinion in *Students for Fair Admissions* is that universities struggled to provide high quality data to show their policy fulfilled its aims. Our data-driven proposal fares better in this regard. However, if one is even more concerned about legal or political repercussions, one could additionally instruct the AI never to heed racial affiliation directly.

³³*Washington v Davis* 426 U.S. 229 (1976).

³⁴Either it never uses race as a criterion, or it is as race-neutral as possible by only using race when nonracial categories fail to explain individual risk levels. In the latter case, the use of race would be narrowly tailored.

³⁵See, for example, Dressel, J., & Farid, H. (2018). The accuracy, fairness, and limits of predicting recidivism. *Science Advances*, 4, eaao5580; Mayson, S. (2019). Bias in, bias out. *Yale Law Journal*, 128, 2218–2300; Obermeyer, Z., Powers, B., Vogeli, C., & Mullainathan, S. (2019). Dissecting racial bias in an algorithm used to manage the health of populations. *Science*, 366, 447–453; Wiens, J., Price, W. N. II, & Sjoding, M. W. (2020). Diagnosing bias in data-driven algorithms for healthcare. *Nature Medicine*, 26, 25–36; Hellman, D. (2020). Measuring algorithmic fairness. *Virginia Law Review*, 106, 811–866; Hedden, B. (2021). On statistical criteria of algorithmic fairness. *Philosophy & Public Affairs*, 49, 209–231.

³⁶Obermeyer, Z., et al., op. cit. note 34.

³⁷Mayson, op. cit. note 34; Johnson, G. M. (2021). Algorithmic bias: On the implicit biases of social technology. *Synthese*, 198, 9941–9961.

³⁸Wiens, J., et al., op. cit. note 34.

³⁹See, for example, Kalish, H., Klumpp-Thomas, C., Hunsberger, S., Baus, H. A., Fay, M. P., Siripong, N., Wang, J., Hicks, J., Mehalko, J., Travers, J., Drew, M., Pauly, K., Spathies, J., Ngo, T., Adusei, K. M., Karkanitsa, M., Croker, J. A., Li, Y., Graubard, B. I., ... Sadtler, K. (2021). Undiagnosed SARS-CoV-2 seropositivity during the first 6 months of the COVID-19 pandemic in the United States. *Science Translational Medicine*, 13, eab3826.

algorithm that we propose can thus operate with health data that are less subjective than might be feared; there is no clear analog to using biased and prejudiced judgments of individuals as input information.

Furthermore, there is another asymmetry. A higher likelihood to receive legal sanctions compounds the initial injustice.⁴⁰ In contrast, our proposal uses information about background injustices to alleviate the sequelae of that initial injustice. That background information is not a biasing “bug” but the central “feature.” The whole point is to remedy a potential resulting health injustice through prioritized access to health care.

7 | REASONS WHY RE-NAM MAY TURN OUT TO BE BETTER TARGETED

Without having implemented a system like RE-NAM, we cannot draw definitive conclusions on how it would compare to alternate allocation schemes like the SVI. The proof should be in the pudding of testing in the field. But some advantages of RE-NAM make it appear promising.

First, RE-NAM is more tailored to the specific public health needs of the COVID-19 pandemic (or, in future uses, the relevant future pandemic) than alternative race-responsive policies. The SVI and the ADI were not created during COVID-19. Naturally, they are less precise for COVID-19 risks than specially designed measures. Although disadvantaged racial minorities were hit harder, COVID-19 had its surprising effects. Early warnings that sub-Saharan Africa would be the worst-hit world region turned out to be false. Concerns about COVID-19 spreading like wildfire in slums in South Asia did not pan out as expected. General disadvantage is often a good proxy—especially when we know little—but attention to outbreak specifics remains valuable. The CCVI, which was created specifically for COVID-19, is a step in the right direction in this respect but it nevertheless inherits some of the SVI's crudeness. Both the SVI and the CCVI give higher preference to geographical areas with a higher share of all non-White racial groups regardless of these groups' respective COVID-19 risks. As mentioned earlier, the evidence on the disproportionate impact on Asian Americans is less straightforward than the evidence for African Americans, Latinx, or Native Americans. The evidence from the United Kingdom shows a large variability among British Asians. Data in the United States are typically collected at coarse-grained census categories. The ADI, in turn, is not sensitive to whether the residents of the relevant area are able to maintain social isolation at work and presumably to many other predictors of COVID-19.

In a similar vein, we think that the British QCovid can be further improved. QCovid uses a single variable to account for social deprivation (the Townsend deprivation score applied to postcodes). Variables about individual circumstances include only medical

information on comorbidities.⁴¹ Adding variables like employment status, the ability to work from home, and so on, which are plausible contenders for capturing COVID-19 risk factors, or even variables that surprisingly are identified as relevant, would presumably increase the accuracy of the model. The choice of the variables has to be, of course, operationally feasible.

A second advantage of RE-NAM is that AI modeling can be continuously updated in a way that the construction of a health index usually cannot. This is helpful not only in order to tailor the model to different allocators but also to reflect the growing and sometimes dynamic⁴² knowledge gained about the virus and infection patterns, and the shifting objective patterns.

Third, epidemiological modeling can be put to multiple uses and thereby assist in prioritization efforts that take into consideration both the risk of severe outcomes for the individual and the risk of infecting others. With models in hand for both outcomes, public health authorities will need to make an assessment about trade-offs between them. Our discussion has focused on the component of the modeling that deals with individual risk. It is a strength of RE-NAM that it can be easily extended to modeling transmission in a way that these health indices cannot readily do.

Fourth, RE-NAM prioritizes those members of racial minorities who bear the greatest individual COVID-19 risk. It is not clear in advance that a priority by race or a priority by living in areas with a high non-White population is the best way to understand and target racial risk factors. Imagine that the model traces, for example, a characteristic narrower than simply a certain racial affiliation (say, the real subset of a racial group in whom severe outcomes for that group concentrate, and that is driving the relatively severe outcomes for the entire racial group of which it is only a part). Again, the case of the United Kingdom is instructive: COVID-19 risk for British Asians is mainly driven by highly elevated risk for British of Bangladeshi and Pakistani ancestry. Our allocation scheme would then prioritize that subset only, and not the larger racial/ethnic group as such. This subset might be a more fine-grained ethnic affiliation, or a combination of racial affiliation and geographic location, or of racial affiliation and age, or any other more complicated combination. Whatever it may be, AI can help us identify which subgroups are driving the disproportionate outcomes without having the predetermined idea that indices measuring health vulnerability necessarily have.

Some supporters of the background injustice rationale would then retort that all African Americans, Latinx people, and Native Americans should get priority, even ones identified to be at low risk of COVID-19 infection and severe disease. Yet, we think of selectivity as an advantage, not a defect, of our proposal. In general, we disagree that a proper response to group inequality is to provide further benefits to already advantaged (in comparison to anyone) members of generally disadvantaged groups.⁴³ Priority to the worse-

⁴¹Clift, A. K., et al., op. cit. note 28.

⁴²Neelon, B., Mutiso, F., Mueller, N. T., Pearce, J. L., & Benjamin-Neelon, S. E. (2021). Spatial and temporal trends in social vulnerability and COVID-19 incidence and death rates in the United States. *PLoS ONE*, 16, e0248702.

⁴⁰See more broadly on the idea of compounding injustice Hellman, D. (2018). Indirect discrimination and the duty to avoid compounding injustice. In H. Collins & T. Khaitan (Eds.), *Foundations of indirect discrimination law* (pp. 105-121). Hart Publishing.

off, one of the justifications invoked for the background injustice rationale, should pertain to individual deprivation, if feasible. Our proposal achieves an individualized assessment of COVID-19 risk that relies on group affiliation only to the extent that that the affiliation is a useful proxy.

8 | CONCLUSION

The disproportionate impact of COVID-19 on racial minorities must be reflected in COVID-19-related allocation decisions. However, we can improve on explicit racial preferences. Epidemiological modeling can help us understand the underlying factors that drive these disproportionate impacts, helping allocation decisions to track these underlying factors nimbly and accurately. When they track race itself, that would be either because race has independent causal contribution to adverse COVID-19 outcomes or because there are causal contributors closely associated with race that no other operationalizable indicator captures.

An epidemiologically informed RE-NAM allocation AI system would be legally defensible and nondiscriminatory, even under a conservative, color-blind conception of nondiscrimination. It signals a way to make resource allocation appropriately race-responsive in effect even under political and legal constraints demanding race-neutral means and aims. This will remain relevant beyond COVID-19 for decision-makers seeking to address racial injustice within the limits of color-blind (symmetric) discrimination laws in the United States and beyond.

ACKNOWLEDGEMENTS

Nir Eyal's work was partly funded by Longview Philanthropy.

CONFLICT OF INTEREST STATEMENT

The authors declare no conflict of interest.

ORCID

Bastian Steuwer  <https://orcid.org/0000-0001-6532-3845>

Nir Eyal  <http://orcid.org/0000-0003-1056-6609>

AUTHOR BIOGRAPHIES

Bastian Steuwer is a political philosopher and bioethicist based at Ashoka University, India. He holds graduate degrees in philosophy and law and has received postgraduate training in population-level bioethics. His current research focuses on discrimination, distributive justice, and the fair allocation of health care.

Nir Eyal is the inaugural Henry Rutgers professor of bioethics at Rutgers University. He founded and directs Rutgers's Center for Population-Level Bioethics, with appointments at the School of Public Health and the Department of Philosophy. His work, which covers many areas of population-level bioethics, appeared in *Science*, *PNAS*, *NEJM*, *Lancet*, *BMJ*, and the leading bioethics venues.

How to cite this article: Steuwer, B., & Eyal, N. (2023). Ethical and legal race-responsive vaccine allocation. *Bioethics*, 37, 814–821. <https://doi.org/10.1111/bioe.13203>

⁴³See Eyal, N. (2018). Inequality in political philosophy and in epidemiology: A remarriage. *Journal of Applied Philosophy*, 35, 149–167; Lippert-Rasmussen, K. (2013). When group measures of health should matter. In N. Eyal, S. A. Hurst, O. F. Norheim, & D. Wikler (Eds.), *Inequalities in health* (pp. 52–65). Oxford University Press.