



Using measures of race to make clinical predictions: Decision making, patient health, and fairness

Charles F. Manski^{a,b,1} , John Mullahy^c, and Atheendar S. Venkataramani^d

Contributed by Charles F. Manski; received February 27, 2023; accepted May 24, 2023; reviewed by Marcella Alsan and Ziad Obermeyer

The use of race measures in clinical prediction models is contentious. We seek to inform the discourse by evaluating the inclusion of race in probabilistic predictions of illness that support clinical decision making. Adopting a static utilitarian framework to formalize social welfare, we show that patients of all races benefit when clinical decisions are jointly guided by patient race and other observable covariates. Similar conclusions emerge when the model is extended to a two-period setting where prevention activities target systemic drivers of disease. We also discuss non-utilitarian concepts that have been proposed to guide allocation of health care resources.

clinical prediction | patient care | utilitarian welfare analysis | race

Recent years have seen considerable debate around two related questions that are sometimes conflated:

Should the manner in which a clinician treats a patient depend on that patient's race?

Should measures of race be included as covariates in clinical prediction models/algorithms?

Numerous arguments for and against the use of race in clinical settings have been advanced. For a concise overview of this debate see the recent contributions by refs. 1 and 2, who summarize the main opposing points of view.

Clear articulation of the goals we wish to achieve with patient care can help clarify the underlying disagreement and perhaps even resolve some of it. Two goals appear prominent. The first concerns the clinician's role in improving the health of the individual patient in front of them. The second concerns a societal objective of eliminating racial disparities in health care and health. Reasonable people may differ in how much weight they place on each goal. This matters because efforts to advance achievement of one goal may not advance, indeed may hinder, achievement of the other. Regardless of one's priorities, conceptual clarity is required to achieve either or both goals.

This study aims to lay a clear and rigorous foundation to inform ongoing debates about the use of measures of race in clinical settings. We develop a simple model of the effects of considering race on clinical prediction and decisions and then generalize it. We begin with a canonical model that uses a single-period utilitarian framework wherein a clinician's sole objective in any clinical encounter is to make the treatment recommendation that maximizes the patient's expected health, conditional on all the information available at the time the treatment recommendation is made. The main result is that failure to use all available information in clinical prediction models or in a particular clinical encounter results in sub-optimal expected health for patients. In particular, statisticians' failure to use observed measures of race in developing clinical prediction algorithms or clinicians' failure to use available measures of race in treatment decisions will generally result in sub-optimal expected health outcomes for patients of all races. This result holds regardless of the extent to which available measures of race correlate with ancestry, socioeconomic status, or other drivers of health. Algorithms that use better measures of underlying drivers (e.g., direct measures of genotype, social deprivation, or biomarkers that capture these processes) may obviate the use of race in specific clinical settings (3–5). However, until such alternate measures are available, algorithms including race would (weakly) outperform those that do not, by virtue of capturing a range of important correlates of health, no matter how imperfectly.

We then extend the canonical model by embedding the clinician's decision in a two-period framework wherein the period-two clinical decision is made in a context where social and environmental factors differentially affect patients in the first period. This model takes into

Significance

The use of race measures in clinical prediction models is contentious. We seek to inform the discourse by evaluating the inclusion of race in probabilistic predictions of illness that support clinical decision making. We show that patients of all races benefit when clinical decisions are jointly guided by patient race and other observable covariates. Similar conclusions emerge when the model is extended to a two-period setting where prevention activities target systemic drivers of disease.

Author affiliations: ^aDepartment of Economics, Northwestern University, Evanston, IL 60208; ^bInstitute for Policy Research, Northwestern University, Evanston, IL 60208; ^cDepartment of Population Health Sciences, University of Wisconsin–Madison, Madison, WI 53726; and ^dDepartment of Medical Ethics & Health Policy, University of Pennsylvania, Philadelphia, PA 19104

Author contributions: C.F.M., J.M., and A.S.V. designed research; performed research; analyzed data; and wrote the paper.

Reviewers: M.A., Harvard University; and Z.O., University of California, Berkeley.

The authors declare no competing interest.

Copyright © 2023 the Author(s). Published by PNAS. This article is distributed under [Creative Commons Attribution-NonCommercial-NoDerivatives License 4.0 \(CC BY-NC-ND\)](https://creativecommons.org/licenses/by-nc-nd/4.0/).

¹To whom correspondence may be addressed. Email: cfmanski@northwestern.edu.

Published August 22, 2023.

account theoretical models of structural or systemic racism, wherein socially produced disadvantages over the patient's life course adversely affect their health and well-being (6–9). As a result, the circumstances of patients seeking care in period two will differ, with some patients being more advantaged than others. Specifically, when period two arises, there will be disparities across the patient population in economic opportunities, healthcare access, and other determinants of health. Despite these disparities, our analysis shows that the clinician's role in period two as a clinician should still be to provide optimal care to each patient in the same manner as recommended by our canonical model. Increasingly, clinicians are advocating and striving for reductions in health disparities. But this is a separate matter than the activities they pursue and decisions they make in a clinical encounter. In our two-period utilitarian model, these activities are best pursued in period one, when prevention activities addressing social and structural drivers of health can reduce disadvantages in period two.

We recognize that the utilitarian framing of our analysis may not appeal to practitioners and policymakers who prioritize non-utilitarian notions of justice and fairness. For example, they may prioritize ensuring groups receiving similar (access to) health care resources as a first order goal. To situate these perspectives, we formally discuss alternate notions of fairness and disparity-aversion proposed to guide societal allocation of resources.

We also include an Appendix that discusses two additional, related literatures. First, we note the rapidly expanding literature on algorithmic fairness in economics, computer science, and related fields. With some exceptions, this recent literature poses criteria for fairness and seeks to empirically measure adherence to them, without embedding them in a problem of welfare maximization as we do here. While our formal analysis does not link explicitly to this literature, we outline some of its noteworthy features (Section A.1 of the Appendix). Second, we consider the practice of race-norming, and illustrate that it is conceptually distinct from the main issue we consider here, which is the use of race measures to inform clinical predictions (Section A.2 of the Appendix).

We do not consider in this paper several additional issues, which are beyond the scope of the present paper but may be promising avenues for future work. First, we do not explicitly model the determination of factors that may influence patients' demand for or health systems' delivery of care, such as cost and other access barriers, cost-effectiveness, and trust. Second, we do not consider multi-stage clinical encounters ("episodes of care"), where clinical prediction models in one stage may sequentially inform patient-specific decisions that in turn ultimately determine health outcomes. Third, we do not explore anti-discrimination regulations that affect risk assessment in other domains, including credit rating, insurance, and justice system outcomes (10, 11). Such regulations do not currently exclude the use of race measures in clinical prediction. although see ref. 12 and discussion in Section 5 of this paper. Fourth, for clarity purposes, we do not explicitly consider mismeasurement in the clinical outcome, an issue which is well-explored by refs. 13 and 14.

We proceed as follows. Section 1 provides background on the recent concern with and several examples of the use of race in medical decision making. Section 2 presents the one-period model of utilitarian treatment choice. Section 3 extends the model to a two-period setting in which preventive care precedes treatment choice. Section 4 discusses non-utilitarian views on fairness and justice. Section 5 concludes.

1. Background

1.1. Decision Making Contexts. Much of clinical decision-making involves the quantitative prediction of disease risk, treatment effectiveness, or other outcomes based on various

sources of data. The canonical empirical prediction model, sometimes called a clinical algorithm, is an estimate $\hat{P}(y|x, z)$ of a probability $P(y|x, z)$ where y is the health outcome of interest and x and z denote vectors of covariates on which the prediction is based (15). For the present discussion, we take x to be covariates that will certainly be included in the model (e.g., patient age) and z to be variables that may be included at the analyst's discretion.*

The consideration that occupies our attention in this paper is the selection of the z variables. We do not endow $P(y|x, z)$ with any causal interpretation. The task at hand is to choose z to generate conditionally optimal predictions. It will be demonstrated formally in Sections 2 and 3 that richer specifications of z generally yield superior clinical predictions than do sparser ones. So long as a candidate covariate z_k has some predictive power, its inclusion in the z vector will result in superior predictions.

1.2. Inclusion of Race in Prediction Models. There is presently considerable debate around the inclusion in prediction models of a particular z_k : patient race. In the clinical literature, (16) and (17) represent notable examples of calls to remove the consideration of race in prediction models, while (3) summarizes a range of arguments. In economics, (1) and (2) summarize the key arguments, with 1) supporting and 2) criticizing the inclusion of measures of race in clinical prediction models.

Manski (1) describes and then questions four assertions that have been advanced as arguments against the inclusion of race in prediction models. These assertions are: i) race is a social, not biological, concept; ii) race should not be considered if there is no established causal link between race and the illness; iii) using race may perpetuate or worsen racial health inequities; and iv) many persons are offended by the use of race in risk assessment. With the stated goal of making clinical decisions that would be expected to yield the best outcome for each patient, (1) concludes with this observation:

If an alternative perspective is to have a compelling foundation, it should explain why society should find it acceptable to make risk assessments using other patient characteristics that clinicians observe, but not race. It should explain why the social benefit of omitting race from risk assessment is sufficiently large that it exceeds the harm to the quality of patient care.

In a somewhat similar vein, (3) states:

There is no time more important than now to understand how race and social and biological factors interact to affect health. Estimation of essential physiologic processes, such as kidney function, with variables that do not incorporate race and are more accurate than race is a worthy aspiration. Those estimating tools should have equal or greater precision, be soundly grounded in evidence on health outcomes, and be acceptable to patients.

A common concern voiced by those arguing against the inclusion of race in prediction is that its very definition is complex and elusive. Many writers note that measured race historically has reflected a social rather than biological concept. When they argue that this precludes the use of race in making clinical predictions, they may have in mind that there is not a one-to-one mapping between a specified racial categorization and elements of ancestry

*Section A.2 of the Appendix discusses the relationship between such clinical prediction models and the so-called race-norming of outcomes.

or epigenetic modifications—e.g., a particular set of genotypes or stress-led differences in patterns of gene expression—that predict disease risk (16). Another argument against the inclusion of race stems from a belief that, no matter how race is defined, including race as a predictor of outcomes will result in inferior health care for racial and ethnic minority populations relative to the care received by others. Health care that is sensitive to a patient's race, it is asserted, risks exacerbating systemic racial biases that are argued to prevail (18). Briggs (2) captures this sentiment concisely:

...while it is difficult to refute the central contention that optimal decision making requires the use of all covariates that are associated with outcome, the assumption that racial covariates, and their application within the medical arena, are sufficiently free from bias (structural, institutional or personal) misses the point of the underlying argument: that race is not the same as every other covariate in our arsenal. It is a covariate that is acting as a proxy for a wide range of other explanatory variables that could be genetic/biological but in many circumstances are more likely to be sociological/socioeconomic.

As we consider the controversy regarding inclusion of race in clinical prediction and decision-making, we think it important to recognize that the two sides of the debate do not simply differ in advocating different strategies for attaining the same goal. Rather, the core differences also concern the attainment of different goals, though these goals are often not explicitly or clearly articulated. Advancing the debate will thus require clarity around the desired goals and the measurements necessary to assess whether or not these goals are being achieved. Improved patient health is a straightforward goal to define and measure. It is less obvious how to conceptualize clearly and measure other objectives that have often been stated, which include the achievement of equity and elimination of disparities in health care (19). It is understandable that different parties may have different goals, reflecting different values. This is all the more reason that a clear articulation of such goals is a necessary antecedent of productive discussion and policymaking.

One notable example of the lack of clarity in the literature concerns the common use of the term “bias.” This term can mean different things to different people in different contexts. Bias in the sense typically encountered in discussions of race has little or nothing to do with how the expectation of a statistical estimator compares to its true value. Bias has a relatively clear meaning in the study of “algorithmic bias”, an area of inquiry kindred to but distinct from what we pursue in this paper (13).[†] Beyond this literature, we find that bias is often weakly conceptualized. Careful scrutiny of statements about the presence of racial bias is essential for advancing understanding.

1.3. Examples. The issues raised above are of more than just academic methodological interest. Across an increasingly broad spectrum of clinical decision-making contexts, debates about whether to include race measures—however defined—in clinical decision making are shaping clinical practice. Examples include the treatment and management of kidney disease (20), the assessment of osteoporosis and fracture risk associated with osteoporosis (21), the use of spirometry to assess lung function (22), and the

determination of appropriate X-ray radiation doses (23).[‡] Vyas et al. (16) and Cerdefia et al. (17) review clinical contexts wherein such considerations have arisen. Decisions to include or exclude race in clinical prediction are already affecting the health care being delivered to patients of all races.

In some fields, leading institutions have formally recommended race-free risk assessment. A notable case is (20), which presents the recommendations of the National Kidney Foundation-American Society of Nephrology Task Force on Reassessing the Inclusion of Race in Diagnosing Kidney Disease. The Task Force considered the prevailing use of race in computation of estimated glomerular filtration rate (eGFR), a measure of kidney function. It recommended removal of race as a determinant of eGFR, writing (pp. 5–6):

For U.S. adults (>85% of whom have normal kidney function), we recommend immediate implementation of the Chronic Kidney Disease Epidemiology Collaboration (CKD-EPI) creatinine equation refit without the race variable in all laboratories in the U.S. because it does not include race in the calculation and reporting, includes diversity in its development, is immediately available to all labs in the U.S., and has acceptable performance characteristics and potential consequences that do not disproportionately affect any one group of individuals.

Research related to this recommendation is documented by (24), the underlying concern being that the use of race predictions may increase eGFR for Black patients, potentially reducing the likelihood of receiving therapy for chronic kidney disease or being listed for transplantation. The recommendation has already been implemented in multiple major medical centers. Powe (3) urged caution, noting that calls to de-adopt race-inclusive algorithms did not consider “all of the ramifications and long-term health consequences”, some of which may introduce harm. For example, the move from using race-free measures of serum creatinine to race-adjusted eGFR measures helped reduce racial disparities in receipt of metformin, a diabetes medication with well-established short- and long-term benefits, by increasing use among patients who otherwise would have been contraindicated from receiving it under race-free measures (25). Some writers have argued that removal of race-adjustment may also increase rates of rejection of Black kidney donor candidates and lead to reductions in doses of chemotherapies (26). A recent analysis found a greater degree of misclassification in many race-free eGFR algorithms relative to algorithms including race, with the exception being a newly developed race-free algorithm that relies on an alternate biomarker, cystatin, instead of the more widely collected biomarker creatinine (4).

Another prominent example is the use of race in models that predict fracture risks, with a particular focus on osteoporosis. For example, the FRAX algorithm (Fracture Risk Assessment Tool; <https://www.sheffield.ac.uk/FRAX/index.aspx>) for the U.S. population has four different versions to assess fracture risk for Black, Caucasian, Hispanic, and Asian patients. The use of these differential algorithms has been criticized for reasons that include the possibility that, *ceteris paribus*, lower fracture risk predictions for Black than White, Asian, and Hispanic patients may result in the former being undertreated to prevent osteoporosis or slow its progression if diagnosed (16). In response to these criticisms, (21) note that there are indeed racial differences in osteoporosis treatment even after adjusting

[†]Briefly, a prominent version of algorithmic bias arises when proxy outcome measures are used instead of true measures in prediction models and when the gap between the proxy and the true measures depends on some variable of interest, e.g. race. See Section A.1 of the Appendix for further discussion.

[‡]In some clinical contexts attention has been paid to whether diagnostic accuracy using traditional diagnostic standards may depend not *per se* on race but rather on skin tone. Examples include pulse oximetry and assessment of pressure injuries or ulcers. Many of the issues raised herein with respect to race inclusion would appear equally applicable to skin tone, with a key distinction being that there are accepted objective measures of skin tone (59).

for fracture risk. They argue that racial differences in treatment gaps—the fraction of a population indicated for but not receiving treatment—are best understood and acted upon when risks are predicted accurately: “The quantification afforded by FRAX has allowed inequalities in the treatment gap to be identified.” While (21) urges that FRAX not be used uncritically, they conclude that its use ultimately “helps to resolve, rather than exacerbate, racial inequalities.”

1.4. The State of the Debate. Across the many clinical contexts where these issues have arisen, perspectives like those offered by (1, 3, 21) are heard far less than calls to proceed with the removal of race from clinical decision-making (2, 16). This dynamic holds even as arguments have been advanced to explicitly consider race in other (more population-level) settings. For example, there have been calls to allocate COVID-19 vaccines on the basis of existing and predicted disparities in disease risk (27) and for greater inclusivity in clinical trials of subjects from sub-populations that have been typically underrepresented (e.g., racial minorities) in such studies (28).⁵ These tensions are borne out in statements such as those by ref. 30:

Research should continue to ascertain and use race, not as an explanatory variable, but rather in a manner that highlights where racial health disparities exist, that reduces enrollment bias, and that maximizes the generalizability of the results.

Recognizing that concerns about the use of race measures in clinical prediction often arises from how race is considered in other domains, such as the criminal justice system or housing markets, (30) urge against overgeneralization in their commentary:

There is clearly a lack of consensus in how to cope with race as a predictor in CPMs [clinical predictive models]... we focused on the narrow but important issue of fairness (i.e., the harms caused to one group vs. other groups), which would militate against the use of race predictively in other contexts but not necessarily in medical decision making. We point out that some of the reluctance about using race predictively in CPMs comes from overgeneralizing from these other contexts, which may substantially misrepresent the concerns in the medical decision making context.

Overall, the lack of clarity and agreement in when and how race should be considered illustrates the need for a common, rigorous foundation to fix ideas and clarify specific concerns and ultimate goals. We demonstrate in the next sections that the application of one conceptually rigorous—and, we would argue, familiar and broadly applicable—foundation leads to the conclusion that the health and welfare of patients of all races is likely to suffer if the use of race in prediction models is proscribed. We offer this framework to help both sides of the debate clearly articulate goals and to identify incorrect, misguided, or unacceptable assumptions. We thereafter describe alternative frameworks whose properties, assumptions, and implications can be similarly discussed and debated.

⁵U.S. Food and Drug Administration (28) states: “Differences in response to medical products (e.g., pharmacokinetics, efficacy, or safety) have been observed in racially and ethnically distinct subgroups of the U.S. population. These differences may be attributable to intrinsic factors (e.g., genetics, metabolism, elimination), extrinsic factors (e.g., diet, environmental exposure, sociocultural issues), or interactions between these factors. Analyzing data on race and ethnicity may assist in identifying population-specific signals.”

2. Utilitarian Patient Care

This section presents a standard medical-economics framing of utilitarian patient care as a single-period problem of treatment choice with predetermined risks of illness. Section 3 extends the framework to encompass preventive efforts that reduce risks of illness.

2.1. Optimal Care with Predetermined Illness Risk. Medical economists have commonly studied clinical decision making in a static setting of individual patient care. This setting supposes that a clinician must choose how to treat each person in a population of patients. The clinician observes certain predetermined covariates for each patient, with associated predetermined risks of illness. The objective is to maximize a utilitarian welfare function, one that sums up the benefits and costs of treatment across the population of patients. A utilitarian welfare function formalizes the idea of “patient-centered care.” The assumption of individualistic care means that the care received by one patient may affect that person but does not affect other members of the population. This assumption is generally realistic when considering non-infectious diseases.

A common problem in clinical decision making is that treatments must be chosen with incomplete knowledge of their potential outcomes. A central idealization of the setting usually studied by medical economists has been to assume that the clinician knows the probability distribution of health and welfare outcomes that may potentially occur if a patient with specified observed attributes is given a specified treatment—i.e., the clinician has rational expectations. The assumption of rational expectations does not assert that the clinician can predict patient outcomes with certainty. Instead, it means that the clinician makes accurate probabilistic predictions conditional on observed patient covariates.

Analysis shows that if a clinician has rational expectations, the problem of optimizing patient care has a simple solution: patients should be divided into groups having the same observed covariates and all patients in a such a group should be given the care that yields the highest within-group mean utility. Our analysis with a utilitarian social welfare function suggests that it is optimal to differentially treat patients with different observed covariates if different treatments maximize their within-group mean utility. In our model, patients with the same observed attributes should be treated uniformly.

Analysis also shows that achievable utilitarian welfare across the population weakly increases as more patient covariates are observed. Observing more covariates enables a clinician to refine the probabilistic predictions of treatment outcomes on which decisions are based. Refining these predictions is beneficial to the extent that doing so affects optimal treatment choices.

These findings have been documented extensively. Abstract analyses, not specific to medical applications, include (31–33). Analyses in the literature on medical economics include (34–36). We prove the findings here in the simple instructive setting of choice between two treatments. Section 2.2 presents well-known findings. Section 2.3 formulates an argument that strengthens the conclusion drawn about the value of covariate information to treatment choice.

2.2. Optimal Choice Between Two Treatments.

2.2.1. Treatments, covariates, and illness probabilities. Suppose that a clinician must choose between treatments A and B. The choice is made without knowing a patient’s illness outcome, $y = 1$ or 0. We assume throughout that y measures accurately the health state on which a patient’s welfare depends; that is, y is not a proxy of the sort considered by (13) in their investigation of algorithmic bias (see Appendix A.1). The clinician observes predetermined patient covariates (x, z). Having rational expectations, the clinician knows a patient’s probability of illness conditional on these covariates.

Thus, the clinician knows $p_x = P(y = 1 | x)$ and $p_{xz} = P(y = 1 | x, z)$. The covariates x may include predetermined variables that predict illness or that affect patient expected utility of treatment (e.g., age, gender, health history, socioeconomic status, trust in medical care). Consider patients who have the same value of x but who vary in their values of z . Assume that z takes values in a finite set Z . Each value of z occurs for a positive fraction of patients; thus, $P(z | x) > 0$ for all $z \in Z$. Assume that p_{xz} varies with z . To relate this setup to the concern with use of race in medical decision making, one may take z to be an observable measure of race. Specifically how race might be “observed”—perceived skin tone; self-designation; category recorded in an electronic medical record; etc.—will clearly be relevant in practice, but it is not essential for this paper’s arguments. Our abstract analysis does not require any specific interpretation of z . It holds when z is any observable covariate.

We show that maximum utilitarian welfare using p_{xz} to predict illness is always at least as large as using p_x . We note that this captures cases such as the prediction problem for kidney disease, where better predictors (of underlying drivers) of kidney function (eGFR) for all groups z have been identified, potentially obviating the need to include z as a covariate. We also note that maximum utilitarian welfare is strictly larger if optimal treatment choice varies with z . In the latter case, we characterize the determinants of the magnitude of the improvement.

2.2.2. Maximizing expected utility. Patient outcomes with each treatment depend on whether a patient has the disease. Patients are heterogeneous, so treatment response may differ across patients. Let $U_x(y, t)$ denote the expected utility that a patient with covariates x would experience with treatment t , should the illness outcome be y . We suppose for simplicity that this expected utility does not vary across patients with different values of z . However, illness probabilities may vary with z . A utilitarian clinician with rational expectations need not know the personal utility function of each individual patient, but is assumed to know expected utility $U_x(y, t)$ for each possible value of (x, t, y) .

A clinician making a treatment decision does not know a patient’s illness outcome but does know the illness probabilities

$$\begin{aligned} &\text{choose treatment A if } p_{xz} \cdot U_x(1, A) + (1 - p_{xz}) \cdot U_x(0, A) \\ &\geq p_x \cdot U_x(1, B) + (1 - p_x) \cdot U_x(0, B), \end{aligned} \quad [2a]$$

$$\begin{aligned} &\text{choose treatment B if } p_{xz} \cdot U_x(1, A) + (1 - p_{xz}) \cdot U_x(0, A) \\ &\leq p_x \cdot U_x(1, B) + (1 - p_x) \cdot U_x(0, B) \end{aligned} \quad [2b]$$

With criterion [1], the maximized value of expected utility for patients with covariates x is

$$\begin{aligned} &\max [p_x \cdot U_x(1, A) + (1 - p_x) \cdot U_x(0, A), \\ &p_x \cdot U_x(1, B) + (1 - p_x) \cdot U_x(0, B)]. \end{aligned} \quad [3]$$

With criterion [2], the maximized value of expected utility for patients with covariates (x, z) is

$$\begin{aligned} &\max [p_{xz} \cdot U_x(1, A) + (1 - p_{xz}) \cdot U_x(0, A), \\ &p_{xz} \cdot U_x(1, B) + (1 - p_{xz}) \cdot U_x(0, B)]. \end{aligned} \quad [4]$$

In this case, the maximized value of expected utility for patients with covariates x is the mean of Eq. 4 with respect to the distribution of z conditional on x ; that is,

$$\begin{aligned} &E_{z|x} \{ \max [p_{xz} \cdot U_x(1, A) + (1 - p_{xz}) \cdot U_x(0, A), \\ &p_{xz} \cdot U_x(1, B) + (1 - p_{xz}) \cdot U_x(0, B)] \}. \end{aligned} \quad [5]$$

2.2.3. Comparison of the criteria using Jensen’s inequality. Jensen’s inequality shows that the magnitude of Eq. 5 weakly exceeds that of Eq. 3, implying that criterion [2] performs at least as well as Eq. 1 from the utilitarian perspective. In particular,

$$\begin{aligned} &E_{z|x} \{ \max [p_{xz} \cdot U_x(1, A) + (1 - p_{xz}) \cdot U_x(0, A), p_{xz} \cdot U_x(1, B) + (1 - p_{xz}) \cdot U_x(0, B)] \} \\ &\geq \max \{ E_{z|x}(p_{xz}) \cdot U_x(1, A) + [1 - E_{z|x}(p_{xz})] \cdot U_x(0, A), E_{z|x}(p_{xz}) \cdot U_x(1, B) + [1 - E_{z|x}(p_{xz})] \cdot U_x(0, B) \} \\ &= \max [p_x \cdot U_x(1, A) + (1 - p_x) \cdot U_x(0, A), p_x \cdot U_x(1, B) + (1 - p_x) \cdot U_x(0, B)]. \end{aligned} \quad [6]$$

p_x and p_{xz} . With this knowledge, the clinician can compute expected utility in two ways, as $p_x \cdot U_x(1, t) + (1 - p_x) \cdot U_x(0, t)$ or as $p_{xz} \cdot U_x(1, t) + (1 - p_{xz}) \cdot U_x(0, t)$. Presuming that the objective is to maximize expected utility, the clinician may therefore use the criterion

$$\begin{aligned} &\text{choose treatment A if } p_x \cdot U_x(1, A) + (1 - p_x) \cdot U_x(0, A) \\ &\geq p_x \cdot U_x(1, B) + (1 - p_x) \cdot U_x(0, B), \end{aligned} \quad [1a]$$

$$\begin{aligned} &\text{choose treatment B if } p_x \cdot U_x(1, A) + (1 - p_x) \cdot U_x(0, A) \\ &\leq p_x \cdot U_x(1, B) + (1 - p_x) \cdot U_x(0, B), \end{aligned} \quad [1b]$$

or the criterion

The first inequality follows from Jensen’s inequality because $\max(\cdot, \cdot)$ is a convex function. The second equality holds because $E_{z|x}(p_{xz}) = p_x$. The first inequality in Eq. 6 is strict if there exist some values of z for which criterion [2] yields a different treatment than criterion (Eq. 1). The inequality is an equality if criteria [2] and [1] have the same solution for all values of z .

2.3. Direct Comparison of the Criteria. Jensen’s inequality provides a simple proof of the qualitative result that a utilitarian clinician should use all observed covariates to predict illness. However, it does not reveal quantitatively the extent to which criterion [2] outperforms criterion [1]. We can do this through direct comparison of the criteria.

Without loss of generality, let treatment A be optimal in Eq. 1. Let A be optimal in Eq. 2 for all $z \in Z_A$ and let Z_B be the complement of Z_A . Thus, inequality (Eq. 2a) holds for $z \in Z_A$, some non-empty proper subset of Z , and does not hold for $z \in Z_B$, also a non-empty proper subset of Z . Criterion [2] yields better outcomes than [1] for persons with $z \in Z_B$ and the same outcomes as [1] for persons with $z \in Z_A$.

Now use the decomposition of Z into (Z_A, Z_B) to rewrite [3] and [5] as

$$\max[p_x \cdot U_x(1, A) + (1 - p_x) \cdot U_x(0, A), p_x \cdot U_x(1, B) + (1 - p_x) \cdot U_x(0, B)] = p_x \cdot U_x(1, A) + (1 - p_x) \cdot U_x(0, A) \quad [7]$$

$$= P(z \in Z_A | x) \cdot E[p_{xz} \cdot U_x(1, A) + (1 - p_{xz}) \cdot U_x(0, A) | x, z \in Z_A] + P(z \in Z_B | x) \cdot E[p_{xz} \cdot U_x(1, A) + (1 - p_{xz}) \cdot U_x(0, A) | x, z \in Z_B],$$

and

$$E_{z|x} \{ \max[p_{xz} \cdot U_x(1, A) + (1 - p_{xz}) \cdot U_x(0, A), p_{xz} \cdot U_x(1, B) + (1 - p_{xz}) \cdot U_x(0, B)] \}$$

$$= P(z \in Z_A | x) \cdot E[p_{xz} \cdot U_x(1, A) + (1 - p_{xz}) \cdot U_x(0, A) | x, z \in Z_A] \quad [8]$$

$$+ P(z \in Z_B | x) \cdot E[p_{xz} \cdot U_x(1, B) + (1 - p_{xz}) \cdot U_x(0, B) | x, z \in Z_B].$$

Subtracting [7] from [8] yields

$$P(z \in Z_B | x) \cdot E\{ [p_{xz} \cdot U_x(1, B) + (1 - p_{xz}) \cdot U_x(0, B)] - [p_{xz} \cdot U_x(1, A) + (1 - p_{xz}) \cdot U_x(0, A)] | x, z \in Z_B \}. \quad [9]$$

The inequality $p_{xz} \cdot U_x(1, B) + (1 - p_{xz}) \cdot U_x(0, B) > p_{xz} \cdot U_x(1, A) + (1 - p_{xz}) \cdot U_x(0, A)$ holds for all $z \in Z_B$. Hence, Eq. 9 is positive. This qualitative finding repeats the earlier one using Jensen's inequality. Moreover, [9] quantifies the extent to which criterion [2] outperforms criterion [1]. The magnitude of Eq. 9 is the product of two factors. One is the fraction $P(z \in Z_B | x)$ of patients for whom treatment B yields strictly larger expected utility than treatment A. The other is the mean gain in expected utility that criterion [2] yields for the subset Z_B of patients.

2.4. Utilitarian Care and Disparities in Treatment and Health.

Utilitarian treatment choice aims to maximize patient well-being, optimizing care within groups of patients who share common observed covariates. In this sense, it embeds a specific, clear idea that clinical decision making should be fair and just, expressing the idea that the primary concern of a clinician is to do best by each individual patient. Nevertheless, it does not imply that patients with different observed covariates receive the same treatments or experience the same health. In this sense it yields treatment, although not necessarily health, disparities across groups of patients (Box 1).

Consider persons with covariates (x, z) . The optimal utilitarian treatment of these persons is A if [2a] holds and is B if [2b] holds. The value of the maximized expected utility of these persons is given by Eq. 4, which varies with the illness probability P_{xz} and expected utility function $U_x(\cdot, \cdot)$. Cross-covariate disparities in treatment and health are well-motivated from the utilitarian perspective if clinicians have rational expectations and act accordingly. Of course, if P_{xz} does not vary with z then there is no compelling reason to include z in a prediction model. For example, use of race measures does not additionally improve predictions of kidney function in newly developed algorithms that use the biomarker cystatin (4).

2.5. Potential Challenges to the Utilitarian Setup. Cross-covariate disparities in treatment and health may not be well-motivated if, not having rational expectations, clinicians act based on imperfect knowledge of illness probabilities and expected utilities. Then clinical decisions may be sub-optimal. Moreover, the degree of sub-optimality may vary across patients with different observed covariates. For example, research suggests that clinicians' assessments of physical pain or cardiac risk among Black patients may be noisier or more prone to bias than those among White

patients (38–40) or that take-up of preventive services among Black patients may increase markedly if their doctor is also Black (40).

It is also possible that barriers to health care access and medical mistrust may reduce the external validity of estimated clinical prediction models (due to incomplete data in medical records, for example), leading to inaccurate assessments whose signs and magnitudes may vary by patient race. Sub-optimality of clinical decision making is always undesirable from the utilitarian perspective, regardless of which specific groups of patients most suffer the consequences.

The utilitarian prescription to improve decision making is to improve knowledge of patient illness probabilities and expected utilities. One could do so by addressing biases in medical education (41), developing enhanced decision-support tools (42), or striving to improve completeness of electronic health records. The utilitarian perspective implies that these solutions would dominate solutions that discard information from clinical prediction and decision-making on race entirely.

This point is underscored by the fact that non-race covariates may themselves be measured with error, often in ways that vary by race and only incompletely capture the predictive nature of race-measures in clinical risk. A recent example demonstrates that race-corrected algorithms for predicting colorectal cancer risk dominate models that use a range of other covariates instead, ostensibly because of measurement error and race-specific missing data (43). An implication of this finding is that researchers or policymakers hoping to rely on measures of covariates that alone may "explain" racial or ethnic disparities in disease risk alone may be less successful in predicting these risks than using race measures as well.

Some have deemed disparities undesirable from non-utilitarian perspectives on fairness and justice, particularly when z measures race or ethnicity (see, for example, the discussion in ref. 45). We discuss these alternate perspectives in Section 4. Additionally, how race is measured in research studies may have implications for the use of such measures in clinical settings, where it may be measured differently. We discuss the implications of measurement issues in Section 5.

Box 1. Optimal Treatments and Health Outcome Disparities

Suppose there is credible evidence that the optimal treatment for Black patients (who are denoted J) is treatment B and optimal treatment for White patients (who are denoted K) is treatment A. Assume that treatments A and B are different dosages of the same drug or intensities of the same intervention, with treatment intensity denoted t . Suppose the respective health production functions are:

$$h_j(t_j) = 8 - (t_j - 4)^2 \text{ and } h_k(t_k) = 10 - (t_k - 2)^2.$$

Then A is $t_k = 2$ and B is $t_j = 4$. Thus, Black patients must receive higher-intensity treatment than White patients to attain optimal health.

Define health disparity as $D(t_j, t_k) = h_j(t_j) - h_k(t_k)$. When no patient receives treatment ($t_j = t_k = 0$), health is worse for Black patients than for White patients, with $D(0, 0) = -14$.

The treatment allocation that is optimal by the utilitarian criteria developed in Section 2 has Black patients receive B ($t_j = 4$) and White patients receive A ($t_k = 2$). The disparity is thus less negative than when no treatment is received, with $D(4, 2) = -2$, although White patients' health is still better than Black patients'.

From this baseline, consider the implications for health of four alternative treatment allocations, each of which treats Black and White patients with a common intensity:

Alternative 1. Everyone is treated with what evidence indicates is best for White patients. Compared with the baseline, the disparity is more negative, with $D(2, 2) = -6$.

Alternative 2. Everyone is treated with what evidence indicates is best for Black patients. The disparity is now positive, with $D(4, 4) = 2$.

Alternative 3. Everyone is treated with a population weighted average of the optimal treatments, $C = pA + (1 - p)B$, where p is Whites' population share). Assume $p = 0.8$. Then the common intensity is 2.4. Now disparity is more negative than at baseline, with $D(2.4, 2.4) = -4.4$.

Alternative 4. Treatments are allocated to eliminate disparity, which occurs if the common intensity is 3.5. Then $D(3.5, 3.5) = 0$. While disparity is eliminated, the health levels of both Black and White patients suffer relative to the baseline.

The implications of these treatment allocations for health levels and disparity are summarized in this table:

Scenario	t		Health		D(t _j , t _k)
	Black (J)	White (K)	Black (J)	White (K)	
No treatment	0	0	-8	6	-14
Baseline: Optimal	4	2	8	10	-2
Alternative 1	2	2	4	10	-6
Alternative 2	4	4	8	6	2
Alternative 3	2.4	2.4	5.44	9.84	-4.4
Alternative 4	3.5	3.5	7.75	7.75	0

3. Optimal Prevention and Treatment in a Two-Period Model

We now consider a two-period problem of utilitarian patient care. Period 2 remains as in Section 2, with a clinician choosing treatments for patients having predetermined risks of illness. The change is that we introduce a period 1, in which society may undertake preventive efforts that reduce illness risk in period 2. In our setup, Period 1 can either precede period 2 or occur simultaneously (i.e., a clinician can engage in a preventive action as well as start treatment). If preventive care were costless, it would be optimal to provide it to all patients. However, such care is costly. Hence, the decision to provide it in period 1 is non-trivial.

To formalize the preventive-care decision in a simple manner, let $s = 1$ if a patient receives such care in period 1 and $s = 0$ otherwise. Suppose that society can personalize such care, choosing s separately for patients with different observed covariates. Let $C_{xz} > 0$ be the social cost of providing preventive care to patients having covariates (x, z) , with cost expressed in units commensurate with patient utility. Let P_{sxz} denote the probability of illness in period 2, which now varies with s as well as with (x, z) . We assume that preventive care reduces the risk of illness; thus, $P_{1xz} < P_{0xz}$.

We have found that, in period 2, it is optimal to condition illness risk on (x, z) when choosing a treatment, yielding [4] as the maximized value of expected utility. From the perspective of period 1, two versions of Eq. 4 are feasible, one setting $s = 1$ and the other setting $s = 0$. The resulting feasible maximized values of expected utility respectively are:

$$\max[p_{1xz} \cdot U_x(1, A) + (1 - p_{1xz}) \cdot U_x(0, A), \quad [10a]$$

$$p_{1xz} \cdot U_x(1, B) + (1 - p_{1xz}) \cdot U_x(0, B)] - C_{xz},$$

$$\max[p_{0xz} \cdot U_x(1, A) + (1 - p_{0xz}) \cdot U_x(0, A), \quad [10b]$$

$$p_{0xz} \cdot U_x(1, B) + (1 - p_{0xz}) \cdot U_x(0, B)].$$

Providing preventive care is optimal if and only if [10a] exceeds [10b]. This inequality holds if preventive care reduces illness risk sufficiently and if such care is not too costly.

We present this model to connect our analysis to the emerging literature on systemic racism and its onward consequences on health and health disparities (6–9). In that literature, long-standing

discriminatory processes harm health over the life-course by increasing stress, reducing economic opportunities and financial security, and reducing access to health care. With this in mind, the term “preventive care” as used here can be interpreted either as direct preventive health care or more broadly to include any number of social activities or policies that reduce disease risk in period 2. Such policies may include broad economic or social interventions (e.g., minimum wage policy, social safety net policies, environmental policies, reparations) or interventions specifically targeted to address discrimination in period 1 (e.g., legal standards, affirmative action).

The key point is that while assessing the costs and benefits of such policies may be valuable, the clinician’s optimal period 2 treatment decision remains as above, given whatever hand period 1 has dealt. Engaging clinicians to work on ameliorating problematic systemic drivers of health is laudable. But our analysis shows that such work is generally separate from the care a clinician provides to a specific patient in period 2.

Consideration of a specific clinical setting can further elucidate these points. A clinician may be interested in prescribing insulin for a hospitalized patient with uncontrolled insulin-dependent diabetes, a condition for which failure to take insulin will result in death. The patient (for the purposes of this example, by virtue of race) faces myriad structural barriers, thereby leading the patient to hold less generous health insurance and face financial barriers to paying for medicines out-of-pocket. The clinician may want to prescribe a contemporary insulin regimen that most closely mimics physiology and thus now represents standard of care. However, despite the clinician’s best efforts, insurance coverage for the contemporary regimen is denied and the co-pays are unaffordable. Instead, coverage for an older and cheaper type of insulin regimen is provided. The clinician correctly discharges the patient with the older regimen, because this particular patient can afford it and thereby is more likely to adhere to it. The clinician may have no ability to affect patient access to insurance and capacity to afford co-pays. However, prior to the clinical encounter in period 2, the clinician may devote effort to address patients’ socio-economic needs and other environmental conditions in period 1, in the expectation that such efforts will reduce illness probabilities for patients in period 2.

4. Non-Utilitarian Concepts of Fairness

We showed in Section 2 that, if clinicians have rational expectations, utilitarian treatment choice optimizes care within groups of patients who share common observed covariates. This expresses a specific version of the idea that clinical decision making should be fair and just, without implying that patients with different observed covariates should receive the same treatments or will experience the same health. We observed that cross-covariate disparities may not be well-motivated from the utilitarian perspective if clinicians make sub-optimal decisions based on imperfect knowledge. The utilitarian prescription to cure sub-optimality is to improve clinical knowledge, thereby improving decisions for all groups of patients.

Writers arguing for removal of race from medical risk assessment appear to have something other than a utilitarian perspective in mind. Vyas et al. (16) call for a general reconsideration of the use of race in risk assessments, stating: “Many of these race-adjusted algorithms guide decisions in ways that may direct more attention or resources to white patients than to members of racial and ethnic minorities.” Cerdeña et al. (17) state: “race-based medicine... perpetuates health-care disparities.” Briggs (2) states: “race is not the same as every other covariate in our arsenal.” Unfortunately, such statements are typically too nebulous to enable readers to

either reconcile or contrast them with the principles of utilitarian decision making.

However, here we do take note of two bodies of work that have sought to generate well-defined non-utilitarian concepts and measures of fairness, with particular application to racial equity. We discuss ideas developed in the economics literature presently. We describe computer-science work on algorithmic fairness in Section A.1 of the Appendix.

4.1. Fairness in the Economics Literature. Economists have long sought to pose and analyze concepts of fairness. For example, (45) cogently posed a concept of “specific egalitarianism.” This concept moves away from utilitarianism, which concerns itself with the overall utility that a person experiences, instead supposing that society desires that (p. 264): “certain specific scarce commodities should be distributed less unequally than the ability to pay for them.” Tobin (45) discussed medical care as a leading case of such a specific commodity.

A simple framework illustrates key points, made originally by ref. 47 and extended by ref. 48. In the present health context, an allocation of resources in an N-member population is an allocation of treatments, \mathbf{T} , where, as above, each individual’s treatment t_n is either A or B. Consider a population of $N = 2$, comprising of one Black patient (denoted J) and one White patient (denoted K), so that $\mathbf{T} = (t_J, t_K)$, where \mathbf{T} is one of (A,A), (A,B), (B,A), or (B,B). Treatments result in scalar health outcomes, or expected outcomes if treatment response is uncertain, via individual-specific health production functions $h = h_n(t_n)$. We initially consider fairness of treatments rather than outcomes since, in general, the clinician—whose fairness is in question—is likely to have greater control over the former.

The treatment that results in the greatest health for J may or may not be the same as the treatment that is best for K; that is, $\text{sign}(h_J(A) - h_J(B)) \not\geq \text{sign}(h_K(A) - h_K(B))$. An algorithm that recommends uniform treatment for J and K when the signs are opposite implies that either J or K is not maximally healthy.

When there are no consumption externalities, \mathbf{T} is said to be envy-free if $t_J \succcurlyeq_J t_K$ and $t_K \succcurlyeq_K t_J$; that is J and K each (weakly) prefers the treatment they receive under \mathbf{T} to the treatment received by the other individual. It is natural to assume that the preference relationship corresponds directly to each individual’s own h , although this is reconsidered below. If \mathbf{T} is also Pareto efficient, then \mathbf{T} is said to be fair. An algorithm that recommends a fair allocation to a clinician might be considered to be a fair algorithm. (Note however that the term “algorithmic fairness” is often used in a different sense—see Section A.1 of the Appendix).

Suppose now that there are externalities in consumption (48) or interdependent preferences (49), wherein J and K care not just about the treatment they receive but also care about the treatment received by the other; that is, they have preferences over the allocations \mathbf{T} . Such externalities might take many different forms, but one that is perhaps particularly relevant to recent discussions of racial equity is where individuals are disparity-averse.

Disparities of concern might imaginably be with respect to treatment allocations themselves or to the health outcomes arising from different treatment allocations. (See Box 1 for an example that underscores the distinction between treatment and outcome disparities.) Suppose for instance that $h_J(A) > h_J(B)$. If J is disparity-averse with respect to outcomes, it may be that J’s preferences over the possible treatments \mathbf{T} nonetheless are:

$$(A, A) \succcurlyeq_J (B, B) \succcurlyeq_J (A, B) \succcurlyeq_J (B, A)$$

(strong disparity aversion, SDA),

or

$$(A, A) \succ_J (A, B) \succ_J (B, B) \succ_J (B, A)$$

(weak disparity aversion, WDA)

Now J 's preference ordering over treatment allocations depends on more than just J 's own health; for example, $(B, B) \succ_J (A, B)$ under SDA.

An alternative characterization of disparity aversion would take a societal perspective rather than the perspective of a given member of society. Without privileging any individual's treatment preferences, one might arrive at a societal ("S") partial ordering of the treatment allocations if social disparity aversion is with respect to treatments

$$\left\{ \begin{array}{l} (A, A) \\ (B, B) \end{array} \right\} \succ_S \left\{ \begin{array}{l} (A, B) \\ (B, A) \end{array} \right\}.$$

Corresponding health outcomes might be reasonable partial tie-breakers if both individuals are, say, better off under A than B:

$$(A, A) \succ_S (B, B) \succ_S \left\{ \begin{array}{l} (A, B) \\ (B, A) \end{array} \right\}.$$

While the Foley-Varian notions of envy and fairness are not obviously applicable in this externality context, recognition that real-world individual and/or social preferences may sometimes have elements of disparity aversion could help us understand why some may view "fair" allocations in clinical contexts to differ from ones that are, by the above definitions, envy-free and Pareto efficient. Indeed, some recent calls for the removal of race measures in clinical predictions hint strongly at authors' disparity aversion. For instance, (16) write: "If doctors and clinical educators rigorously analyze algorithms that include race correction, they can judge, with fresh eyes, whether the use of race or ethnicity is appropriate.... They can discern whether the correction is likely to relieve or exacerbate inequities. If the latter, then clinicians should examine whether the correction is warranted."

5. Conclusions

One characterization of patient-centered care and personalized medicine (50) holds that:

Not only are care plans customized, but medications are often customized as well. A patient's individual genetics, metabolism, biomarkers, immune system, and other "signatures" can now be harnessed in many disease states—especially cancer—to create personalized medications and therapies, as well as companion diagnostics that help clinicians better predict the best drug for each patient.

Marshaling data that provide informative measures of these signatures is an important empirical task. These signatures' measures may not be perfect. Nevertheless, as we show formally in Sections 2 and 3, they have the potential to improve clinical predictions and, therefore, patients' health and well-being.⁶ We emphasized in the introductory paragraphs that this paper does not explicitly model

⁶Kanis et al. (21) write: "...risk factors should be chosen according to established criteria irrespective of our understanding of their basis or their accuracy. A good example is consumption of alcohol, which is notorious for being inaccurately reported. In general, people who drink alcohol tend to neglect or underestimate their alcohol consumption.... It matters not whether the return is accurate—only that it provides a consistent indication of risk, which it does. Thus, we are more interested in association than causality. The same goes for race, location, and ethnicity."

the determination of factors that may influence patients' demand for or health systems' delivery of care, such as cost and other access barriers, cost-effectiveness, and trust. Currently these factors are often not well measured, and if mismeasured may be mismeasured in race-specific ways (43). As such, inclusion of race might actually better address such issues until more informative data are available.

A familiar mantra conceives healthcare quality as the delivery of the right care at the right time in the right place (51). Strongly implied in this statement is that "the right care" means "the particular care that's best for each patient." Consider a thought experiment wherein two otherwise-identical healthcare delivery systems (say R and S) differ only in how clinical prediction models are used to inform clinicians' treatment decisions. In R treatment decisions are guided by race-inclusive models $P(y|x,z)$, while in S treatment decisions are guided by race-exclusive models $P(y|x)$. In which system would fully informed patients choose to enroll? In which system would fully informed clinicians choose to practice? While we cannot pretend to understand patients' and clinicians' motives for the choices they would make, we submit that questions of this nature are worth contemplating.

If implemented in practice, the arguments we have offered here ultimately support better quality care being delivered to patients of all races. While the model we outline meets a well-known mantra of healthcare quality, in practice, the way forward when making decisions around whether and how to consider race in clinical-decision making will require a much clearer sense of policymaker and public goals than has been achieved through the public debate to date. The present contribution is a step towards putting these debates on firmer and more transparent footing.

Our results have immediate and potentially widespread policy implications. For example, the U.S. Dept. of Health and Human Services has recently undertaken efforts to revise Section 1557 of the Affordable Care Act ("Nondiscrimination"). Among other things, these efforts include proposing a new § 92.210 that focuses on discrimination and clinical algorithms:

Proposed § 92.210 states that a covered entity must not discriminate against any individual on the basis of race, color, national origin, sex, age, or disability through the use of clinical algorithms in its decision-making.... The intent of proposed § 92.210 is not to prohibit or hinder the use of clinical algorithms but rather to make clear that discrimination that occurs through their use is prohibited.... The Department notes that the use of algorithms that rely upon race and ethnicity-conscious variables may be appropriate and justified under certain circumstances, such as when used as a means to identify, evaluate, and address health disparities (12).

It is for lawyers and regulators to offer suitable definitions of "discriminate against" and "identify, evaluate, and address health disparities." Our analysis suggests that if the goal of the health care system is to deliver optimal care by the standards developed in Section 2, clinical appropriateness—what is best for the patient—must figure centrally in the deliberations. To this end clear distinctions must be drawn between treatments and outcomes when assessing disparities and discrimination.

We note in closing that there has been and remains considerable controversy regarding the definition of race. Race has been defined by ancestry, biology, social context, or a combination thereof. Race has also been defined based on self-identification or external perception. Moreover, these definitions may be fluid over time, even at the individual level. The challenges that definitional considerations like these pose for the preceding analysis are subtle but potentially important.

Specifically, suppose a clinical trial is used as the basis of the evidence on which a clinician's decisions will be made, and suppose further that in the trial's data each participant's race was coded in some manner (the particular manner is unimportant). The evidence at hand for the clinician will be a set of estimated prediction models $\hat{P}(y|x, z = z_t)$, one for each race category z_t and each x represented in the trial. The clinician now encounters a patient seeking treatment, observing their x characteristics and conceiving their race in some manner, say z_c . The question is whether the race category that serves as the basis of the clinician's treatment decision, z_c , is the same race category that would have been coded for this patient had they been a participant in the clinical trial. If so, the preceding analysis goes through without modification. If not, a more complex analysis must be pursued that is beyond this paper's scope.

Appendix

A.1. Algorithmic Bias and Algorithmic Fairness. In recent years considerable attention has been devoted to notions of algorithmic bias. Obermeyer et al. (52) have authored a plain language “playbook” that decision makers might consult to determine whether algorithms on which they rely are susceptible to bias. For these authors, biased algorithms are defined as follows:

More generally, in many important social sectors, algorithms guide decisions about who gets what. In these situations, we believe that if an algorithm scores two people the same, those two people should have the same basic needs—no matter the color of their skin, or other sensitive attributes... We consider algorithms that fail this test to be biased [page 1].

Algorithmic fairness is defined not merely by the absence of bias in an algorithm. Barocas et al. (53) outline challenges in defining algorithmic fairness:

We should not expect work on fairness in machine learning to deliver easy answers. And we should be suspicious of efforts that treat fairness as something that can be reduced to an algorithmic stamp of approval. At its best, this work will make it far more difficult to avoid the hard questions when it comes to debating and defining fairness, not easier. It may even force us to confront the meaningfulness and enforceability of existing approaches to discrimination in law and policy, expanding the tools at our disposal to reason about fairness and seek out justice [page 34].

These authors underscore such definitional challenges by offering nineteen distinct algorithmic fairness criteria (table 6 on page 75).

To investigate how algorithms may be unfair, (54) suggest a decomposition of the difference in an algorithm's estimated predictions between two groups, $\hat{E}[\tilde{Y}|G_1] - \hat{E}[\tilde{Y}|G_2]$, where \tilde{Y} is the measured outcome predicted by the algorithm. They write:

A raw difference in predictions across groups may arise for three reasons: differences in base rates, differences in

measurement error, or differences in estimation error across groups. Intervening at the level of the algorithm may address only the last two components of the disparity by investing in better training data and collecting a better proxy for the outcome of interest. In contrast, the base rate difference is a product of the underlying socioeconomic context itself, not the algorithm [page 92].

While the growing literature on algorithmic fairness is challenging to summarize succinctly, we mention here a few recent contributions. Paulus and Kent (55) provide a valuable comprehensive overview that stresses the importance of distinguishing algorithmic bias from algorithmic fairness. Barocas et al. (53) have written an e-book that explores a large set of criteria for algorithmic fairness. Another is (56) who approach algorithmic fairness (which they call fairness in AI systems) via welfare optimization approaches that involve definition of relevant social welfare functions and corresponding considerations of computational issues involved in their optimization. Liang et al. (57) focus on algorithmic fairness and algorithmic predictive accuracy, exploring in particular when there will and will not tend to be tradeoffs (e.g., when greater fairness can be attained only at the expense of reduced predictive accuracy).

A.2. Race-Normed Outcomes. Distinct from this paper's focus is the issue of race norming of outcomes. While widely used, the term race norming does not to our knowledge have a commonly accepted technical definition. It does, however, claim a Wikipedia page (<https://en.wikipedia.org/wiki/Race-norming>) that offers this definition: “Race-norming, more formally called within-group score conversion and score adjustment strategy, is the practice of adjusting test scores to account for the race or ethnicity of the test-taker.” This definition suggests transformation of an observed health outcome y that is measured in a uniform manner across races into a race-specific adjusted outcome, say $w(y, z)$ for a specified function $w(\cdot, \cdot)$. Then the conditional probability distribution of interest becomes $P[w(y, z)|x, z]$ rather than the $P(y|x, z)$ that has been our concern.

Race norming of clinical outcomes has been controversial. A prominent recent example is the National Football League's financial settlement with former players for concussion-related brain injuries (58). The original settlement that denied compensation to some Black players was based on race-normed cognitive test outcomes that indicated Black players to have different baseline cognitive capabilities than non-Blacks. Original settlements have been reconsidered given the questionable credibility of these baselines.

An important yet distinct issue in the NFL case is whether cognitive test scores are appropriate outcome measures on which to base compensation. It may be that what these tests measure is at best a proxy for true health status with the gaps between proxy and truth being race dependent. The use of such problematic proxies is one of the main concerns in the algorithmic bias work discussed earlier (13).

Data, Materials, and Software Availability. All study data are included in the main text.

ACKNOWLEDGMENTS. The authors are grateful to many colleagues for valuable comments and suggestions.

1. C. Manski, Patient-centered appraisal of race-free clinical risk assessment. *Health Econ.* **31**, 2109–2114 (2022).
2. A. Briggs, Healing the past, reimagining the present, investing in the future: What should be the role of race as a proxy covariate in health economics informed health care policy? *Health Econ.* **31**, 2115–2119 (2022).
3. N. Powe, Black kidney function matters—Use or misuse of race? *J. Am. Med. Assoc.* **324**, 737–738 (2020).

4. C. Hsu et al., Race, genetic ancestry, and estimating kidney function in CKD. *N. Engl. J. Med.* **385**, 1750–1760 (2021).
5. L. Inker et al., New creatinine- and cystatin C-based equations to estimate GFR without race. *N. Engl. J. Med.* **385**, 1737–1749 (2021).
6. Z. Bailey et al., Structural racism and health inequities in the USA: Evidence and interventions. *The Lancet* **389**, 1453–1463 (2017).

7. R. O'Brien, T. Neman, N. Seltzer, L. Evans, A. Venkataramani, Structural racism, economic opportunity, and racial health disparities: Evidence from U.S. counties. *SSM Popul. Health* **11**, e100564 (2020).
8. J. Bohren, P. Hull, A. Imas, Systemic discrimination: Theory and measurement (National Bureau of Economic Research, Working paper no. 29820, 2022).
9. W. Darity, Positions and possessions: Stratification economics and intergroup inequality. *J. Econ. Lit.* **60**, 400–426 (2022).
10. D. Pope, J. Sydnor, Implementing anti-discrimination policies in statistical profiling models. *Am. Econ. J. Econ. Policy* **3**, 206–231 (2011).
11. D. Arnold, W. Dobbie, P. Hull, Measuring racial discrimination in bail decisions. *Am. Econ. Rev.* **112**, 2992–3038 (2022).
12. U.S. Dept. of Health and Human Services, Nondiscrimination in health programs and activities. *Federal Register*, August 4, 2022, 47824–47920, Docket ID: HHS-OS-2022-0012.
13. Z. Obermeyer *et al.*, Dissecting racial bias in an algorithm used to manage the health of populations. *Science* **366**, 447–453 (2019).
14. S. Mullainathan, Z. Obermeyer, On the inequality of predicting A while hoping for B. *Am. Econ. Rev. Papers Proc.* **111**, 37–42 (2021).
15. C. Manski, *Patient Care under Uncertainty* (Princeton University Press, Princeton, 2019).
16. D. Vyas, L. Eisenstein, D. Jones, Hidden in plain sight—Reconsidering the use of race correction in clinical algorithms. *N. Engl. J. Med.* **383**, 874–882 (2020).
17. J. Cerdeña, M. Plaisime, J. Tsai, From race-based to race-conscious medicine: How anti-racist uprisings call us to act. *Lancet* **396**, 1125–1128 (2020).
18. D. Jones, Moving beyond race-based medicine. *Ann. Int. Med.* **174**, 1745–1746 (2021).
19. A. Sen, Why health equity? *Health Econ.* **11**, 659–666 (2002).
20. C. Delgado *et al.*, A unifying approach for GFR estimation: Recommendations of the NKF-ASN task force on reassessing the inclusion of race in diagnosing kidney disease. *J. Am. Soc. Nephrol.* **32**, 2994–3015 (2021).
21. J. Kanis *et al.*, On behalf of the International Osteoporosis Foundation, FRAX and ethnicity. *Osteoporosis Int.* **31**, 2063–2067 (2020).
22. S. Bonner, E. Wakeam, The end of race correction in spirometry for pulmonary function testing and surgical implications. *Ann. Surg.* **276**, e3–e5 (2022), 10.1097/SLA.0000000000005431.
23. L. Bavli, D. Jones, Race correction and the X-ray machine—The controversy over increased radiation doses for black Americans in 1968. *N. Engl. J. Med.* **387**, 947–952 (2022).
24. W. Williams, J. Hogan, J. Ingelfinger, Time to eliminate health care disparities in the estimation of kidney function. *N. Engl. J. Med.* **385**, 1804–1806 (2021).
25. J. Shin *et al.*, The FDA metformin label change and racial and sex disparities in metformin prescription among patients with CKD. *J. Am. Soc. Nephrol.* **31**, 1847–1858 (2020).
26. A. Levey, S. Tigan, N. Powe, J. Coresh, L. Inker, Kidney disease, race, and GFR estimation. *Clin. J. Am. Soc. Nephrol.* **15**, 1203–1212 (2020).
27. H. Schmidt, L. Gostin, M. Williams, Is it lawful and ethical to prioritize racial minorities for COVID-19 vaccines? *J. Am. Med. Assoc.* **324**, 2023–2024 (2020).
28. U. S. Food and Drug Administration, Enhancing the diversity of clinical trial populations—Eligibility criteria, enrollment practices, and trial designs: Guidance for industry, Center for Drug Evaluation and Research (CDER), Center for Biologics Evaluation and Research (CBER) (U.S. Food and Drug Administration, Silver Spring, MD, 2020), November 2020.
29. N. Bhakta *et al.*, Addressing race in pulmonary function testing by aligning intent and evidence with practice and perception. *Chest* **161**, 288–297 (2022).
30. J. K. Paulus, D. M. Kent, Race and ethnicity: A part of the equation for personalized clinical decision making? *Circ. Cardiovasc. Qual. Outcomes* **10**, e003823 (2017), 10.1161/CIRCOUTCOMES.117.003823.
31. I. Good, On the principle of total evidence. *Br. J. Philos. Sci.* **17**, 319–321 (1967).
32. C. Manski, *Identification for Prediction and Decision* (Harvard University Press, Cambridge, MA, 2007).
33. J. Kadane, M. Shervish, T. Seidenfeld, Is ignorance bliss? *J. Philos.* **105**, 5–36 (2008).
34. C. Phelps, A. Mushlin, Focusing technology assessment using medical decision theory. *Med. Decision Making* **8**, 279–289 (1988).
35. A. Basu, D. Meltzer, Value of information on preference heterogeneity and individualized care. *Med. Decision Making* **27**, 112–27 (2007).
36. C. Manski, Diagnostic testing and treatment under ambiguity: Using decision analysis to inform clinical practice. *Proc. Natl. Acad. Sci. U.S.A.* **110**, 2064–2069 (2013).
37. K. Schulman *et al.*, The effect of race and sex on physicians recommendations for cardiac catheterization. *N. Engl. J. Med.* **340**, 618–626 (1999).
38. K. Hoffman, S. Trawalter, J. Axt, M. Oliver, Racial bias in pain assessment and treatment recommendations, and false beliefs about biological differences between blacks and whites. *Proc. Natl. Acad. Sci. U.S.A.* **113**, 4296–4301 (2016).
39. M. Sun, T. Oliwa, M. Peek, E. Tung, Negative patient descriptors: Documenting racial bias in the electronic health record. *Health Affairs* **41**, 203–211 (2022).
40. M. Alsan, O. Garrick, G. Graziani, Does diversity matter for health? Experimental evidence from Oakland. *Am. Econ. Rev.* **109**, 4071–4111 (2019).
41. D. Burgess, M. van Ryn, J. Dovidio, S. Saha, Reducing racial bias among health care providers: Lessons from social-cognitive psychology. *J. Gen. Intern. Med.* **22**, 882–887 (2007).
42. E. Pierson, D. Cutler, J. Leskovic, S. Mullainathan, Z. Obermeyer, An algorithmic approach to reducing unexplained pain disparities in underserved populations. *Nat. Med.* **27**, 136–140 (2021).
43. A. Zink, Z. Obermeyer, E. Pierson, Race corrections in clinical models: Examining family history and cancer risk. medRxiv [Preprint] (2023). <https://doi.org/10.1101/2023.03.31.23287926> (Accessed 30 April 2023).
44. M. Kasy, R. Abebe, "Fairness, equality, and power in algorithmic decision-making" in *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, (2021), pp. 576–586.
45. J. Tobin, On limiting the domain of inequality. *J. Law Econ.* **13**, 263–278 (1970).
46. D. Foley, Resource allocation and the public sector. *Yale Econ. Essays* **7**, 45–98 (1967).
47. H. Varian, Equity, envy, and efficiency. *J. Econ. Theory* **9**, 63–91 (1974).
48. W. Thomson, "Fair allocation rules" in *Handbook of Social Choice and Welfare*, K. Arrow, A. Sen, and K. Suzumura, Eds. (Elsevier, Amsterdam, 2011), vol. II, chap. 21.
49. R. Pollak, Interdependent preferences. *Am. Econ. Rev.* **8**, 309–320 (1976).
50. NEJM Catalyst, What is patient-centered care? (2017). <https://catalyst.nejm.org/doi/full/10.1056/CAT.17.0559>. (1 January 2017).
51. W. Nowak *et al.*, Right care, right time, right place, every time. *Healthcare Financ. Manage.* **66**, 82–88 (2012).
52. Z. Obermeyer *et al.*, *Algorithmic Bias Playbook* (University of Chicago, Chicago Booth Center for Applied Artificial Intelligence, 2021).
53. S. Barocas, M. Hardt, A. Narayanan, Fairness and machine learning—Limitations and opportunities. <https://fairmlbook.org/> (Accessed 10 June 2022).
54. A. Rambachan *et al.*, An economic perspective on algorithmic fairness. *Am. Econ. Rev. Papers Proc.* **110**, 91–95 (2020).
55. J. Paulus, D. Kent, Predictably unequal: Understanding and addressing concerns that algorithmic clinical prediction may increase health disparities. *npj Digital Med.* **3**, 99 (2020).
56. V. Chen, J. Hooker, *Welfare-Based Fairness through Optimization* (Carnegie-Mellon University, 2021).
57. A. Liang, J. Lu, X. Mu, *Algorithmic Design: Fairness versus Accuracy* (Northwestern University, 2022).
58. W. Hobson, How 'race-norming' was built into the NFL concussion settlement. *Washington Post* (2021). [washingtonpost.com/sports/2021/08/02/race-norming-nfl-concussion-settlement/](https://www.washingtonpost.com/sports/2021/08/02/race-norming-nfl-concussion-settlement/) (Accessed 11 August 2023).
59. H. McCreath *et al.*, Use of munsell color charts to measure skin tone objectively in nursing home residents at risk for pressure ulcer development. *J. Adv. Nursing* **72**, 2077–2085 (2016).