

Assessing Disparate Impacts of Personalized Interventions: Identifiability and Bounds

Nathan Kallus
Cornell University
kallus@cornell.edu

Angela Zhou
Cornell University
az434@cornell.edu

Abstract

Personalized interventions in social services, education, and healthcare leverage individual-level causal effect predictions in order to give the best treatment to each individual or to prioritize program interventions for the individuals most likely to benefit. While the sensitivity of these domains compels us to evaluate the fairness of such policies, we show that actually auditing their disparate impacts per standard observational metrics, such as true positive rates, is impossible since ground truths are unknown. Whether our data is experimental or observational, an individual's actual outcome under an intervention different than that received can never be known, only predicted based on features. We prove how we can nonetheless point-identify these quantities under the additional assumption of monotone treatment response, which may be reasonable in many applications. We further provide a sensitivity analysis for this assumption by means of sharp partial-identification bounds under violations of monotonicity of varying strengths. We show how to use our results to audit personalized interventions using partially-identified ROC and xROC curves and demonstrate this in a case study of a French job training dataset.

1 Introduction

The expanding use of predictive algorithms in the public sector for risk assessment has sparked recent concern and study of fairness considerations [3, 9, 10]. One critique of the use of predictive risk assessment argues that the discussion should be reframed to instead focus on the role of *positive interventions* in distributing beneficial resources, such as directing pre-trial services to prevent recidivism, rather than in meting out pre-trial detention based on a risk prediction [8]; or using risk assessment in child welfare services to provide families with additional childcare resources rather than to inform the allocation of harmful suspicion [24, 51]. However, due to limited resources, interventions are necessarily targeted. Recent research specifically investigates the use of models that predict an intervention's benefit in order to efficiently target their allocation, such as in

developing triage tools to target homeless youth [36, 47]. Both ethics and law compel such personalized interventions to be fair and to avoid disparities in how they impact different groups defined by certain protected attributes, such as race, age, or gender.

The delivery of interventions to better target those individuals deemed most likely to respond well, even if a prediction or policy allocation rule does not have access to the protected attribute, might still result in disparate impact (with regards to social welfare) for the same reasons that these disparities occur in machine learning classification models [18]. (See Appendix C for an expanded discussion on our use of the term “disparate impact.”) However, in the problem of personalized interventions, the “fundamental problem of causal inference,” that outcomes are not observed for interventions not administered, poses a fundamental challenge for evaluating the fairness of any intervention allocation rule, as the true “labels” of intervention efficacy of any individual are never observed in the dataset. Metrics commonly assessed in the study of fairness in machine learning, such as group true positive and false positive rates, are therefore conditional on potential outcomes which are not observed in the data and therefore cannot be computed as in standard classification problems.

The problem of personalized policy learning has surfaced in econometrics and computer science [35, 41], gaining renewed attention alongside recent advances in causal inference and machine learning [4, 23, 52]. In particular, [15] analyze optimal treatment allocations for malaria bednets with nonparametric plug-in estimates of conditional average treatment effects, accounting for budget restrictions; [22] use the generalized random forests method of [53] to evaluate heterogeneity of causal effects in a program matching at-risk youth in Chicago with summer jobs on outcomes and crime; and [36] use BART [27] to analyze heterogeneity of treatment effect for allocation of homeless youth to different interventions, remarking that studying fairness considerations for algorithmically-guided interventions is necessary.

In this paper, we address the challenges of assessing the disparate impact of such personalized intervention rules in the face of unknown ground truth labels. We show that we can actually obtain point identification of common observational fairness metrics under the assumption of *monotone treatment response*. We motivate this assumption and discuss why it might be natural in settings where interventions only either help or do nothing. Recognizing nonetheless that this assumption is not actually testable, we show how to conduct sensitivity analyses for fairness metrics. In particular, we show how to obtain sharp partial identification bounds on the metrics of interest as we vary the strength of violation of the assumption. We then show to use these tools to visualize disparities using partially identified ROC and xROC curves. We illustrate all of this in a case study of personalized job training based on a dataset from a French field experiment.

2 Problem Setup

We suppose we have data on individuals (X, A, T, Y) consisting of:

- Prognostic features $X \in \mathcal{X}$, upon which interventions are personalized;
- Sensitive attribute $A \in \mathcal{A}$, against which disparate impact will be measured;
- Binary treatment indicator $T \in \{0, 1\}$, indicating intervention exposure; and
- Binary response outcome $Y \in \{0, 1\}$, indicating the benefit to the individual.

Our convention is to identify $T = 1$ with an active intervention, such as job training or a homeless prevention program, and $T = 0$ with lack thereof. Similarly, we assume that a positive outcome, $Y = 1$, is associated with a beneficial event for the individual, e.g., successful employment or non-recidivation. Using the Neyman-Rubin potential outcome framework [29], we let $Y(0), Y(1) \in \{0, 1\}$ denote the potential outcomes of each treatment. We let the observed outcome be the potential outcome of the assigned treatment, $Y = Y(T)$, encapsulating non-interference and consistency assumptions, also known as SUTVA [49]. Importantly, for any one individual, we never *simultaneously* observe $Y(0)$ and $Y(1)$. This is sometimes termed the fundamental problem of causal inference. We assume our data either came from a randomized controlled trial (the most common case) or an unconfounded observational study so that the treatment assignment is ignorable, that is, $Y(1), Y(0) \perp\!\!\!\perp T \mid X, A$.

When both treatment and potential outcomes are binary, we can exhaustively enumerate the four possible realizations of potential outcomes as $(Y(0), Y(1)) \in \{0, 1\}^2$. We call units with $(Y(0), Y(1)) = (0, 1)$ responders, $(Y(0), Y(1)) = (1, 0)$ anti-responders, and $Y(0) = Y(1)$ non-responders. Such a decomposition is also common in instrumental variable analysis [2] where the binary outcome is take-up of treatment with the analogous nomenclature of compliers, never-takers, always-takers, and defiers. In the context of talking about an actual outcome, following [42], we replace this nomenclature with the notion of response rather than compliance. We remind the reader that due to the fundamental problem of causal inference, response type is *unobserved*.

We denote the conditional probabilities of each response type by

$$p_{ij} = p_{ij}(X, A) = \mathbb{P}(Y(0) = i, Y(1) = j \mid X, A).$$

By exhaustiveness of these types, $p_{00} + p_{01} + p_{10} + p_{11} = 1$. (Note p_{ij} are *random variables*.)

We consider evaluating the fairness of a personalized intervention policy $Z = Z(X, A) \in \{0, 1\}$, which assigns interventions based on observable features X, A (potentially just X). Note that by definition, the intervention has zero effect on non-responders, negative effect on anti-responders, and a positive effect only on responders. Therefore, in seeking to benefit individuals with limited

resources, the personalized intervention policy should seek to target only the responders. Naturally, response type is unobserved and the policy can only mete out interventions based on observables.

In classification settings, minimum-error classifiers on the efficient frontier of type-I and -II errors are given by Bayes classifiers that threshold the probability of a positive label. In personalized interventions, policies that are on the efficient frontier of social welfare (fraction of positive outcomes, $\mathbb{P}(Y(Z) = 1)$) and program cost (fraction intervened on, $\mathbb{P}(Z = 1)$) are given by thresholding ($Z = \mathbb{I}[\tau \geq \theta]$) the *conditional average treatment effect* (CATE):

$$\begin{aligned}\tau &= \tau(X, A) = \mathbb{E}[Y(1) - Y(0) \mid X, A] = p_{01} - p_{10} \\ &= \mathbb{P}(Y = 1 \mid T = 1, X, A) - \mathbb{P}(Y = 1 \mid T = 0, X, A),\end{aligned}$$

where the latter equality follows by the assumed ignorable treatment assignment. Estimating τ from unconfounded data using flexible models has been the subject of much recent work [27, 50, 53].

We consider observational fairness metrics in analogy to the classification setting, where the “true label” of an individual is their *responder status*, $R = \mathbb{I}[Y(1) > Y(0)]$. We define the analogous true positive rate and true negative rate for the intervention assignment Z , conditional on the (unobserved) events of an individual being a responder or non-responder, respectively:

$$\begin{aligned}\text{TPR}_a &= \mathbb{P}(Z = 1 \mid A = a, Y(1) > Y(0)), \\ \text{TNR}_a &= \mathbb{P}(Z = 0 \mid A = a, Y(1) \leq Y(0)).\end{aligned}\tag{1}$$

2.1 Interpreting Disparities for Personalized Interventions

The use of predictive models to deliver interventions can induce disparate impact if responding (respectively, non-responding) individuals of different groups receive the intervention at disproportionate rates under the treatment policy. This can occur even with efficient policies that threshold the true CATE τ and can arise from the disparate predictiveness of X, A of response type (i.e., how far p_{ij} are from 0 and 1). This is problematic because the choice of features X is usually made by the intervening agent (e.g., government agency, etc.).

We discuss one possible interpretation of TPR or TNR disparities in this setting when the intervention is the bestowal of a benefit, like access to job training or case management. From the point of view of the intervening agent, there are specific program goals, such as employment of the target individual within 6 months. Therefore, false positives are costly due to program cost and false negatives are missed opportunities. But outcomes also affect the individual’s utility. Discrepancies in TPR across values of A are of concern since they suggest that the needs of those who could actually benefit from intervention (responders) in one group are not being met at the same rates as in other groups. Arguably, for benefit-bestowing interventions, TPR discrepancies are of greater concern. Nonetheless, from the point of view of the individual, the intervention may always grant some positive resource (e.g., from the point of view of well-being),

regardless of responder status, since it corresponds to access to a good (and the individual can gain other benefits from job training that may not necessarily align with the intervener’s program goals, such as employment in 1 year or personal enrichment). If so, then TNR discrepancies across values of A imply a “disparate benefit of the doubt” such that the policy disparately over-benefits one group over another using the limited public resource without the cover of advancing the public program’s goal, which may raise fairness and envy concerns, especially since this “waste” is at the cost of more slots for responders.

Beyond assessing disparities in TPR and TNR for one fixed policy, we will also use our ability to assess these over varying CATE thresholds in order to compute xAUC metrics [33] in Section 6. These give the disparity between the probabilities that a non-responder from group a is ranked above a responder from group b and vice-versa. Thus, they measure the disproportionate access one group gets relative to another in *any* allocation of resources that is non-decreasing in CATE.

We emphasize that the identification arguments and bounds that we present on fairness metrics are primarily intended to facilitate the *assessment* of disparities, which may require further inquiry as to their morality and legality, not necessarily to promote statistical parity via adjustments such as group-specific thresholds, though that is also possible using our tools. We defer a more detailed discussion to Section 8 and re-emphasize that assessing the distribution of outcome-conditional model errors are of central importance both in machine learning [10, 25, 45] and in the economic efficiency of targeting resources [14, 16, 44].

3 Related Work

[40] consider estimating joint treatment effects of race and treatment under a deep latent variable model to reconstruct unobserved confounding. For evaluating fairness of policies derived from estimated effects, they consider the gap in population accuracy $\text{Acc}_a = \mathbb{P}(Z = Z^* \mid A = a)$, where $Z^* = \mathbb{I}[\tau(X) > 0]$ is the (identifiable) optimal policy. In contrast, we highlight the unfairness of even optimal policies and focus on outcome-conditional error rates (TPR, TNR), where the non-identifiability of responder status introduces challenges regarding identifiability.

The issue of model evaluation under the censoring problem of selective labels has been discussed in situations such as pretrial detention, where detention censors outcomes [32, 38]. Sensitivity analysis is used in [30] to account for possible unmeasured confounders. The distinction is that we focus on the targeted delivery of interventions with unknown (but estimated) causal effects, rather than considering classifications that induce one-sided censoring but have definitionally known effects.

Our emphasis is distinct from other work discussing fairness and causality that uses graphical causal models to decompose predictive models along causal pathways and assessing the normative validity of path-specific effects [34, 37],

such as the effect of probabilistic hypothetical interventions on race variables or other potentially immutable protected attributes. When discussing treatments, we here consider interventions corresponding to allocation of concrete resources (e.g., give job training), which are in fact physically manipulable by an intervening agent. The correlation of the intervention’s *conditional average* treatment effects by, say, race and its implications for downstream resource allocation are our primary concern.

There is extensive literature on partial identification in econometrics, e.g. [43]. In contrast to previous work that analyzes partial identification of average treatment effects when data is confounded and using monotonicity to improve precision [6, 13, 43], we focus on unconfounded (e.g., RCT) data and achieve full identification by assuming monotonicity and consider sensitivity analysis bounds for *nonlinear* functionals of partially identified sets, namely, true positive and false positive rates.

4 Identifiability of Disparate Impact Metrics

Since the definitions of the disparate impact metrics in Eq. (1) are conditioned on an unobserved event, such as the response event $Y(1) > Y(0)$, they actually cannot be identified from the data, even under ignorable treatment. That is, the values of $\text{TPR}_a, \text{TNR}_a$ can vary even when the joint distribution of (X, A, T, Y) remains the same, meaning the data we see cannot possibly tell us about the specific value of $\text{TPR}_a, \text{TNR}_a$.

Proposition 1. $\text{TPR}_a, \text{TNR}_a$ (or discrepancies therein over groups) are generally not identifiable.

Essentially, Proposition 1 follows because the data only identifies the marginals $p_{10} + p_{11}, p_{01} + p_{11}$ while $\text{TPR}_a, \text{TNR}_a$ depend on the joint via p_{01} , which can vary even while marginals are fixed. Since this can vary independently across values of A , discrepancies are not identifiable either.

4.1 Identification under Monotonicity

We next show identifiability if we impose the additional assumption of monotone treatment response.

Assumption 1 (Monotone treatment response). $Y(1) \geq Y(0)$. (Equivalently, $p_{10} = 0$.)

Assumption 1 says that anti-responders do not exist. In other words, the treatment either does nothing (e.g., an individual would have gotten a job or not gotten a job, regardless of receiving job training) or it benefits the individual (would get a job if and only if receive job training), but it never harms the individual. This assumption is reasonable for positive interventions. As [31] points out, policy learning in this setting is equivalent to the binary classification problem of predicting responder status.

Proposition 2. *Under Assumption 1,*

$$\begin{aligned} \text{TPR}_a &= \frac{\mathbb{E}[\tau \mid A = a, Z = 1] \mathbb{P}(Z = 1 \mid A = a)}{\mathbb{E}[\tau \mid A = a]}, \\ \text{TNR}_a &= \frac{\mathbb{E}[(1 - \tau) \mid A = a, Z = 0] \mathbb{P}(Z = 0 \mid A = a)}{\mathbb{E}[(1 - \tau) \mid A = a]}. \end{aligned} \quad (2)$$

Since the quantities on the right hand sides in Eq. (2) are in terms of identified quantities (functions of the distribution of (X, A, T, Y)), this proves identifiability. Given a sample and an estimate of τ , it also provides a simple recipe for estimation by replacing each average or probability by a sample version, since both A and Z are discrete.

Thus, Proposition 2 provides a novel means of assessing disparate impact of personalized interventions under monotone response. This is relevant because monotonicity is a defensible assumption in the case of many interventions that bestow an additional benefit, good, or resource, such as the ones mentioned in Section 1. Nonetheless, the validity of Assumption 1 is itself not identifiable. Therefore, should it fail even slightly, it is not immediately clear whether these disparity estimates can be relied upon. We therefore next study a sensitivity analysis by means of constructing partial identification bounds for $\text{TPR}_a, \text{TNR}_a$.

5 Partial Identification Bounds for Sensitivity Analysis

We next study the partial identification of disparate impact metrics when Assumption 1 fails, i.e., $p_{10} \neq 0$. We first state a more general version of Proposition 2. For any $\eta = \eta(X, A)$, let

$$\begin{aligned} \rho_a^{\text{TPR}}(\eta) &:= \frac{\mathbb{E}[\tau + \eta \mid A = a, Z = 1] \mathbb{P}(Z = 1 \mid A = a)}{\mathbb{E}[\tau + \eta \mid A = a]}, \\ \rho_a^{\text{TNR}}(\eta) &:= \frac{\mathbb{E}[1 - (\tau + \eta) \mid A = a, Z = 0] \mathbb{P}(Z = 0 \mid A = a)}{\mathbb{E}[1 - (\tau + \eta) \mid A = a]}. \end{aligned}$$

Proposition 3. $\text{TPR}_a = \rho_a^{\text{TPR}}(p_{10}), \text{TNR}_a = \rho_a^{\text{TNR}}(p_{10})$.

Since the anti-responder probability p_{10} is unknown, we cannot use Proposition 3 to identify $\text{TPR}_a, \text{TNR}_a$. We instead use Proposition 3 to compute bounds on them by restricting p_{10} to be in an uncertainty set. Formally, given an uncertainty set \mathcal{U} for p_{10} (i.e., a set of functions of x, a), we define the simultaneous identification region of the TPR and TNR for all groups $a \in \mathcal{A}$ as:

$$\Theta = \left\{ (\rho_a^{\text{TPR}}(\eta), \rho_a^{\text{TNR}}(\eta))_{a \in \mathcal{A}} : \eta \in \mathcal{U} \right\} \subseteq \mathbb{R}^{2 \times |\mathcal{A}|}.$$

For brevity, we will let $\rho_a(\eta) = (\rho_a^{\text{TPR}}(\eta), \rho_a^{\text{TNR}}(\eta))$ and $\rho(\eta) = (\rho_a(\eta))_{a \in \mathcal{A}}$.

The set Θ describes all possible simultaneous values of the group-conditional true positive and true negative rates. As long as $\forall \eta \in \mathcal{U}$ we have $0 \leq \eta(X, A) \leq$

$\min(\mathbb{P}(Y = 1 | T = 0, X, A), \mathbb{P}(Y = 0 | T = 1, X, A))$ (which is identified from the data) by Proposition 3 this set is necessarily sharp [43] given only the restriction that $p_{10} \in \mathcal{U}$. (In particular, this bound on η can be achieved by just point-wise clipping \mathcal{U} with this identifiable bound as necessary.) That is, given a joint on (X, A, T, Y) , on the one hand, every $\rho \in \Theta$ is realized by some full joint distribution on $(X, A, T, Y(0), Y(1))$ with $p_{10} \in \mathcal{U}$, and on the other hand, every such joint gives rise to a $\rho \in \Theta$. In other words, Θ is an *exact* characterization of the in-fact possible simultaneous values of the group-conditional TPRs and TNRs.

Therefore, if, for example, we are interested in the minimal and maximal possible values for the true (unknown) TPR discrepancy between groups a and b , we should seek to compute $\inf_{\rho \in \Theta} \rho_a^{\text{TPR}} - \rho_b^{\text{TPR}}$ and $\sup_{\rho \in \Theta} \rho_a^{\text{TPR}} - \rho_b^{\text{TPR}}$. More generally, for any $\mu \in \mathbb{R}^{2 \times |\mathcal{A}|}$, we may wish to compute

$$h_{\Theta}(\mu) := \sup_{\rho \in \Theta} \mu^{\top} \rho. \quad (3)$$

Note that this, for example, covers the above example since for any μ we can also take $-\mu$. The function h_{Θ} is known as the *support function* of Θ [48]. Not only does the support function provide the maximal and minimal contrasts in a set, it also exactly characterizes its convex hull. That is, $\text{Conv}(\Theta) = \{\rho : \mu^{\top} \rho \leq h_{\Theta}(\mu) \forall \mu\}$. So computing h_{Θ} allows us to compute $\text{Conv}(\Theta)$.

Our next result gives an explicit program to compute the support function when \mathcal{U} has a product form of within-group uncertainty sets:

$$\mathcal{U} = \{\eta : \eta(\cdot, a) \in \mathcal{U}_a \forall a \in \mathcal{A}\}, \quad (4)$$

which leads to $\Theta = \prod_{a \in \mathcal{A}} \Theta_a$ where $\Theta_a = \{\rho_a(\eta_a) : \eta_a \in \mathcal{U}_a\}$.

Proposition 4. *Let $r_a^z := \mathbb{P}(Z = z | A = a)$ and $\tau_a^z := \mathbb{E}[\tau | A = a, Z = z]$. Suppose \mathcal{U} is as in (4). Then Eq. (3) can be reformulated as:*

$$\begin{aligned} h_{\Theta}(\mu) &= \sum_{a \in \mathcal{A}} h_{\Theta_a}(\mu_a) \\ h_{\Theta_a}(\mu_a) &= \sup_{\omega_a, t_a} \mu_a^{\text{TPR}} r_a^1 (t_a \tau_a^1 + \mathbb{E}[\omega_a(X) | A = a, Z = 1]) \\ &\quad + \frac{\mu_a^{\text{TNR}} r_a^0}{t_a - 1} (t_a (1 - \tau_a^0) + \mathbb{E}[\omega_a(X) | A = a, Z = 0]) \\ \text{s.t. } \omega_a(\cdot) &\in t_a \mathcal{U}_a, \quad t_a (r_a^0 \tau_a^0 + r_a^1 \tau_a^1) + \mathbb{E}[\omega_a | A = a] = 1. \end{aligned}$$

For a fixed value of t_a , the above program is a linear program, given that \mathcal{U}_a is linearly representable. Therefore a solution may be found by grid search on the univariate t_a . Moreover, if $\mu_a^{\text{TPR}} = 0$ or $\mu_a^{\text{TNR}} = 0$, the above remains a linear program even with t_a as a variable [17]. With this, we are able to express group-level disparities through assessing the support function at specific contrast vectors μ .

5.1 Partial Identification under Relaxed Monotone Treatment Response

We next consider the implications of the above for the following relaxation of the monotone treatment response assumption:

Assumption 2 (B -relaxed monotone treatment response). $p_{10} \leq B$.

Note that Assumption 2 with $B = 0$ recovers Assumption 1 and Assumption 2 with $B = 1$ is a vacuous assumption. In between these two extremes we can consider milder or stronger violations of monotone response and the partial identification bounds they corresponds to. This provides us with a means of sensitivity analysis of the disparities we measure, recognizing that monotone response may not hold exactly and that disparities may not be exactly identifiable. For the rest of the paper, we focus solely on partial identification under Assumption 2. Note that Assumption 2 corresponds exactly to the uncertainty set $\mathcal{U}_B = \{\eta : 0 \leq \eta(X, A) \leq \min(B, \mathbb{P}(Y = 1 | T = 0, X, A), \mathbb{P}(Y = 0 | T = 1, X, A))\}$. We define $\Theta_B = \prod_{a \in \mathcal{A}} \Theta_{B,a}$ to be the corresponding identification region.

Under Assumption 2, our bounds take on a particularly simple form. Let

$$\mathcal{B}_a^z(B) = \mathbb{E}[\min(B, \mathbb{P}(Y = 1 | T = 0, X, A), \mathbb{P}(Y = 0 | T = 1, X, A)) | A = a, Z = z]$$

and define

$$\begin{aligned} \bar{\rho}_a^{\text{TPR}}(B) &= \frac{(\tau_a^1 + \mathcal{B}_a^1(B))r_a^1}{\tau_a^0 r_a^0 + (\tau_a^1 + \mathcal{B}_a^1(B))r_a^1}, \\ \underline{\rho}_a^{\text{TPR}}(B) &= \frac{\tau_a^1 r_a^1}{(\tau_a^0 + \mathcal{B}_a^0(B))r_a^0 + \tau_a^1 r_a^1}, \\ \bar{\rho}_a^{\text{TNR}}(B) &= \frac{(1 - \tau_a^0)r_a^0}{(1 - \tau_a^0)r_a^0 + (1 - \tau_a^1 - \mathcal{B}_a^1(B))r_a^1}, \\ \underline{\rho}_a^{\text{TNR}}(B) &= \frac{(1 - \tau_a^0 - \mathcal{B}_a^0(B))r_a^0}{(1 - \tau_a^0 - \mathcal{B}_a^0(B))r_a^0 + (1 - \tau_a^1)r_a^1}. \end{aligned}$$

Proposition 5. *Suppose Assumption 2 holds. Then $[\underline{\rho}_a^{\text{TPR}}(B), \bar{\rho}_a^{\text{TPR}}(B)]$ and $[\underline{\rho}_a^{\text{TNR}}(B), \bar{\rho}_a^{\text{TNR}}(B)]$ are the sharp identification intervals for TPR_a and TNR_a , respectively. Moreover, $(\underline{\rho}_a^{\text{TPR}}(B), \underline{\rho}_a^{\text{TNR}}(B)) \in \Theta_{B,a}$ and $(\bar{\rho}_a^{\text{TPR}}(B), \bar{\rho}_a^{\text{TNR}}(B)) \in \Theta_{B,a}$, i.e., the two extremes are simultaneously achievable.*

6 Partial Identification of Group Disparities and ROC and xROC Curves

We discuss diagnostics to summarize possible impact disparities across a range of possible policies.

TPR and TNR disparity. Discrepancies in model errors (TPR or TNR) are of interest when auditing classification performance on different groups with a given, fixed policy Z . Under Assumption 1, they are identified by Proposition 2. Under violations of Assumption 1, we can consider their partial identification bounds. If the *minimal* disparity remains nonzero, that provides strong evidence of disparity. Similarly, if the *maximal* disparity is large, a responsible decision maker should be concerned about the possibility of a disparity.

Under Assumption 2, Proposition 5 provides that the sharp identification intervals of $\text{TPR}_a - \text{TPR}_b$ and $\text{TNR}_a - \text{TNR}_b$ are, respectively, given by

$$\begin{aligned} & [\underline{\rho}_a^{\text{TPR}}(B) - \bar{\rho}_b^{\text{TPR}}(B), \bar{\rho}_a^{\text{TPR}}(B) - \underline{\rho}_b^{\text{TPR}}(B)], \\ & [\underline{\rho}_a^{\text{TNR}}(B) - \bar{\rho}_b^{\text{TNR}}(B), \bar{\rho}_a^{\text{TNR}}(B) - \underline{\rho}_b^{\text{TNR}}(B)]. \end{aligned} \quad (5)$$

Given effect scores τ , we can then use this to plot *disparity curves* by plotting the endpoints of Eq. (5) for policies $Z = \mathbb{I}[\tau \geq \theta]$ for varying thresholds θ .

Robust ROC Curves We first define the analogous *group-conditional ROC curve* corresponding to a CATE function τ . These are the parametric curves traced out by the pairs $(1 - \text{TNR}_a, \text{TPR}_a)$ of policies that threshold the CATE for varying thresholds. To make explicit that we are now computing metrics for different policies, we use the notation $\rho(\eta; \tau \geq \theta)$ to refer to the metrics of the policy $Z = \mathbb{I}[\tau \geq \theta]$. Under Assumption 1, Proposition 2 provides point identification of the group-conditional ROC curve:

$$\text{ROC}_a(\tau) := \{(1 - \rho_a^{\text{TNR}}(0; \tau \geq \theta), \rho_a^{\text{TPR}}(0; \tau \geq \theta)) : \theta \in \mathbb{R}\}$$

When Assumption 1 fails, we cannot point identify $\text{TPR}_a, \text{TNR}_a$ and correspondingly we cannot identify $\text{ROC}_a(\tau)$. We instead define the *robust ROC curve* as the union of all partially identified ROC curves. Specifically:

$$\Theta_a^{\text{ROC}}(\tau) := \{(1 - \rho_a^{\text{TNR}}(\eta_a; \tau \geq \theta), \rho_a^{\text{TPR}}(\eta_a; \tau \geq \theta)) : \theta \in \mathbb{R}, \eta_a \in \mathcal{U}_a\}.$$

Plotted, this set provides a visual representation of the region that the true ROC curve can lie in. We next prove that under Assumption 2, we can easily compute this set as the area between two curves.

Proposition 6. *Let $\mathcal{U} = \mathcal{U}_B$. Then $\Theta_a^{\text{ROC}}(\tau)$ is given as the area between the two parametric curves $\underline{\text{ROC}}_a(\tau) := \{(1 - \underline{\rho}_a^{\text{TNR}}(B; \tau \geq \theta), \underline{\rho}_a^{\text{TPR}}(B; \tau \geq \theta)) : \theta \in \mathbb{R}\}$ and $\bar{\text{ROC}}_a(\tau) := \{(1 - \bar{\rho}_a^{\text{TNR}}(B; \tau \geq \theta), \bar{\rho}_a^{\text{TPR}}(B; \tau \geq \theta)) : \theta \in \mathbb{R}\}$.*

This follows because the extremes are simultaneously achievable as noted in Proposition 5. We highlight, however, that the lower (resp., upper) ROC curve may not be simultaneously realizable as an ROC curve of any single policy.

Robust xROC Curves Comparison of group-conditional ROC curves may not necessarily show impact disparities as, even in standard classification settings ROC curves can overlap despite disparate impacts [25, 33]. At the same time,

comparing disparities for fixed policies Z with fixed thresholds may not accurately capture the impact of using τ for rankings. [33] develop the xAUC metric for assessing the *bipartite ranking* quality of risk scores, as well as the analogous notion of a xROC curve which parametrically plots the TPR of one group vs. the FPR of *another group*, at any fixed threshold. This is relevant if effect scores τ are used for downstream decisions by different facilities with different budget constraints or if the score is intended to be used by a “human-in-the-loop” exercising additional judgment, e.g., individual caseworkers as in the encouragement design of [12].

Under Assumption 1, we can point identify $\text{TPR}_a, \text{TNR}_a$, so, following [33], we can define the point-identified xROC curve as

$$\text{xROC}_{a,b}(\tau) = \{(1 - \rho_b^{\text{TNR}}(0; \tau \geq \theta), \rho_a^{\text{TPR}}(0; \tau \geq \theta)) : \theta \in \mathbb{R}\}.$$

Without Assumption 1, we analogously define the *robust xROC* curve as the union of all partially identified xROC curves:

$$\Theta_{a,b}^{\text{xROC}}(\tau) = \{(1 - \rho_b^{\text{TNR}}(\eta_a; \tau \geq \theta), \rho_a^{\text{TPR}}(\eta_a; \tau \geq \theta)) : \theta \in \mathbb{R}, \eta_a \in \mathcal{U}_a\}.$$

Proposition 7. *Let $\mathcal{U} = \mathcal{U}_B$. Then $\Theta_{a,b}^{\text{xROC}}(\tau)$ is given as the area between the two parametric curves $\underline{\text{xROC}}_{a,b}(\tau) := \{(1 - \underline{\rho}_b^{\text{TNR}}(B; \tau \geq \theta), \underline{\rho}_a^{\text{TPR}}(B; \tau \geq \theta)) : \theta \in \mathbb{R}\}$ and $\overline{\text{xROC}}_{a,b}(\tau) := \{(1 - \overline{\rho}_b^{\text{TNR}}(B; \tau \geq \theta), \overline{\rho}_a^{\text{TPR}}(B; \tau \geq \theta)) : \theta \in \mathbb{R}\}$.*

This follows because \mathcal{U}_B takes the form of a product set over $a \in \mathcal{A}$.

7 Case Study: Personalized Job Training (Behaghel et al.)

We consider a case study from a three-armed large randomized controlled trial that randomly assigned job-seekers in France to a control-group, a job training program managed by a public vendor, and an out-sourced program managed by a private vendor [11]. While the original experiment was interested in the design of contracts for program service delivery, we consider a task of heterogeneous causal effect estimation, motivated by interest in personalizing different types of counseling or active labor market programs that would be beneficial for the individual. Recent work in policy learning has also considered personalized job training assignment [35, 52] and suggested excluding sensitive attributes from the input to the decision rule for fairness considerations, but without consideration of fairness in the causal effect estimation itself and how significant impact disparities may still remain after excising sensitive attributes because of it.

We focus on the public program vs. control arm, which enrolled about 7950 participants in total, with $n_1 = 3385$ participants in the public program. The treatment arm, $T = 1$, corresponds to assignment to the public program. The original analysis suggests a small but statistically significant positive treatment effect of the public program, with an ATE of 0.023. We omit further details on




Figure 1: TPR and TNR disparity curves and bounds on French job training dataset (Eq. (5))

the data processing to Appendix B. We consider the group indicators: *nationality* (0, 1 denoting French nationals vs. non-French, respectively), *gender* (denoting woman vs. non-woman), and *age* (below the age of 26 vs. above). (Figures for gender appear in Appendix B.)

In Fig. 1, we plot the identified “disparity curves” of Eq. (5) corresponding to the maximal and minimal sensitivity bounds on TPR and TNR disparity between groups. Levels of shading correspond to different values of B , with color legend at right. We learn τ by the Generalized Random Forests method of [5, 53] and use sample splitting, learning τ on half the data and using our methods to assess bounds on ρ^{TPR} , ρ^{TNR} and other quantities with out-of-sample estimates on the other half of the data. We bootstrap over 50 sampled splits and average disparity curves to reduce sample uncertainty.

In general, the small probability of being a responder leads to increased sensitivity of TPR estimates (wide identification bands). The curves and sensitivity bounds suggest that with respect to nationality and gender, there is small or no disparity in true positive rates but the true negative rates for nationality, gender, and age may differ significantly across groups, such that non-women would have




Figure 2: ROC and xROC for $A = \text{nationality, age}$ on French job training dataset

a higher chance of being bestowed job-training benefits when they are in fact not responders. However, TPR disparity by age appears to hold with as much as -0.1 difference, with older actually-responding individuals being less likely to be given job training than younger individuals. Overall, this suggests that differences in heterogeneous treatment effects across age categories could lead to significant adverse impact on older individuals.

This is similarly reflected in the robust ROC, xROC curves (Fig. 2). Despite possibly small differences in ROCs, the xROCs indicate strong disparities: the sensitivity analysis suggests that the likelihood of ranking a non-responding young individual above a responding old individual (xAUC [33]) is clearly larger than the symmetric error, meaning that older individuals who benefit from the treatment may be disproportionately shut out of it as seats are instead given to non-responding younger individuals.

8 Discussion and Conclusion

We presented identification results and bounds for assessing disparate model errors of causal-effect maximizing treatment policies, which can lead disparities

in access to those who stand to benefit from treatment across groups. Whether this is “unfair” would naturally rely on one’s normative assumptions. One such is “claims across outcomes,” that individuals have a claim to the public intervention if they stand to benefit, which can be understood within [1]’s axiomatic justification of fair distribution. There may also be other justice-based considerations, e.g. minimax fairness. We discuss this more extensively in Appendix C.

With the new ability to *assess* disparities using our results, a second natural question is whether these disparities warrant adjustment, which is easy to do given our tools combined with the approach of [25]. This question again is dependent both on one’s viewpoint and ultimately on the problem context, and we discuss it further in Appendix C. Regardless of normative viewpoints, auditing allocative disparities that would arise from the implementation of a personalized rule must be a crucial step of a responsible and convincing program evaluation. We presented fundamental identification limits to such assessments but provided sensitivity analyses that can support reliable auditing.

References

- [1] M. Adler. *Well-Being and Fair Distribution*.
- [2] J. D. Angrist, G. W. Imbens, and D. B. Rubin. Identification of causal effects using instrumental variables. *Journal of the American statistical Association*, 1996.
- [3] J. Angwin, J. Larson, S. Mattu, and L. Kirchner. Machine bias. Online., May 2016.
- [4] S. Athey. Beyond prediction: Using big data for policy problems. *Science*, 2017.
- [5] S. Athey, J. Tibshirani, S. Wager, et al. Generalized random forests. *The Annals of Statistics*, 47(2):1148–1178, 2019.
- [6] A. Balke and J. Pearl. Bounds on treatment effects from studies with imperfect compliance. *Journal of the American Statistical Association*, 92(439):1171–1176, 1997.
- [7] A. Banerjee, E. Duflo, N. Goldberg, D. Karlan, R. Osei, W. Parienté, J. Shapiro, B. Thuysbaert, and C. Udry. A multifaceted program causes lasting progress for the very poor: Evidence from six countries. *Science*, 348(6236):1260799, 2015.
- [8] C. Barabas, K. Dinakar, J. Ito, M. Virza, and J. Zittrain. Interventions over predictions: Reframing the ethical debate for actuarial risk assessment. *Proceedings of Machine Learning Research*, 2017.

- [9] S. Barocas and A. Selbst. Big data’s disparate impact. *California Law Review*, 2014.
- [10] S. Barocas, M. Hardt, and A. Narayanan. *Fairness and Machine Learning*. fairmlbook.org, 2018. <http://www.fairmlbook.org>.
- [11] L. Behaghel, B. Crépon, and M. Gurgand. Private and public provision of counseling to job seekers: Evidence from a large controlled experiment. *American Economic Journal: Applied Economics*, 2014.
- [12] S. Behncke, M. Frölich, and M. Lechner. Targeting labour market programmes: results from a randomized experiment. *Work. Pap. 3085, IZA (Inst. Study Labor)*, 2007.
- [13] A. Beresteanu, I. Molchanov, and F. Molinari. Partial identification using random set theory. *Journal of Econometrics*, 166(1):17–32, 2012.
- [14] M. Berger, D. A. Black, and J. A. Smith. *Econometric evaluation of labour market policies*, chapter Evaluating Profiling as a Means of Allocating Government Service, pages 59–84. 2000.
- [15] D. Bhattacharya and P. Dupas. Inferring welfare maximizing treatment assignment under budget constraints. *Journal of Econometrics*, 2012.
- [16] C. Brown, M. Ravallion, and D. van de Walle. A poor means test? econometric targeting in africa. *Policy Research Working Paper 7915: World Bank Group, Development Research Group, Human Development and Public Services Team*, 2016.
- [17] A. Charnes and W. W. Cooper. Programming with linear fractional functionals. *Naval research logistics (NRL)*, 9(3-4):181–186, 1962.
- [18] I. Chen, F. Johansson, and D. Sontag. Why is my classifier discriminatory? *In Advances in Neural Information Processing Systems 31*, 2018.
- [19] A. Chouldechova. Fair prediction with disparate impact: A study of bias in recidivism prediction instruments. *In Proceedings of FATML*, 2016.
- [20] S. Corbett-Davies and S. Goel. The measure and mismeasure of fairness: A critical review of fair machine learning. *ArXiv preprint*, 2018.
- [21] B. Crepon and G. J. van den Berg. Active labor market policies. *Annual Review of Economics*, Vol. 8:521-546, 2016.
- [22] J. M. Davis and S. B. Heller. Using causal forests to predict treatment heterogeneity: An application to summer jobs. *American Economic Review: Papers and Proceedings*, 107(5): 546–550, 2017.
- [23] M. Dudik, D. Erhan, J. Langford, and L. Li. Doubly robust policy evaluation and optimization. *Statistical Science*, 2014.

- [24] V. Eubanks. *Automating inequality: How high-tech tools profile, police, and punish the poor*. St. Martin’s Press, 2018.
- [25] M. Hardt, E. Price, N. Srebro, et al. Equality of opportunity in supervised learning. In *Advances in Neural Information Processing Systems*, pages 3315–3323, 2016.
- [26] H. Heidari, C. Ferrari, K. Gummadi, and A. Krause. Fairness behind a veil of ignorance: A welfare analysis for automated decision making. In *Advances in Neural Information Processing Systems*, pages 1265–1276, 2018.
- [27] J. L. Hill. Bayesian nonparametric modeling for causal inference. *Journal of Computational and Graphical Statistics*, 20(1):217–240, 2011.
- [28] L. Hu and Y. Chen. Fair classification and social welfare. *arXiv preprint arXiv:1905.00147*, 2019.
- [29] G. Imbens and D. Rubin. *Causal Inference for Statistics, Social, and Biomedical Sciences*. Cambridge University Press, 2015.
- [30] J. Jung, R. Shroff, A. Feller, and S. Goel. Algorithmic decision making in the presence of unmeasured confounding. *ArXiv*, 2018.
- [31] N. Kallus. Classifying treatment responders under causal effect monotonicity. *Proceedings of International Conference on Machine Learning*, 2019.
- [32] N. Kallus and A. Zhou. Residual unfairness in fair machine learning from prejudiced data. *Forthcoming at ICML*, 2018.
- [33] N. Kallus and A. Zhou. The fairness of risk scores beyond classification: Bipartite ranking and the xauc metric. *arXiv preprint, <https://arxiv.org/pdf/1902.05826.pdf>*, 2019.
- [34] N. Kilbertus, M. Rojas-Carulla, G. Parascandolo, M. Hardt, D. Janzing, and B. Schölkopf. Avoiding discrimination through causal reasoning. *Advances in Neural Information Processing Systems 30, 2017*, 2017.
- [35] T. Kitagawa and A. Tetenov. Empirical welfare maximization. 2015.
- [36] A. Kube and S. Das. Allocating interventions based on predicted outcomes: A case study on homelessness services. *Proceedings of the AAAI Conference on Artificial Intelligence*, 2019.
- [37] M. J. Kusner, J. R. Loftus, C. Russell, and R. Silva. Counterfactual fairness. *NIPS*, 2017.
- [38] H. Lakkaraju, J. Kleinberg, J. Leskovec, J. Ludwig, and S. Mullainathan. The selective labels problem: Evaluating algorithmic predictions in the presence of unobservables. *Proceedings of KKD2017*, 2017.

- [39] L. T. Liu, S. Dean, E. Rolf, M. Simchowitz, and M. Hardt. Delayed impact of fair machine learning. *Proceedings of the 35th International Conference on Machine Learning (ICML), Stockholm, Sweden, 2018*.
- [40] D. Madras, E. Creager, T. Pitassi, and R. Zemel. Fairness through causal awareness: Learning latent-variable models for biased data. *ACM Conference on Fairness, Accountability, and Transparency (ACM FAT*) 2019, 2019*.
- [41] C. Manski. *Social Choice with Partial Knowledge of Treatment Response*. The Econometric Institute Lectures, 2005.
- [42] C. F. Manski. Monotone treatment response. *Econometrica: Journal of the Econometric Society*, pages 1311–1334, 1997.
- [43] C. F. Manski. *Partial identification of probability distributions*. Springer Science & Business Media, 2003.
- [44] L. McBride and A. Nichols. Retooling poverty targeting using out-of-sample validation and machine learning. *Policy Research Working Paper 7849 (World Bank Group, Development Economics Vice Presidency Operations and Strategy Team)*, 2016.
- [45] S. Mitchell, E. Potash, and S. Barocas. Prediction-based decisions and fairness: A catalogue of choices, assumptions, and definitions. *arXiv*, 2018.
- [46] T. Mkandawire. Targeting and universalism in poverty reduction. *Social Policy and Development*, 2005.
- [47] E. Rice. The tay triage tool: A tool to identify homeless transition age youth most in need of permanent supportive housing. 2013.
- [48] R. T. Rockafellar. *Convex analysis*. Princeton university press, 2015.
- [49] D. B. Rubin. Comments on “randomization analysis of experimental data: The fisher randomization test comment”. *Journal of the American Statistical Association*, 75(371):591–593, 1980.
- [50] U. Shalit, F. Johansson, and D. Sontag. Estimating individual treatment effect: generalization bounds and algorithms. *Proceedings of the 34th International Conference on Machine Learning, 2017*.
- [51] R. Shroff. Predictive analytics for city agencies: Lessons from children’s services. *Big data*, 5(3):189–196, 2017.
- [52] S. Wager and S. Athey. Efficient policy learning. 2017.
- [53] S. Wager and S. Athey. Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association*, (just-accepted), 2017.

A Omitted proofs

Proof of Proposition 1. To prove this we exhibit a simple example satisfying ignorability where both $\text{TPR}_a, \text{TNR}_a$ and differences therein varies while the joint distribution of (X, A, T, Y) does not.

Let $\mathcal{X} = \{0, 1\}$, $Z = \mathbb{I}[X = 1]$, $\mathbb{P}(T = t, X = x | A = a) = \frac{1}{4}$, $\mathbb{P}(A = a) = 1/|\mathcal{A}|$. To specify a joint distribution of $(X, A, T, Z, Y(1), Y(0))$ that satisfies ignorable treatment, it only remains to specify p_{ij} .

Note that in this case

$$\text{TPR}_a = \frac{p_{01}(1, a)}{p_{01}(0, a) + p_{01}(1, a)}, \quad \text{TNR}_a = \frac{1 - p_{01}(0, a)}{2 - p_{01}(0, a) - p_{01}(1, a)}.$$

The result follows by noting that where the corresponding joint distribution of (X, A, T, Z, Y) is completely specified by $p_{01} + p_{11}$, $p_{10} + p_{11}$, while p_{01} could vary as long as these sums are neither 0 nor 1. Since we can vary this independently across values of A , differences are not identifiable either. \square

Proof of Proposition 2.

$$\begin{aligned} & \mathbb{P}(Z = 1 | A = a, Y(1) > Y(0)) \\ &= \frac{\mathbb{P}(Y(1) > Y(0) | A = a, Z = 1)\mathbb{P}(Z = 1 | A = a)}{\mathbb{P}(Y(1) > Y(0) | A = a)} \\ &= \frac{\mathbb{E}[\mathbb{E}[Y(1) = 1 | \underset{A=a, X=x}{T=1}] - \mathbb{E}[Y(0) = 1 | \underset{A=a, X=x}{T=0}] | \underset{A=a}{Z=1}]\mathbb{P}(Z = 1 | A = a)}{\mathbb{E}[\mathbb{P}(Y(1) = 1 | \underset{A=a, X=x}{T=1}) - \mathbb{P}(Y(0) = 1 | \underset{A=a, X=x}{T=0}) | A = a]} \\ &= \frac{\mathbb{E}[\tau(X, A) | \underset{A=a}{Z=1}]\mathbb{P}(Z = 1 | A = a)}{\mathbb{E}[\tau(X, A) | A = a]} \end{aligned}$$

where the first equality holds by Bayes' rule, the second by iterating expectations on X and Assumption 1, and the third by unconfoundedness and consistency of potential outcomes. The proof for identification of TNR is identical for the quantity $\mathbb{P}(Z = 0 | A = a, Y(1) \leq Y(0))$. \square

Proof of Proposition 3. Recalling that CATE identifies, under violations of Assumption 1

$$\begin{aligned} \tau &= \mathbb{E}[Y(1) - Y(0) | X, A] = p_{01} - p_{10}, \\ &= \frac{\mathbb{E}[\tau + \eta | \underset{A=a}{Z=1}]\mathbb{P}(Z = 1 | A = a)}{\mathbb{E}[\tau + \eta | A = a]} = \frac{(p_{01} - p_{10} + p_{10})\mathbb{P}(Z = 1 | A = a)}{\mathbb{E}[(p_{01} - p_{10} + p_{10}) | A = a]} \\ &= \mathbb{P}(Z = 1 | A = a, Y(1) > Y(0)) \end{aligned}$$

\square

Proof of Proposition 4. The support function evaluated at μ is:

$$\max_{\eta} \sum_{a \in \mathcal{A}} \mu_a^{\text{TPR}} \frac{\mathbb{E}[\tau + \eta | A = a, Z = 1] r_a^1}{\mathbb{E}[\tau + \eta | A = a]} + \mu_a^{\text{TNR}} \frac{\mathbb{E}[1 - (\tau + \eta) | A = a, Z = 0] r_0^a}{\mathbb{E}[1 - (\tau + \eta) | A = a]}$$

s.t. $0 \leq \eta(x, a) \leq \mathcal{B}_a^z(B)$, $\forall x \in \mathcal{X}, \forall a \in \mathcal{A}$

We apply the Charnes-Cooper transformation [17] with the bijection $t_a = \frac{1}{\mathbb{E}[\tau + \eta | A = a]}$, $\omega_a = \eta t_a$. The denominator of the second term under this bijection is equivalently

$$\mathbb{E}[1 - (\tau + \eta) | A = a] = 1 - \frac{1}{t_a}$$

such that we can rewrite the second term as

$$\begin{aligned} & \mu_a^{\text{TNR}} r_a^0 \left(\frac{1}{1 - 1/t_a} \mathbb{E}[1 - \tau | \frac{A=a}{Z=0}] + \frac{1/t_a}{1 - 1/t_a} \mathbb{E}[\omega_a | \frac{A=a}{Z=0}] \right) \\ &= \frac{\mu_a^{\text{TNR}} r_a^0}{t_a - 1} (t \mathbb{E}[1 - \tau | \frac{A=a}{Z=0}] + \mathbb{E}[\omega_a | \frac{A=a}{Z=0}]) \end{aligned}$$

and the objective function overall as:

$$\max_{\eta} \sum_{a \in \mathcal{A}} (\mu_a^{\text{TPR}} r_a^1) (t_a \tau_a^1 + \mathbb{E}[\omega_a | \frac{A=a}{Z=1}]) + \frac{\mu_a^{\text{TNR}} r_a^0}{t_a - 1} (t_a (1 - \tau_a^0) + \mathbb{E}[\omega_a | \frac{A=a}{Z=0}])$$

The new constraint set (including the constraint yielding the definition of t_a) is:

$$\mathcal{U} = \{\mathbb{E}[\tau t_a + \omega_a | A = a] = 1, \omega_a(x, a) \leq t_a \mathcal{B}_a^z(B), \forall x \in \mathcal{X}, \forall a \in \mathcal{A}\}$$

□

Proof of Proposition 5. We first consider the case of maximizing or minimizing the TPR.

We leverage the invariance in the objective function under the surjection on $\eta(x, a)$ to its marginal expectation over a $Z = z, A = a$ partition.

$$w(x, a) = \begin{cases} \mathbb{E}[\eta | Z = 1, A = a] & \text{if } Z = 1 \\ \mathbb{E}[\eta | Z = 0, A = a] & \text{if } Z = 0 \end{cases}$$

Therefore we can reparametrize the program as optimizing over coefficients x, y of the optimal solution, $w(x, y) = x\mathbb{I}[Z = 0] + y\mathbb{I}[Z = 1]$. Define the fractional objective

$$\begin{aligned} g(\alpha, \beta) &= \frac{\mathbb{E}[\tau + x\mathbb{I}[Z = 0] + y\mathbb{I}[Z = 1] | A = a, Z = 1] \mathbb{P}(Z = 1 | A = a)}{\mathbb{E}[\tau + x\mathbb{I}[Z = 0] + y\mathbb{I}[Z = 1] | A = a]} \\ &= \frac{(\mathbb{E}[\tau | A = a, Z = 1] + y)r_a^1}{\mathbb{E}[\tau | A = a] + xr_a^0 + yr_a^1} \end{aligned}$$

First note that without loss of generality that when maximizing, we can set $x = 0$ since this decreases the objective regardless of the value of y . We can

consider the constrained problem $\max_{y \leq B} h(y)$ where $h(y) = g(0, y)$. Then we have the first and second derivatives,

$$\begin{aligned}\frac{\partial h}{\partial y} &= \frac{r_a^1(\mathbb{E}[\tau | A = a] - \mathbb{E}[\tau | Z = 1, A = a])}{(yr_a^1 + \mathbb{E}[\tau | A = a])^2}, \\ \frac{\partial^2 h}{\partial \beta^2} &= \frac{(r_a^1)^2(\mathbb{E}[\tau | A = a] - \mathbb{E}[\tau | Z = 1, A = a])}{(yr_a^1 + \mathbb{E}[\tau | A = a])^3}\end{aligned}$$

By inspection, since $y \geq 0$ we have that $\frac{\partial^2 h}{\partial y^2} \geq 0$ so the function is convex. So when maximizing h on the constraints for y , it attains optimal value at the boundary (since h is increasing). When minimizing, note that the derivative is not vanishing anywhere on the constraint set so it suffices to check the endpoints, where the minimum is achieved at $\beta = 0$.

We now consider the case of minimizing or maximizing the TNR.

Now consider a generic $f(x) = \frac{a-bx}{c-bx-dy}$ which represents the TNR sensitivity bound with $\omega = x\mathbb{I}[Z = 0] + y\mathbb{I}[Z = 1]$, and the constants

$$\begin{aligned}a &= r_a^0 - \mathbb{E}[\tau | Z = 0, A = a], & c &= 1 - \mathbb{E}[\tau | A = a] \\ b &= r_a^0, & d &= r_a^1\end{aligned}$$

Without loss of generality we know that we can set y to its upper bound B when maximizing as we are only increasing the objective value; then $c' = c - Br_a^1$. We verify that the second derivative is negative, so that the function is concave:

$$\frac{\partial^2 f}{\partial x^2} = \frac{2b^2(a-c')}{(c'-bx)^3} = \frac{2(r_a^0)^2(r_a^0 - \mathbb{E}[\tau | Z = 0, A = a]) - (1 - \mathbb{E}[\tau | A = a]) - Br_a^1}{(1 - \mathbb{E}[\tau | A = a]) - Br_a^1 - xr_a^0}$$

Checking the sign of the numerator simplifies to checking the sign of

$$a - c' = (-r_a^1 + \mathbb{E}[\tau + B | Z = 1, A = a])$$

which is negative. The denominator is lower bounded by $1 - \mathbb{E}[\tau | A = a] - Br_a^1$ which is always positive: therefore the problem is concave. The first derivative $\frac{\partial f}{\partial x} = \frac{b(a-c')}{(c'-bx)^2}$ is negative on the domain; therefore the maximum is achieved at $x = 0$. Therefore, when maximizing, $\omega = B\mathbb{I}[Z = 1]$.

For minimizing the TPR, we take a similar approach: analogously, we can set y to its lower bound without loss of generality. Following the same analysis, the function is still concave $\frac{\partial^2 f}{\partial x^2} = \frac{2b^2(a-c')}{(c'-bx)^3}$ since $-r_a^1 - \mathbb{E}[\tau | A = a] < 0$ and decreasing with nonzero first-derivative; so the minimum is achieved at $\omega = \mathbb{I}[Z = 0]B$. □

B Behaghel et al. Job Training

We processed the data using replication files available with the AEJ: Applied Economics journal electronic supplement. For the sake of simplicity, we analyze




Figure 3: Diagnostics for gender protected attribute for Section 7 (not-woman vs. woman)

the trial as if it were a randomized controlled trial (without accounting for noncompliance or different randomization probabilities that differ by region). Thus, we consider intention-to-treat effects (as intention to treat is ultimately the policy lever available). We further restricted some covariates, omitting some where personalized allocation based on these covariates seemed unlikely for fairness reasons. The covariates we retain include: length of previous employment, salary, education level, reason for unemployment, region, years of experience at previous job, statistical risk level, job search type (full-time or non-full time), wage target, time of first unemployment spell, job type, and number of children.

An interacted linear model indicates potential heterogeneity of treatment effect with significance on college education, economic layoff, those seeking work due to fixed term contracts or those with previous layoffs.

C Substantive Discussion: Fairness vs. Justice

We first caveat our use of “disparate impact”: while our selection of protected attributes parallels choices of protected attributes that appear elsewhere in the




Figure 4: ROC curves under Assumption 1 for Section 7

literature on fair machine learning, for the case of interventions, there may not be precedent from discrimination case law, nonetheless assessing fairness with respect to these social groups may be of concern. We view disparate impact in this domain as assessing fairness of outcome rates under a personalization model.

Should true positive rates be adjusted for? Our presentation of an identification strategy of fairness metrics for allocating interventions with unknown causal effects begs the question: should disparities in TPR and FNR be adjusted for in the interventional welfare setting? Is responder-accuracy parity a meaningful prescriptive notion of fairness?

One critique of outcome-conditional fair classification metrics recognizes the dependence of false positive rates on the underlying *base rate*, $\mathbb{P}(Y = 1 \mid A = a)$, [19, 20]. The equivalent situation occurs when the within-group ATE varies by the protected attribute, e.g. $\mathbb{E}[\tau \mid A = a]$ differs.

Ultimately, external domain knowledge is required to adjudicate whether group-wide disparities in ATE should be adjusted for, or to decide which normative notion of distributive justice or fairness is appropriate. For example, consider the case of job training. From an economic perspective, multiple mechanisms could explain heterogeneity in CATE by race. Active labor market programs (see [21]) may be less effective for one group vs. another group due to the presence of labor-market discrimination. Alternatively, they could be less effective due to correlation of group status and efficacy that is mediated by occupation choice: one group may be more interested in labor markets where the primary benefits of job search counseling, in reducing search frictions, are not barriers to employment in the first place relative to other factors such as skills gaps. Intuitively, the former mechanism of ATE variation by group reflects a notion of “disparity” which remains problematic, while the latter may seem to reflect an unproblematic causal mechanism. While mediation analysis and fairness defined in terms of path-specific effects could further decompose the treatment effect along these stated mechanisms, in policy settings, collecting all of the relevant information can be burdensome, and deciding on a causal graph can be difficult.

Claims Across Outcomes We first outline different frameworks for thinking

about fairness/equity of algorithms and interventions. Analogous to the proposals arising from metrics proposed in fairness in machine learning, one might view the decision-maker’s concern to be of ensuring *accuracy* parity, that the decisions meted out are overall beneficial to individual. We view a theory of fairness that assesses disparities in outcome-conditional error rates in the context of a theory of normative claims arising from “claims across outcomes”. [1] develops a “claims across outcomes” framework of fairness and social welfare, in the context of an overall welfarist theory of justice.

On the one hand, fair classification from the point of view of assessing or equalizing TPR or TNR disparities may be interpreted in a claims context as: for an individual with “true outcome” Y and covariates X , an individual with the true label $Y = 1$ as having a comparative claim for $\hat{Y} = 1$, if the predictor \hat{Y} is an allocation tool. We can map the setting of personalized interventions to the “claims across outcomes” setting: the potential outcomes framework posits for each individual the random variables of outcomes $Y(0), Y(1)$. In the responder setting, the true label is responder status $Y(1) > Y(0)$. However, since these are *jointly unobservable*, in situations where heterogeneous treatment effects are plausible, the best guess is an individual-level treatment effect conditional on covariates, $\mathbb{E}[Y(0) | X = x], \mathbb{E}[Y(1) | X = x]$. In this interventional setting, one can think of individuals having claims in favor of favorable outcomes, e.g. a claim in favor of $Y(1)$ if $Y(1) > Y(0)$.

For the case of interventions, classification decisions Z are allocative of real interventions, and we argue that implicitly, the consideration of social welfare (balancing efficiency and program costs) is an important factor in the original design of social programs or personalized interventions. This is in sharp contrast to the literature on fair classification which considers settings such as lending in finance, or risk prediction in the criminal justice system, where overriding concerns are primarily those of *vendor* utility.

On the other side of the spectrum, we can recall axiomatically justified social welfare functions that apply to the case of *deterministic* resource allocation, where outcomes are generally known. A decision-maker might also be concerned with equity considerations, adopting a min-max welfare criterion, appealing to Rawlsian justice frameworks. Another approach is simply assessing the population cardinal welfare of the allocation, e.g. the policy value $\mathbb{E}[Y(\pi(X))]$ or a social-welfare transformation thereof, $\mathbb{E}[g(Y(\pi(X)))]$. The literature on policy learning addresses welfare functionals that are linear functionals of potential outcomes, see [35]. Cardinal welfare constraints such as those studied in [26] can be applied with an imputed CATE function.

Comparison to other work on fair classification and welfare. [39] study the implications of classifier-based decisions, as well as proposals for statistical parity, on group welfare. Their work addresses selection rules that have known marginal impacts by group. [28] studies the welfare weights implied by classification parity metrics and shows that enforcing classification parity metrics are not Pareto-improving. Rather than studying the welfare implications of classification

parity, we are concerned with assessing non-identifiable model errors in the causal-effect personalized intervention setting. Since in the personalized intervention setting, welfare is a primary objective for the Planner (e.g. social services, or social protection more broadly), modulo cost considerations, combining the distributional information from identification of classification errors with other social welfare objectives is of possible interest.

We next aim to provide concrete examples of discussions regarding the distributional impacts of interventions, in order to provide additional context on different contexts wherein different notions of “fairness” from the fairness in machine learning literature map onto welfare or justice concerns, as stated in discussions on interventional outcomes.

Lexicographic fairness or maximin (Rawlsian) fairness.

In a large multi-site graduation trial on testing an intensive, composite intervention targeted at the "ultra-poor", which comprised wraparound services including coaching and revenue-generating resources, still the poorest seemed to benefit least from the intervention in terms of sustained revenue [7]. In this setting, concerns about maximin fairness (Rawlsian justice) might override considerations of efficiency insofar as one might be willing to invest resources to help the worst-off on humanitarian grounds.

Universalism.

Criticisms of targeted policies in general note practical difficulties introduced by imposing and enforcing eligibility guidelines. [46]. Although discussion of resource constraints may be used to justify a targeting scheme, critics of targeting argue that the most efficient targeting is not as welfare-improving as simply advocating for greater resources [24].

Additional distributional preferences on $Y(Z)$ with respect to equitable or redistributive aims of the policy.

[14] consider profiling based on covariates as a means of allocating government services, in the example of allocating predicting unemployment duration to allocate reemployment services. They outline competing equity vs. efficiency concerns, in the case that unemployment duration is correlated with treatment efficacy (e.g efficacy of reemployment services), and conclude that “ tradeoffs between alternative social goals in designing profiling systems are likely to be empirically important... the form and extent of these tradeoffs may depend on empirical relationships between the impacts of the program being allocated and the equity-related characteristics of potential participants." While outcome-conditional true positive rates or true negative rates compare model performance across binary protected attributes, program designers may remain concerned regarding the distribution of benefits. [?] consider “removing the veil of ignorance” under the simplifying of constant treatment response to consider distributional (quantile) treatment effects, as a relaxation of the anonymity axiom of cardinal social welfare. Distributional preferences are relevant when program designers are concerned about model performance at finer-grained levels than discrete protected attribute.