Fairness in Risk Assessment Instruments: Post-Processing to Achieve Counterfactual Equalized Odds

Alan Mishler
Department of Statistics
Carnegie Mellon University

Edward H. Kennedy
Department of Statistics
Carnegie Mellon University

Alexandra Chouldechova
Heinz College
Carnegie Mellon University

ABSTRACT

In domains such as criminal justice, medicine, and social welfare, decision makers increasingly have access to algorithmic Risk Assessment Instruments (RAIs). RAIs estimate the risk of an adverse outcome such as recidivism or child neglect, potentially informing high-stakes decisions such as whether to release a defendant on bail or initiate a child welfare investigation. It is important to ensure that RAIs are fair, so that the benefits and harms of such decisions are equitably distributed.

The most widely used algorithmic fairness criteria are formulated with respect to observable outcomes, such as whether a person actually recidivates, but these criteria are misleading when applied to RAIs. Since RAIs are intended to inform interventions that can reduce risk, the prediction itself affects the downstream outcome. Recent work has argued that fairness criteria for RAIs should instead utilize potential outcomes, i.e. the outcomes that *would* occur in the absence of an appropriate intervention [11]. However, no methods currently exist to satisfy such fairness criteria.

In this paper, we target one such criterion, *counterfactual equalized odds*. We develop a post-processed predictor that is estimated via doubly robust estimators, extending and adapting previous post-processing approaches [16] to the counterfactual setting. We also provide doubly robust estimators of the risk and fairness properties of arbitrary fixed post-processed predictors. Our predictor converges to an optimal fair predictor at fast rates. We illustrate properties of our method and show that it performs well on both simulated and real data.

CCS CONCEPTS

• Computing methodologies → Machine learning.

KEYWORDS

fairness, risk assessment, post-processing, counterfactual

ACM Reference Format:

Alan Mishler, Edward H. Kennedy, and Alexandra Chouldechova. 2021. Fairness in Risk Assessment Instruments: Post-Processing to Achieve Counterfactual Equalized Odds. In *Conference on Fairness, Accountability, and Transparency (FAccT '21)*, March 3–10, 2021, Virtual Event, Canada. ACM, New York, NY, USA, 15 pages. <https://doi.org/10.1145/3442188.3445902>

This work is licensed under a Creative Commons Attribution International 4.0 License. *FAccT '21*, March 3–10, 2021, Virtual Event, Canada
© 2021 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-8309-7/21/03.
<https://doi.org/10.1145/3442188.3445902>

1 INTRODUCTION

Machine learning is increasingly involved in high stakes decisions in domains such as healthcare, criminal justice, and consumer finance. In these settings, ML models often take the form of Risk Assessment Instruments (RAIs): given covariates such as demographic information and an individual's medical/criminal/financial history, the model predicts the likelihood of an adverse outcome, such as a dangerous medical event, recidivism, or default on a loan. Rather than rendering an automatic decision, the model produces a "risk score," which a decision maker may take into account when deciding whether to prescribe a medical treatment, release a defendant on bail, or issue a personal loan.

The proliferation of machine learning has raised concerns that learned models may be discriminatory with respect to sensitive features like race, sex, age, and socioeconomic status. For example, there has been vigorous debate about whether a widely used recidivism prediction tool called COMPAS is biased against black defendants [1, 2, 12, 28, 30]. Concerns have also been raised about risk assessments used to identify high risk medical patients [36] and about common credit scoring algorithms such as FICO [37], among many others. Collectively, these types of algorithms directly impact a large and growing swath of the global population.

These concerns have led to an explosion of methods in recent years for developing fair models and auditing the fairness of existing models. The most widely discussed fairness criteria impose constraints on the joint distribution of a sensitive feature, an outcome, and a predictor. These "observational" fairness criteria are inappropriate for RAIs, however. RAIs are not concerned with the *observable outcomes* in the training data ("Did patients of this type historically experience serious complications?"), which are themselves a product of historical treatment decisions. Rather, they are concerned with the *potential outcomes* associated with available treatment decisions ("Would patients of this type experience complications *if not treated*?"). Because treatments are not assigned at random—doctors naturally treat the patients they think are at high risk—these are distinct questions.

Coston et al. [11] showed how RAIs that are optimized to predict observable rather than potential outcomes systematically underestimate risk for units that have historically been receptive to treatment, leading to suboptimal treatment decisions. They further showed how evaluations of the performance and fairness properties of RAIs with respect to observable outcomes are misleading. They proposed that RAIs should instead target counterfactual versions of standard performance and fairness metrics. However, they left open the question of how to develop predictors that satisfy such fairness notions.

In this paper, we develop a method to generate predictors that satisfy the fairness criterion *approximate counterfactual equalized*

odds. While many existing methods target observational fairness criteria [7, 13, 16, 21, 25, 34, 47] and various types of causally motivated fairness [24, 27, 32, 33], no methods currently exist that target counterfactual versions of standard observable fairness criteria like equalized odds. Our method post-processes an arbitrary existing predictor, extending previous post-processing methods [16] to the counterfactual setting.

Our contributions are as follows. We first define approximate counterfactual equalized odds (§2). After discussing related work (§3) and motivating the use of equalized odds over other candidate criteria (§4), we present a linear program that produces a loss-optimal post-processed predictor that satisfies this criterion (§5). We provide theoretical results that our post-processed predictor is consistent in a particular sense at rates that depend on certain nuisance parameters. We show that our method performs well on both simulated and real data (§6).

2 NOTATION AND FAIRNESS DEFINITIONS

A table listing all notational choices can be found in Appendix F.

Let A, D, Y denote a sensitive feature, decision, and outcome, respectively. We consider the setting in which all three are binary, though most of the definitions below extend readily to continuous settings. We define the counterfactual quantities of interest via the potential outcomes framework of [17, 35, 39]. Denote by Y^0, Y^1 the potential (equivalently, “counterfactual”) outcomes $Y^{D=0}, Y^{D=1}$. Y_i^d is the outcome that would be observed for unit i if, possibly contrary to fact, the decision were set to $D_i = d$. We refer to the two levels of the sensitive feature A as the two “groups,” and we use “treatment” and “intervention” synonymously with “decision.” Let S be any random variable that takes values in $\{0, 1\}$.

In most RAI settings, one of the decision options is a natural baseline corresponding to “no intervention” ($D = 0$). Examples include the risk of recidivism if a defendant is released pretrial, or the risk of neglect or abuse if a child welfare call is not screened in for further investigation. Many or most RAIs do not generate a separate risk score for the outcome associated with intervention. In the case of child welfare, for example, call screeners must screen in any case in which a child is in apparent danger of neglect or abuse, regardless of the chances that a subsequent intervention will successfully prevent that neglect or abuse.

Denote the observational and counterfactual false positive rates of S for group a by $\text{FPR}(S, a) = \mathbb{P}(S = 1 \mid Y = 0, A = a)$ and $\text{cFPR}(S, a) = \mathbb{P}(S = 1 \mid Y^0 = 0, A = a)$. For example, $\text{cFPR}(S, 0)$ could represent the chance of being falsely labeled high-risk, among those black defendants who would not actually go on to recidivate if released pretrial, while $\text{cFPR}(S, 1)$ could represent the corresponding error rate for white defendants who would not recidivate if released pretrial. Let FNR, cFNR , denote the corresponding false negative rates.

Definition 2.1. A predictor S satisfies *observational equalized odds* (oEO) with respect to A and Y if $S \perp\!\!\!\perp A \mid Y$. It satisfies *counterfactual equalized odds* (cEO) if $S \perp\!\!\!\perp A \mid Y^0$.

When A, Y , and S are all binary, equalized odds is equivalent to requiring that the corresponding false positive and false negative rates be equal for the two levels of A . Our post-processed predictor will target a relaxation of this criterion, defined below.

Definition 2.2. The *counterfactual error rate differences* for a predictor S are the differences Δ^+ and Δ^- in the cFPR and cFNR for the two groups $A = 0, A = 1$, defined as follows:

$$\Delta^+(S) = \text{cFPR}(S, 0) - \text{cFPR}(S, 1)$$

$$\Delta^-(S) = \text{cFNR}(S, 0) - \text{cFNR}(S, 1)$$

Definition 2.3. When A, Y , and S are all binary, S satisfies *approximate counterfactual equalized odds* with *fairness constraints* $\epsilon^+, \epsilon^- \in [0, 1]$ if

$$|\Delta^+(S)| \leq \epsilon^+$$

$$|\Delta^-(S)| \leq \epsilon^-$$

In general, a fairness-constrained predictor would not outperform an optimal unconstrained predictor, and in some cases, satisfying cEO exactly might degrade performance to the point that the RAI is no longer useful. This relaxation of cEO allows RAI designers to negotiate this tradeoff. This is similar in spirit to notions of approximate fairness that appear throughout the literature [13, 22, 31].

3 RELATED WORK

3.1 Observational fairness criteria

Equalized odds is one of several popular fairness criteria that impose constraints on the joint distribution of (A, Y, S) [3]. These criteria appear under a variety of names. Equalized odds is known more generally as *separation*, a term which covers settings in which these variables are not necessarily binary. The other two popular criteria in this class are *independence* ($S \perp\!\!\!\perp A$) and *sufficiency* ($Y \perp\!\!\!\perp A \mid S$). Independence also manifests as *demographic parity*, *statistical parity*, and *group fairness*. Sufficiency is equivalent to *calibration* or *predictive parity* when all three variables are binary. Variants of all three criteria may be defined for example by conditioning on additional variables.

The counterfactual versions of these criteria simply replace Y with the potential outcome Y^0 that is of interest [11]. Note that these definitions cannot accommodate more than one potential outcome, such as the vector (Y^0, Y^1) , because only one of these outcomes is observed for each unit. This is the “fundamental problem of causal inference” [17].

Except in highly constrained, unrealistic conditions, these three criteria are pairwise unsatisfiable, regardless of whether they are defined with respect to Y or Y^0 [3, 9, 26]¹. We must therefore choose and justify which criterion we wish to target.

3.2 Causal fairness criteria

The counterfactual fairness criteria just described consider potential outcomes with respect to a decision D . There is a distinct set of causally motivated fairness criteria that consider counterfactuals of the sensitive feature, or a proxy for the sensitive feature. They characterize a decision or prediction as fair if the sensitive feature or proxy does not “cause” the decision or prediction, either directly or along a prohibited pathway [24, 27, 32, 33, 45, 48]. There is some controversy over whether it is meaningful to discuss a counterfactual of a feature like race or gender [14, 18, 44]. Additionally,

¹See [19] for a set of sufficient conditions under which these unsatisfiability results disappear.

satisfying these metrics typically precludes use of most of the features that go into risk assessment, like prior history, which is not tenable in practice [11]. Finally, it is not clear that counterfactuals of the sensitive feature are useful or appropriate to consider in the context of risk assessment. For example, in the child welfare setting, workers are compelled to screen in calls whenever a child is in danger of neglect or abuse. While it is important to ensure that risk is assessed accurately for different groups, it would be inappropriate to make screen-in decisions based on what a child's risk of neglect or abuse would be *if they had been of a different race* their whole life, even if such an assessment were possible.

3.3 Ways of achieving fairness

There are three broad approaches to developing fair models: (1) preprocessing the input data to remove bias [7, 21], (2) constraining the learning process (aka “in-processing”) [13, 34, 47], and (3) post-processing a model to satisfy fairness constraints [16, 25].

Our approach can be viewed as belonging to class (3). We refer to the predictor that our method returns equivalently as a “post-processed” or “derived” predictor. Each approach has advantages and disadvantages. Many widely used RAIs are proprietary tools developed by for-profit companies, so they are not amenable to internal tinkering. Developing new, fair(er) RAIs would be costly and perhaps infeasible from a policy perspective. The advantage of post-processing in this setting is that it can be applied to models that are already in use. The predictor that our method returns requires access at runtime only to the sensitive feature and the output of the existing predictor, so in principle, it could easily be incorporated into existing risk assessment pipelines.

In particular, our approach extends the work of Hardt et al. [16], who proposed a method to post-process binary predictors to satisfy observational equalized odds (oEO) while minimizing loss with respect to observable Y . Their post-processed predictor is the solution to a simple linear program. We adapt their method to the counterfactual setting, in which the fairness criterion is approximate cEO and the loss function is weighted classification error with respect to Y^0 . Because Y^0 is not observable when $D \neq 0$, we require tools from causal inference to solve this problem. Hardt et al.'s analysis treats the joint distribution of (A, S, Y) as known and frames post-processing primarily as an optimization problem. We build on their results by not making this assumption and treating post-processing as a statistical estimation problem.

3.4 Why equalized odds?

When evaluating a predictive system, it seems natural to focus on its real-world impact rather than its outputs per se. One desirable property of a decision process is the avoidance of disparate impact. Disparate impact is a legal doctrine enshrined in U.S. law that prohibits practices which have an unjustifiable adverse impact on people who share a protected characteristic, regardless of discriminatory intent. By way of shorthand, we will say that if $D \not\perp A \mid Y^0$, then the system exhibits discriminatory disparate impact². In recidivism prediction, for example, this could mean that black defendants ($A = 0$) who would not recidivate if released ($Y^0 = 0$) are more

likely to be detained pretrial ($D = 1$) than white defendants who would not recidivate if released ($A = 1, Y^0 = 0$).

In the context of RAIs, decision makers typically have wide latitude in how they interpret and act on the risk scores, so constraining the RAI does not enforce fairness with respect to their decisions. However, if decision makers, after the introduction of the RAI, make their decisions only on the basis of the RAI scores and other variables U which are independent of the RAI and A given Y^0 , then counterfactual equalized odds will imply $D \perp A \mid Y^0$. That is, let $D = f(S, U)$ represent the function f describing the decision process after the RAI S is introduced. If cEO is satisfied and $U \perp (S, A) \mid Y^0$, then it follows that $D \perp A \mid Y^0$. Even if $U \not\perp (S, A) \mid Y^0$, it is easy to see that if the conditional independence statement *nearly* holds, or if f depends primarily on S rather than U , then discriminatory disparate impact can be small.

No such guarantees hold for predictors satisfying either independence or sufficiency. Chouldechova [9] in particular showed how predictors which satisfy sufficiency (predictive parity) are likely to yield decisions such that $D \not\perp A \mid Y$; these arguments are unchanged when we substitute Y^0 for Y . Though there is no consensus about how to quantify fairness, this is at least one consideration in favor of equalized odds.

4 MOTIVATING EXAMPLE

Having motivated equalized odds over predictive parity or independence, we now motivate the use of counterfactual rather than observational equalized odds.

Consider a school district that assigns tutors to students who are believed to be at risk of academic failure. The school district wishes to develop a RAI, S , to better identify students who need tutors while ensuring that this resource is allocated fairly across two levels of the sensitive feature A . Let $D \in \{0, 1\}$ represent the decision to assign (1) or not assign (0) a tutor, and let $Y \in \{0, 1\}$ represent academic success (0) or failure (1).

A cEO predictor S satisfies $\mathbb{P}(S \mid Y^0, A) = \mathbb{P}(S \mid Y^0)$, while an oEO predictor S satisfies $\mathbb{P}(S \mid Y, A) = \mathbb{P}(S \mid Y)$. Divergence in these predictors is driven by the extent to which $Y \neq Y^0$ in the training data. In order to parameterize this divergence, we introduce the following definitions.

Definition 4.1. The *need rate* for group a is $\mathbb{P}(Y^0 = 1 \mid A = a)$, the probability that a student from group a would fail without a tutor.

Definition 4.2. The *opportunity rate* for group a is $\mathbb{P}(D = 1 \mid Y^0 = 1, A = a)$, the probability that a student in group a who needs a tutor receives one.

Definition 4.3. The *intervention strength* for group a is $\mathbb{P}(Y^1 = 0 \mid Y^0 = 1, A = a)$, the probability that a student in group a who would fail without a tutor would succeed with a tutor.

We simulate a simple data generating process in which we allow the intervention strength to vary, while constraining it to be equal for the two groups. We fix all other parts of the distribution. In particular, we set $\mathbb{P}(A = 1) = 0.7$, set the need rates to 0.4 and 0.2 for groups 0 and 1, and set the opportunity rates to 0.6 and 0.4. We set the probabilities that a tutor is assigned when it is not needed to $\mathbb{P}(D = 1 \mid Y^0 = 0, A = 0) = 0.3$ and $\mathbb{P}(D = 1 \mid Y^0 = 0, A = 1) = 0.2$.

²Some authors use “disparate impact” to refer to the criterion $S \not\perp A$, i.e. *independence* [47].

This represents a scenario in which the minority group has greater need, perhaps due to socioeconomic factors or prior educational opportunities, and also is likelier than the majority group to receive resources at baseline (prior to the development of the RAI). Finally, we set $\mathbb{P}(Y^1 = 0 \mid Y^0 = 0) = 1$, meaning that tutoring never increases the risk of failure.

We consider a hypothetical oEO predictor S with fixed false positive rate $\mathbb{P}(S = 1 \mid Y = 0, A) = \mathbb{P}(S = 1 \mid Y = 0) = 0.1$ and false negative rate $\mathbb{P}(S = 0 \mid Y = 1, A) = \mathbb{P}(S = 0 \mid Y) = 0.2$. We assume $S \perp\!\!\!\perp Y^0 \mid A, Y$, as would be the case for example when S is a high quality predictor of Y . Figure 1 shows the cTPRs for this predictor as a function of intervention strength, relative to the baseline opportunity rates for the two groups. When the intervention has no effect (strength 0), the cTPRs are equal because $Y \equiv Y^0$, so the cTPR and TPR are identical. (Of course, a strength of 0 means the tutoring is worthless.) For all strength values > 0 , the cTPR of the minority group is lower than for the majority group. The difference in error rates increases as intervention strength increases. A cEO predictor avoids this problem by design: the cTPRs for the two groups are constrained to be equal.

This example makes it clear that oEO predictors in general will not prevent discriminatory disparate impact, whereas, as discussed in section 3.4, counterfactual EO predictors have at least the potential to mitigate or avoid it.

This example also illustrates how oEO predictors can reduce rates of appropriate intervention. For example, suppose that decision makers, after the introduction of the RAI, set $D \equiv S$, i.e. they assign tutors precisely to students whom the RAI labels as high risk. Then, for any intervention strength > 0.5 , the opportunity rate for the minority group decreases below baseline: the RAI *harms the minority group*.

5 AN OPTIMAL FAIR DERIVED PREDICTOR

Having motivated counterfactual equalized odds, we now develop a method to generate predictors which satisfy it.

5.1 Estimand

We expand our notation in order to fully describe our problem setting. Consider a random vector $Z = (A, X, D, S, Y) \sim \mathbb{P}$, where in addition to the binary sensitive feature A , decision D , and outcome Y , we have covariates $X \in \mathbb{R}^p$ and a previously trained binary predictor $S \in \{0, 1\}$. We require only that S is observable; we do not require access to its inputs or internal structure. S in practice could represent a RAI that is already in use, such as a recidivism prediction tool. The covariates X may or may not overlap with the inputs to S . Their role in the analysis is to render counterfactual quantities identifiable.

Our target is a derived predictor that satisfies approximate cEO. As in the case of observable equalized odds considered by [16], we achieve this by randomly flipping S with probabilities that depend only on S and A . Consider a column vector $\theta = (\theta_{0,0}, \theta_{0,1}, \theta_{1,0}, \theta_{1,1}) \in [0, 1]^4$. We define an associated derived predictor S_θ :

$$S_\theta \sim \text{Bern}(\theta_{A,S})$$

$$\text{where } \theta_{A,S} = \sum_{a,s \in \{0,1\}} \mathbb{1}\{A = a, S = s\} \theta_{a,s}$$

Figure 1: Counterfactual true positive rates (cTPRs; solid lines) for a RAI satisfying observational equalized odds (oEO), as a function of the intervention strength $\mathbb{P}(Y^1 = 0 \mid Y^0 = 1)$. Dashed lines indicate opportunity rates $\mathbb{P}(D = 1 \mid Y^0 = 0)$ prior to the development of the RAI. The more effective the tutoring (the higher the intervention strength), the worse the RAI is at identifying students who need it, and the greater the disparity in its performance between the minority and the majority group. When tutoring is more effective, the RAI may reduce the appropriate assignment of tutors below the baseline opportunity rates.

In other words, the $\theta_{a,0}$ parameters represent conditional probabilities that S flips, while the $\theta_{a,1}$ parameters represent conditional probabilities that S doesn't flip. Notice that for $\hat{\theta} = (0, 1, 0, 1)$, we have $S_{\hat{\theta}} = S$: the derived predictor is equal to the input predictor.

Our target is a loss-optimal fair predictor S_{θ^*} , where the fairness criterion is approximate cEO. The loss function we consider is weighted classification error. For fixed θ , denote the loss³ by $\mathcal{L}(S_\theta; w^+, w^-) = w^+ \mathbb{P}(Y^0 = 0, S_\theta = 1) + w^- \mathbb{P}(Y^0 = 1, S_\theta = 0)$, where w^+, w^- are chosen by the user to capture the relative importance of false positives and false negatives. (Note that we will generally suppress the dependence of $\mathcal{L}$ on w^+, w^- .) The estimand is

$$\begin{aligned} \theta^* &\in \arg \min_{\theta} \mathcal{L}(S_\theta) \\ &\text{subject to } \theta \in [0, 1]^4 \\ &|\Delta^+(S_\theta)| \leq \epsilon^+ \\ &|\Delta^-(S_\theta)| \leq \epsilon^- \end{aligned}$$

where Δ^+, Δ^- are given above in Definition 2.2, and the fairness constraints $\epsilon^+, \epsilon^- \in [0, 1]$ are chosen by the user. Setting both these constraint parameters to 0 requires cEO to be satisfied exactly, while setting them to 1 allows S_{θ^*} to be arbitrarily unfair. Setting ϵ^+ to 0 regardless of ϵ^- forces S_{θ^*} to satisfy *counterfactual equal opportunity*; see [16] for the observational definition of this criterion.

³We refer to this quantity as “loss” instead of the conventional “risk” in order to avoid confusion between risk assessment and the error rate of a predictor.

Remark 1. Note that the full vector Z is required only to estimate the parameter θ^* that defines the optimal fair derived predictor. Once θ^* has been estimated, the resulting derived predictor requires access at runtime only to the sensitive feature A and the input predictor S .

Since our estimands involve counterfactual quantities, distributional assumptions are required in order to equate them to observable quantities.

5.2 Identification

In this subsection we show that the counterfactual error rates and loss can be identified under standard causal inference assumptions. All the quantities to be identified can be written in terms of the loss and the counterfactual error rates of the input predictor S . For ease of notation, we first define two nuisance parameters that appear in the estimand and associated estimators, namely the outcome regression and propensity score function. We generally drop the arguments from these functions in subsequent usage for the sake of conciseness.

$$\begin{aligned}\mu_0(A, X, S) &= \mathbb{E}[Y \mid A, X, S, D = 0] \\ \pi(A, X, S) &= \mathbb{P}(D = 1 \mid A, X, S)\end{aligned}$$

We make the following standard “no unmeasured confounding”-type causal inference assumptions:

1. (Consistency) $Y = DY^1 + (1 - D)Y^0$
2. (Positivity) $\exists \delta \in (0, 1)$ s.t. $\mathbb{P}(\pi(A, X, S) \leq 1 - \delta) = 1$
3. (Ignorability) $Y^0 \perp\!\!\!\perp D \mid A, X, S$

The consistency assumption means that the outcome observed for each individual is precisely the potential outcome corresponding to the treatment received. This implies that one person’s treatment assignment does not affect another person’s outcomes, meaning, for example, that an individual’s recidivism behavior does not depend on whether other individuals are detained or released. The positivity or *overlap* assumption requires that within strata of (A, X, S) of measure > 0 , individuals have some chance of receiving no intervention. Finally, the ignorability or *no unmeasured confounding* assumption requires that within strata of (A, X, S) , the treatment D is essentially random with respect to Y^0 . Satisfying ignorability assumptions typically requires collecting a rich enough set of deconfounding covariates. In the present case, even if X is low dimensional, the ignorability assumption is plausible if the input predictor S substantially drives decision making, or if it happens to be an accurate (if not necessarily fair) predictor of Y^0 .

Before giving the identifying expressions for the loss $\mathcal{L}(S_\theta)$ and the error rate differences Δ^+ , Δ^- , we give identifying expressions for the error rates of the input predictor S , which themselves appear in the expressions for Δ^+ , Δ^- .

Proposition 1. Under assumptions 1-3, the counterfactual error rates of the input predictor S are identified as follows:

$$\begin{aligned}\text{cFPR}(S, a) &= \frac{\mathbb{E}[S(1 - \mu_0)\mathbb{1}\{A = a\}]}{\mathbb{E}[(1 - \mu_0)\mathbb{1}\{A = a\}]} \\ \text{cFNR}(S, a) &= \frac{\mathbb{E}[(1 - S)\mu_0\mathbb{1}\{A = a\}]}{\mathbb{E}[\mu_0\mathbb{1}\{A = a\}]}\end{aligned}$$

Proofs of propositions are given in Appendix A. We now define several quantities that appear in the identifying expressions for $\mathcal{L}(S_\theta)$, Δ^+ , and Δ^- :

$$\begin{aligned}\beta_{a,s} &= \mathbb{E}[\mathbb{1}\{A = a, S = s\} (w^+ - (w^+ + w^-)\mu_0)], \text{ for } a, s \in \{0, 1\} \\ \beta &= (\beta_{0,0}, \beta_{0,1}, \beta_{1,0}, \beta_{1,1}) \\ \beta^+ &= (1 - \text{cFPR}(S, 0), \text{cFPR}(S, 0), \text{cFPR}(S, 1) - 1, -\text{cFPR}(S, 1)) \\ \beta^- &= (-\text{cFNR}(S, 0), \text{cFNR}(S, 0) - 1, \text{cFNR}(S, 1), 1 - \text{cFNR}(S, 1))\end{aligned}$$

Proposition 2. Under assumptions 1-3, the loss and error rates of the derived predictor S_θ are identified as:

$$\begin{aligned}\mathcal{L}(S_\theta) &= \theta^T \beta + w^- \mathbb{E}[\mu_0] \\ \Delta^+(S_\theta) &= \theta^T \beta^+ \\ \Delta^-(S_\theta) &= \theta^T \beta^-\end{aligned}$$

Since the term $w^- \mathbb{E}[\mu_0]$ in the loss is fixed, we can drop it without changing the minimizer of the loss. We can therefore rewrite the estimand as

$$\begin{aligned}\theta^* &\in \arg \min_{\theta} \theta^T \beta \\ \text{subject to } \theta &\in [0, 1]^4 \\ |\theta^T \beta^+| &\leq \epsilon^+ \\ |\theta^T \beta^-| &\leq \epsilon^-\end{aligned} \tag{1}$$

In other words, the optimal fair derived predictor is the solution to a linear program (LP). We refer to this as the “true LP” since it defines the estimand. We now define an estimator $\hat{\theta}$ as the solution to an “estimated LP.”

5.3 Estimation

An estimator for θ^* is derived by computing estimates $\hat{\beta}, \hat{\beta}^+, \hat{\beta}^-$ of the true LP coefficients and then solving the resulting estimated LP:

$$\begin{aligned}\hat{\theta} &\in \arg \min_{\theta} \theta^T \hat{\beta} \\ \text{subject to } \theta &\in [0, 1]^4 \\ |\theta^T \hat{\beta}^+| &\leq \epsilon^+ \\ |\theta^T \hat{\beta}^-| &\leq \epsilon^-\end{aligned}$$

Any solution $\hat{\theta}$ suffices. There are many possible estimators of the true LP coefficients. Theorems 1 and 2 below show that the rates at which these coefficients are estimated propagate to the loss and fairness properties of $S_{\hat{\theta}}$. Arguably the simplest estimators involve plugging in an estimate of the regression function μ_0 . We propose instead using doubly robust estimators, which yield faster rates of convergence than plugin estimators in general nonparametric settings [41, 43].

For ease of notation, let

$$\phi = \frac{1 - D}{1 - \pi} (Y - \mu_0) + \mu_0$$

denote the uncentered efficient influence function for $\mathbb{E}(Y^0)$, and let $\hat{\phi}$ denote an estimate constructed from estimates $\hat{\mu}_0$ and $\hat{\pi}$ [4, 15, 23, 42]. To minimize the use of indices, let $\mathbb{P}_n(f(Z)) = n^{-1} \sum_{i=1}^n f(Z_i)$

denote the sample average of any function of Z . The doubly robust estimators for individual coefficients are:

$$\widehat{\beta}_{a,s} = \mathbb{P}_n[\mathbb{1}\{A = a, S = s\}(\mathbf{w}^+ - (\mathbf{w}^+ + \mathbf{w}^-)\widehat{\phi})] \quad (2)$$

$$\widehat{\text{cFPR}}(S, a) = \frac{\mathbb{P}_n[\mathbb{1}\{A = a\}S(1 - \widehat{\phi})]}{\mathbb{P}_n[\mathbb{1}\{A = a\}(1 - \widehat{\phi})]} \quad (3)$$

$$\widehat{\text{cFNR}}(S, a) = \frac{\mathbb{P}_n[\mathbb{1}\{A = a\}(1 - S)\widehat{\phi}]}{\mathbb{P}_n[\mathbb{1}\{A = a\}\widehat{\phi}]} \quad (4)$$

These estimates are assembled into the corresponding vectors $\widehat{\beta}, \widehat{\beta}^+, \widehat{\beta}^-$.

In order to obtain optimal convergence rates, it is generally necessary to estimate the nuisance functions $\widehat{\mu}_0$ and $\widehat{\pi}$ on one sample and then compute the sample mean $\mathbb{P}_n$ on an independent sample conditional on those estimates. To obtain full sample size efficiency, one can swap the folds, repeat the procedure, and average the results, an approach that is popularly called cross-fitting [5, 8, 38, 49]. A k -fold version of cross-fitting with $k > 2$ is also possible. If $\widehat{\mu}_0$ and $\widehat{\pi}$ are assumed to be sufficiently “well-behaved,” i.e. if they belong to Donsker classes, then no such sample splitting is necessary. We prefer to avoid this assumption and utilize sample splitting. See Appendix C for a schematic of the sample splitting procedure.

We now show that $S_{\widehat{\theta}}$ approaches optimal behavior at rates that depend on the rates at which the LP coefficient vectors are estimated⁴. We define two quantities of interest: the *loss gap* and the *excess unfairness*, and give accompanying theorems.

Following standard usage, we say that an estimator $\widehat{\psi}$ of a parameter ψ is consistent at rate $f(n)$ for some real-valued function $f(n)$ if $\|\widehat{\psi} - \psi\| = O_{\mathbb{P}}(1/f(n))$ for a suitable norm $\|\cdot\|$. For example, if $f(n) = \sqrt{n}$, then we say $\widehat{\psi}$ converges at $\sqrt{n}$ rates. We say that an estimator converges *faster* than $f(n)$ if $\|\widehat{\psi} - \psi\| = o_{\mathbb{P}}(1/f(n))$.

Definition 5.1. The *loss gap* is $\mathcal{L}(S_{\widehat{\theta}}) - \mathcal{L}(S_{\theta^*})$, the difference in loss between the derived predictor and the optimal derived predictor.

We use the term *loss gap* rather than *excess loss* to acknowledge that the loss of $S_{\widehat{\theta}}$ can be less than the loss of S_{θ^*} , if $\widehat{\theta}$ falls outside the true constraints. Of course, this can only occur if $S_{\widehat{\theta}}$ violates the true fairness constraints, which can happen because the constraints are estimated.

Theorem 1. (Loss gap.) Suppose that $\widehat{\beta}, \widehat{\beta}^+$, and $\widehat{\beta}^-$ are all consistent at rate $f(n)$. Under Assumptions 1-3:

$$\mathcal{L}(S_{\widehat{\theta}}) - \mathcal{L}(S_{\theta^*}) = O_{\mathbb{P}}(1/f(n))$$

Definition 5.2. The *excess unfairness* of $S_{\widehat{\theta}}$ in the cFPR is

$$\text{UF}^+(S_{\widehat{\theta}}) := \max\{|\text{cFPR}(S_{\widehat{\theta}}, 0) - \text{cFPR}(S_{\widehat{\theta}}, 1)| - \epsilon^+, 0\},$$

and the *excess unfairness* of $S_{\widehat{\theta}}$ in the cFNR is

$$\text{UF}^-(S_{\widehat{\theta}}) := \max\{|\text{cFNR}(S_{\widehat{\theta}}, 0) - \text{cFNR}(S_{\widehat{\theta}}, 1)| - \epsilon^-, 0\}$$

Since the estimated constraints should fluctuate around the true constraints, it’s possible for $S_{\widehat{\theta}}$ to have error rate differences that are smaller than ϵ^+, ϵ^- , which motivates bounding these quantities below by 0.

⁴We ignore optimization error, since this is a function of the number of optimization iterations and can be made arbitrarily small [6].

Theorem 2. (Excess unfairness.) Suppose that $\widehat{\beta}, \widehat{\beta}^+$, and $\widehat{\beta}^-$ are all consistent at rate $f(n)$. Under assumptions 1-3:

$$\max\{\text{UF}^+(S_{\widehat{\theta}}), \text{UF}^-(S_{\widehat{\theta}})\} = O_{\mathbb{P}}(1/f(n))$$

Remark 2. (The behavior of $\widehat{\theta}$ vs. $S_{\widehat{\theta}}$.) Without assumptions about how the loss and fairness of $S_{\widehat{\theta}}$ depend on $\widehat{\theta}$, there is no guarantee about the rate at which $\widehat{\theta}$ will approach θ^* . This is not a concern, however, since the object of interest is not θ^* per se but a predictor that behaves like S_{θ^*} .

5.4 Estimating performance of the derived predictor

Once $\widehat{\theta}$ has been computed, it is of interest to check both the loss $\mathcal{L}(S_{\widehat{\theta}})$ and the error rate differences $\Delta^+(S_{\widehat{\theta}})$, $\Delta^-(S_{\widehat{\theta}})$ of the resulting derived predictor $S_{\widehat{\theta}}$, for example to understand the performance “cost” of fairness and to check whether the procedure successfully controlled the error rate differences.

These estimates should be computed on a test set that is independent of the sample used to estimate $\widehat{\theta}$. Within the test set, the same sample splitting considerations apply: unless they are assumed to belong to Donsker classes, the nuisance parameters $\widehat{\mu}_0$ and $\widehat{\pi}$ should be estimated on separate folds from the folds used to compute the relevant sample means $\mathbb{P}_n$.

Since the estimators below are conditional on a fixed $\widehat{\theta}$, they can in fact be applied to any fixed parameter value $\theta \in [0, 1]^4$. We define one additional quantity of interest.

Definition 5.3. The *loss change* for a derived predictor S_{θ} relative to an input predictor S is $\Gamma(S_{\theta}) = \mathcal{L}(S_{\theta}) - \mathcal{L}(S)$.

We refer to a *loss change* rather than an *increase in loss* because it is possible for S_{θ} to have smaller loss than S . This is not a typical expectation: in fair prediction problems, the set of fair classifiers is necessarily smaller than the set of fair and unfair classifiers, so there is a fairness-accuracy tradeoff. In the RAI setting, however, since predictors are typically trained to predict observable outcomes, their performance may be arbitrarily bad with respect to the potential outcome Y^0 . It is therefore not implausible than a derived fair predictor could have higher accuracy than the input predictor.

Once again, we propose using doubly robust estimators. These estimators are essentially identical to the estimators of the LP coefficients used in the previous section.

$$\widehat{\mathcal{L}}(S_{\theta}) = \theta^T \widehat{\beta} + \mathbf{w}^- \mathbb{P}_n(\widehat{\phi}) \quad (\text{loss})$$

$$\widehat{\Gamma}(S_{\theta}) = (\theta - \tilde{\theta})^T \widehat{\beta} \quad (\text{loss change})$$

$$\widehat{\text{cFPR}}(S_{\theta}, a) = \theta_{a,0} \left(1 - \widehat{\text{cFPR}}(S, a)\right) + \theta_{a,1} \widehat{\text{cFPR}}(S, a) \quad (\text{cFPR})$$

$$\widehat{\text{cFNR}}(S_{\theta}, a) = (1 - \theta_{a,0}) \left(\widehat{\text{cFNR}}(S, a)\right) + \theta_{a,1} \left(1 - \widehat{\text{cFNR}}(S, a)\right) \quad (\text{cFNR})$$

$$\widehat{\Delta}^+(S_{\theta}) = \theta^T \widehat{\beta}^+ \quad (\text{error rate difference in cFPR})$$

$$\widehat{\Delta}^-(S_{\theta}) = \theta^T \widehat{\beta}^- \quad (\text{error rates difference in cFNR})$$

where recall $\tilde{\theta} = (0, 1, 0, 1)$, so that $S_{\tilde{\theta}} = S$ (i.e., the derived predictor is simply the input predictor). Note that the loss estimator adds back in the portion of the loss that doesn’t depend on θ and that

we consequently removed from the LP in (1). When the nuisance parameter estimators $\hat{\mu}_0$ and $\hat{\pi}$ are estimated at faster than $n^{1/4}$ rates, these doubly robust estimators are $\sqrt{n}$ -consistent and asymptotically normal, so we can use the central limit theorem to derive asymptotically valid confidence intervals.

6 RESULTS

There is no previous method designed to achieve counterfactual equalized odds or related fairness criteria that we can compare our method to. However, we compare our method to an approach that uses plugin estimators for the LP coefficients.

6.1 Simulations

We use one set of simulations to illustrate Theorems 1 and 2 and another set to explore fairness-performance tradeoffs. We use equal misclassification weights $w^+ = w^- = 1$, so that false positives and false negatives contribute equally to the loss.

Each estimation procedure was run 500 times for each sample size $n \in \{100, 200, 500, 1000, 5000, 20000\}$. Since μ_0 is known here, the “true” loss and fairness values were computed on a separate validation set of size 500,000, using plugin estimators with the true μ_0 . These values showed negligible variation over many repetitions.

6.1.1 Setup. First, we define a *pre-RAI* data generating process. Using this data, we train a predictor S to predict observable outcomes Y , mirroring how RAIs are typically constructed in practice. We then define a *post-RAI* data generating process, which only differs in that the predictor S now affects the decisions D . This emulates the way RAIs are intended to work in practice; for example, a criminal defendant labeled high-risk ($S = 1$) be a RAI might be less likely to be released pre-trial ($D = 0$) than they would have been prior to the introduction of the RAI. The data generating process is designed to meet assumptions 1-3, with $\pi(D \mid A, X, S)$ upper bounded at 0.975. It is described fully in Appendix D.

6.1.2 Theorems 1 and 2. To simulate the estimation of the LP coefficient vectors at a particular rate, we add random noise ϵ of magnitude $\sigma_{\mathbb{P}}(1/n^{1/4})$ to the nuisance parameters μ_0 and π^5 . Note that $\sqrt{n}$ rates for regression functions cannot be attained in general nonparametric settings. These slower rates *can* be attained in such settings, under relatively weak assumptions [43].

The principal advantage of doubly robust estimators over plugin estimators is that they converge at a rate that is like the product of the nuisance parameter rates. When $\hat{\mu}_0$ and $\hat{\pi}$ both converge faster than $n^{1/4}$, the estimators $\hat{\beta}^+, \hat{\beta}^-, \hat{\beta}^-$ all converge at $\sqrt{n}$ rates, which is the fastest rate that can be attained for these parameters in general nonparametric settings. Plugin versions of $\hat{\beta}^+, \hat{\beta}^-, \hat{\beta}^-$ would inherit the slow rates of $\hat{\mu}_0$.

Figure 2 shows the loss $\mathcal{L}(S_{\hat{\theta}})$ and the excess unfairness values $UF^+(S_{\hat{\theta}})$, $UF^-(S_{\hat{\theta}})$ for the post-processed predictor $S_{\hat{\theta}}$ with fairness constraints $\epsilon^+ = 0.10$, $\epsilon^- = 0.20$. As expected, when doubly robust estimators are used, the loss and excess unfairness values converge at $\sqrt{n}$ rates to the loss $\mathcal{L}(S_{\theta^*})$ of the optimal derived predictor and 0, respectively. When plugin estimators are used, the rates are slower than $\sqrt{n}$.

⁵The noise is added on the logit scale to ensure that $\hat{\mu}_0, \hat{\pi}$ remain in $[0, 1]$, and $\hat{\pi}$ is again truncated to 0.975.

Figure 2: (Illustration of Theorems 1 and 2). Loss $\mathcal{L}$ and excess unfairness values UF^+ , UF^- for the derived predictor $S_{\hat{\theta}}$ for samples of size 100 to 20,000. Each vertical line represents a mean ± 1 sd over 500 simulations. Orange horizontal lines represent the loss of the optimal derived predictor S_{θ^*} (top left panel) or 0. The top row represents our doubly robust (DR) procedure and shows that the loss and excess unfairness converge to their target values. The bottom two rows represent values from the DR procedure or a plugin (PI) procedure, transformed by $\psi(S_{\hat{\theta}}) \mapsto \sqrt{n}(\psi(S_{\hat{\theta}}) - \psi(S_{\theta^*}))$, where ψ is $\mathcal{L}$ or UF^+ or UF^- , as appropriate. These rows illustrate that $\sqrt{n}$ -convergence is only guaranteed for $\hat{\theta}_{DR}$: the scaled values for $\hat{\theta}_{DR}$ do not grow in n , while the scaled values for $\hat{\theta}_{PI}$ begin to diverge.

6.1.3 Fairness-performance tradeoffs. Figure 3 shows the loss change $\Gamma(S_{\theta^*}) = \mathcal{L}(S_{\theta^*}) - \mathcal{L}(S)$ for each point in a grid of fairness constraints ϵ^+ , ϵ^- . Here, S is the Bayes-optimal predictor of Y^0 in our data generating scenario, meaning $S(A, X) = \mathbb{E}[Y^0 \mid A, X]$. Since any derived predictor necessarily has greater loss than the Bayes-optimal predictor, we refer to the loss change here equivalently as the *performance cost*.

In the data generating process used in the previous section, the Bayes-optimal predictor has absolute error rate differences of only 0.05 (Δ^+) and 0.04 (Δ^-), which leaves little room to illustrate the potential cost of fairness. For these simulations, therefore, we alter the data generating process slightly. (See Appendix D). This results in a Bayes-optimal predictor with absolute error rate differences of 0.23 (Δ^+) and 0.40 (Δ^-) and a loss of 0.24, which are plausible values for a real predictor.

As expected, when $\epsilon^+ \geq \Delta^+(S)$ or $\epsilon^- \geq \Delta^-(S)$, the performance cost is 0: the input predictor already satisfies the fairness constraints, so our method simply returns the input predictor. As the tolerances tighten towards 0, the performance declines, though never substantially. For $\epsilon^+ = \epsilon^- = 0$, when the derived predictor is constrained to satisfy exact cEO, the loss increases by 0.10, to

Figure 3: (Fairness-performance tradeoffs). Loss change $\Gamma(\theta^*) = \mathcal{L}(S_{\theta^*}) - \mathcal{L}(S)$ for the Bayes-optimal input predictor $S(A, X) = \mathbb{E}[Y^0 | A, X]$ and θ^* corresponding to different fairness constraints ϵ^+ , ϵ^- . The black area represents fairness constraints that are looser than the error rate differences of the input predictor ($\Delta^+(S) = 0.24$, $\Delta^-(S) = 0.40$), which incur no performance cost. The highest performance cost (0.10) occurs when the error rates differences are both constrained to be 0, meaning the derived predictor S_{θ} satisfies cEO exactly.

0.34. The different values for $\Delta^+(S)$ and $\Delta^-(S)$ are reflected in the differing costs of satisfying fairness along the two axes: the cost of controlling $\Delta^+(S_{\theta})$ are lower than the costs of controlling $\Delta^-(S_{\theta})$.

[46] showed that post-processing can result in predictors with poor performance, but it is unclear how likely this is to be a problem in practice. While the fairness-accuracy tradeoff naturally depends on the data generating process, our example illustrates that fairness can in some cases be achieved without substantial performance costs.

6.2 COMPAS data

We illustrate our method on the COMPAS recidivism dataset gathered by ProPublica [2, 29]. COMPAS refers to a collection of tools designed to assess the risk of recidivism. The dataset comprises public arrest records, criminal records, and COMPAS RAI scores from Broward County, Florida, spanning 2013–2016. After filtering the data in the same manner as [29] and restricting to defendants who are labeled African-American ($A = 0$) or Caucasian ($A = 1$), we are left with data for 5278 individuals (3175 African-American, 2103 Caucasian).

We utilize the COMPAS scores for general, as opposed to violent, recidivism. The scores are given in risk deciles. Since our method operates on a binary input predictor S , we follow ProPublica and set scores of 1–4 to $S = 0$ (“low risk”) and scores of 5–10 to $S = 1$ (“high risk”). The outcome Y is recidivism within a two-year time period. (See [29] for how recidivism is operationalized.) ProPublica’s analysis focuses on the use of COMPAS to inform pretrial release decisions. The dataset includes dates in and out of jail but does not

indicate whether defendants were released pretrial, so we set the treatment D to 0 if defendants left jail within three days of being arrested, and 1 otherwise. This yields 3645 released individuals (2158 African-American, 1487 Caucasian) and 1633 incarcerated individuals (1017 African-American, 616 Caucasian). Note that this threshold is somewhat arbitrary. Florida state law generally requires individuals to be brought before a judge for a bail hearing within 48 hours of arrest, but it may take time for individuals to post bail if they are required and able to do so.

The covariates X consist of gender (coded male or female), age (coded categorically for < 25 , between 25 and 45, and > 45), the number of prior crimes, and charge degree (misdemeanor or felony). Without consulting with domain experts, it is difficult to assess the plausibility of the positivity and ignorability assumptions given these covariates. Hence we intend our analysis primarily to be illustrative of our method, and we resist drawing strong substantive conclusions about COMPAS.

We weight false positives and false negatives equally, i.e. we set $w^+ = w^- = 1$. We randomly split the data into training and test sets of equal size. For $\epsilon \in \{0, 0.01, 0.05, 0.10, 0.20, \dots, 0.90, 1\}$, we set the fairness constraints to $\epsilon^+ = \epsilon^- = \epsilon$, computing the corresponding estimate $\hat{\theta}$ on the training set and estimating properties of the post-processed predictor $S_{\hat{\theta}}$ on the test set. We also estimate properties of the binarized COMPAS score S on the test set. We use random forests to estimate both the propensity scores $\hat{\pi}$ and the outcome regression $\hat{\mu}_0$. To reduce the variance of the estimates, we employ 5-fold cross-fitting: within the train set, we compute five estimates $\hat{\theta}_j$, $j = 1, \dots, 5$, using four folds at a time to estimate the nuisance parameters and the held-out fold to compute $\hat{\theta}_j$. Then $\hat{\theta} := \frac{1}{5} \sum_j \hat{\theta}_j$. We utilize the test set in an analogous fashion for the remaining estimators.

We also estimate the “predictive change” $\mathbb{P}(S_{\hat{\theta}} \neq S)$, the proportion of predictions that the post-processed predictor flips. A point estimate and confidence interval for this quantity can be straightforwardly computed via a formula given in Appendix E.

Table 1 contains estimates and confidence intervals for COMPAS and for the post-processed predictor corresponding to fairness constraints of $\epsilon^+ = \epsilon^- = 0.05$. The loss for COMPAS is 0.36, and the differences in the cFPR and cFNR are -0.24 and 0.16, respectively. The signs of these differences are consistent with what ProPublica found in their analysis with respect to observable Y : the false positive rates are higher for African-American defendants, while the false negative rates are higher for Caucasian defendants. The post-processing procedure successfully shrinks these differences to -0.05 and -0.03, which fall within the target range of $[-0.05, 0.05]$. This reduction corresponds to flipping 9% of the COMPAS scores, and it incurs an increase in risk of only 0.03.

The value of $\hat{\theta}$ corresponding to $S_{\hat{\theta}}$ here is (0, 0.91, 0.23, 1). The 0 and the 1 indicate that $S_{\hat{\theta}}$ does not change the COMPAS scores for African-American defendants who receive a “low-risk” score or Caucasian defendants who receive a “high-risk” score. The scores for high-risk African-American defendants are flipped to low-risk $1 - 0.91 = 9\%$ of the time, while the scores for low-risk Caucasian defendants are flipped to high-risk 23% of the time. This has the effect of increasing the false positive rate and decreasing the false

Figure 4: Convergence of the estimated loss $\widehat{\mathcal{L}}(S_{\widehat{\theta}})$, predictive change $\widehat{\mathbb{P}}(S_{\widehat{\theta}} \neq S)$, and error rate differences $\widehat{\Delta}^+(S_{\widehat{\theta}}), \widehat{\Delta}^-(S_{\widehat{\theta}})$, for post-processed versions of the binarized COMPAS predictor. Fairness constraints are set to $\epsilon^+ = \epsilon^- = \epsilon$ over a range of values ϵ . Vertical lines are 95% CIs. Horizontal orange lines indicate the reference values for COMPAS, or 0 in the case of predictive change. The dashed blue lines $y = x$ and $y = -x$, mark the target fairness constraints.

negative rate for Caucasians, while moving the rates in the opposite directions for African-Americans.

Figure 4 shows the loss, error rate differences, and predictive change for fairness constraints ranging from 0 (requiring no gap in error rates) to 1 (imposing no fairness constraints). Each constraint induces an estimate $\widehat{\theta}$ and a corresponding post-processed predictor $S_{\widehat{\theta}}$. The estimated fairness gaps fall along or within the lines $y = \pm x$, indicating that each $S_{\widehat{\theta}}$ satisfies its target constraints. At the most stringent setting of 0, the loss for $S_{\widehat{\theta}}$ is approximately 0.40, which compares favorably with the estimated baseline loss of 0.36 for COMPAS. This $S_{\widehat{\theta}}$ flips slightly less than 20% of the scores.

For $\epsilon > 0.24$, the fairness constraints are essentially no longer active, since COMPAS itself satisfies these constraints. Indeed, as expected, the $\widehat{\theta}$ values for $\epsilon > 0.24$ are all essentially $[0, 1, 0, 1]$, meaning that $S_{\widehat{\theta}} = S$, and the estimated risk and fairness values all fall close to the estimated values for COMPAS. (There is still some variation in the estimated values due to randomness in the k-fold cross-fitting procedure.)

These results illustrate that our approach performs as intended on a real dataset: if these data were indeed generated from a distribution satisfying the identifying assumptions, then our post-processed predictor would satisfy approximate counterfactual equalized odds while incurring little cost in performance.

Table 1: Estimates and 95% confidence intervals for the loss $\mathcal{L}$, loss change Γ , error rates cFPR and cFNR for groups 0 and 1, error rate differences Δ^+, Δ^- , and predictive change $\mathbb{P}(S_{\widehat{\theta}} \neq S)$ for the binarized COMPAS predictor S and the post-processed predictor $S_{\widehat{\theta}}$, with ϵ^+ and ϵ^- set to 0.05.

	S	$S_{\widehat{\theta}}$
$\widehat{\mathcal{L}}(\cdot)$	0.36 (0.32, 0.41)	0.39 (0.35, 0.42)
$\widehat{\Gamma}(\cdot)$	–	0.03 (0.01, 0.04)
$\widehat{\text{cFPR}}(\cdot, 0)$	0.43 (0.36, 0.49)	0.39 (0.33, 0.45)
$\widehat{\text{cFPR}}(\cdot, 1)$	0.24 (0.18, 0.31)	0.42 (0.37, 0.47)
$\widehat{\text{cFNR}}(\cdot, 0)$	0.30 (0.25, 0.35)	0.36 (0.31, 0.40)
$\widehat{\text{cFNR}}(\cdot, 1)$	0.53 (0.46, 0.60)	0.41 (0.35, 0.46)
$\widehat{\Delta}^+(\cdot)$	-0.24 (-0.32, -0.15)	-0.05 (-0.12, 0.02)
$\widehat{\Delta}^-(\cdot)$	0.18 (0.09, 0.28)	-0.03 (-0.10, 0.05)
$\widehat{\mathbb{P}}(\cdot \neq S)$	–	0.09 (0.09, 0.09)

6.3 Child welfare data

Cost-sensitive loss functions can drive $\widehat{\theta}$ to a trivial classifier that always predicts one class. We illustrate this phenomenon on a dataset representing calls to a child-welfare hotline in Allegheny County, Pennsylvania. The data comprises over 30,000 calls and contains over 1,000 features. The features describe allegations made in the call, assessments of risk made by hotline workers, and features pertaining to individuals associated with the call. Workers must decide whether to *screen in* a call, which means opening an investigation into the allegations. The baseline decision $D = 0$ is to screen out, meaning no investigation takes place. The outcome Y is re-referral to the hotline within a six month period. For further details about the child welfare setting and this dataset in particular, see [10, 11].

Unlike the COMPAS dataset, this dataset does not include a previously trained predictor. We therefore first build a predictor S that predicts Y^0 , and then we post-process S . In this setting, we have reason to believe that the identification assumptions in section 5.2 are plausible, once cases with the highest propensity for screen-in are removed; see [11]. (RAIs are not necessary or useful for cases that are already guaranteed to be screened in.) In order to accomplish this filtering, we first build a propensity score model using random forests on roughly one third of the data. The model appears well-calibrated, so we filter out the approximately 20% of the cases with estimated propensity scores greater than 0.99. Note that downstream results did not change substantially when these cases were left in.

We then train a classification random forest S to predict Y conditional on $A, X, D = 0$, using the same third of the data. Under the identifying assumptions, $Y|X, A, D = 0$ is equal in distribution to $Y^0|A, X$, so S is indeed an estimate of the target Y^0 . Following recommended usage in this setting, we set the classification threshold to capture the top 25% riskiest cases [10].

The predictor S has estimated error rate differences and 95% confidence intervals of $\widehat{\Delta}^+ = -0.02 \pm 0.01$ and $\widehat{\Delta}^- = 0.09 \pm 0.08$. It is unsurprising that these differences are small, given that rereferral rates are similar for Black (0.24) and White (0.27) cases. (See [9] for

Figure 5: Cost-sensitive post-processing for the child welfare predictor over a range of cost ratios, with fairness constraints $\epsilon^+ = \epsilon^- = 0.01$. Each column represents a single $\hat{\theta}$, with the four components $\theta_{a,s}$ in rows. False positives are weighted between 1.25 and 3 times as heavily as false negatives to the left of the dashed line, and vice versa to the right. Extreme cost ratios push the post-processed classifier to a trivial classifier that always predicts 0 (to the left of the orange lines) or 1 (to the right). Between these, post-processing essentially returns the input predictor.

an examination of the relationship between base rates and error rates.)

In order to have nontrivial (active) fairness constraints, we set $\epsilon^+ = \epsilon^- = 0.01$. Figure 5 shows the value of $\hat{\theta}$ over a range of cost ratios w^+/w^- and w^-/w^+ . When false positives are weighted more than 1.5 times as heavily as false negatives, post-processing returns classifiers that are very close to the simple majority classifier $S_{(0,0,0,0)} \equiv 0$. When false negatives are weighted more than 2 times as heavily as false positives, post-processing returns the simple minority classifier $S_{(1,1,1,1)} \equiv 1$. Since the input classifier is approximately fair, between those ranges, post-processing returns classifiers that are very close to the input classifier $S = S_{(0,1,0,1)}$, with only the fourth component $\hat{\theta}_{1,1}$ deviating slightly from 1.

This behavior is expected. Note that a simple majority or minority classifier always satisfies counterfactual equalized odds, since the error rate differences are 0. Since the post-processed predictor only has access at runtime to two binary features, as either false positives or false negatives become sufficiently important, one of these simple classifiers will at some point become the lowest risk option. This is possible in principle when w^+ and w^- are equal, but it is guaranteed as their ratio grows.

Since this dataset did not include a pretrained predictor of Y^0 , it would be preferable to adopt an in-processing approach, i.e. to train a predictor that satisfies the desired fairness constraints in a single stage, rather than training an unconstrained predictor and then post-processing it. We pursue this task in ongoing work.

7 DISCUSSION AND CONCLUSION

In this paper we considered fairness in risk assessment instruments (RAIs), which are naturally concerned with potential outcomes rather than strictly observable outcomes. We defined the fairness criterion *approximate counterfactual equalized odds* (approximate cEO), which allows users to negotiate the tradeoff between fairness and performance. We argued that this fairness criterion is likelier than other candidate criteria to reduce discriminatory disparate impact, which we defined as $D \not\perp A \mid Y^0$.

We presented a method to post-process an existing binary predictor to satisfy approximate cEO using doubly robust estimators, and we showed that our method has favorable convergence properties. Our rate results translate readily to the post-processing setting of [16], in which the outcome of interest is the observable Y and the fairness criterion is (approximate) observational equalized odds.

Once it is constructed, the post-processed predictor requires access at runtime only to the sensitive feature and the input predictor, making it relatively feasible to implement on top of existing RAIs. A predictor trained from scratch would be constrained by the set of covariates available in deployment, whereas the post-processing approach allows researchers to devise a set of suitable deconfounding covariates and then collect an appropriate dataset on a one-time basis.

In closing, we note that from our perspective, notions of fairness in predictive systems ought to be subordinate to notions of fairness grounded in the actual decisions or events that those systems inform, and the impact that those decisions have on people's lives. Though little is currently known about how decision makers respond to RAIs, there is some evidence that judges do not have much faith in recidivism predictions and that RAIs can have little impact on decisions [20, 40]. As RAIs and the general public's understanding of how they function co-evolve, it is likely that the ways in which decision makers respond to them will evolve as well.

Nevertheless, it seems plausible that some fairness criteria for RAIs are likelier than others to lead to increased (un)fairness with respect to decisions and outcomes. While this is ultimately an empirical question, we believe that this kind of consideration ought to ground discussions of fairness in RAIs and predictive systems generally. As long as there are domains involving high stakes decisions that we do not wish to fully automate, RAIs will remain relevant, and so will the task of ensuring that they lead to a society that is more fair, not less.

ACKNOWLEDGMENTS

Edward Kennedy gratefully acknowledges support from NSF DMS Grant 1810979. Additionally, we are grateful to the Block Center for Technology and Society at Carnegie Mellon University for funding this work, and to the Allegheny County Department of Human Services for furnishing the child welfare data.

REFERENCES

- [1] Julia Angwin and Jeff Larson. 2016. Bias in Criminal Risk Scores Is Mathematically Inevitable, Researchers Say. *ProPublica* (Dec 2016). <https://www.propublica.org/article/bias-in-criminal-risk-scores-is-mathematically-inevitable-researchers-say>
- [2] Julia Angwin, Jeff Larson, Surya Mattu, and Lauren Kirchner. 2016. Machine Bias. *ProPublica* (May 2016). <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>
- [3] Solon Barocas, Moritz Hardt, and Arvind Narayanan. 2018. Fairness and Machine Learning. (2018). <http://www.fairmlbook.org>
- [4] Peter J. Bickel, Chris A.J. Klaassen Ya'acov Ritov, and Jon A. Wellner. 1993. *Efficient and adaptive estimation for semiparametric models*. Johns Hopkins University Press, Baltimore.
- [5] Peter J. Bickel and Ya'acov Ritov. 1988. Estimating Integrated Squared Density Derivatives : Sharp Best Order of Convergence Estimates. *Sankhyā: The Indian Journal of Statistics, Series A* 50, 3 (1988), 381–393. <https://www.jstor.org/stable/25050710>
- [6] Stephen P. Boyd and Lieven Vandenbergh. 2004. *Convex optimization*. Cambridge University Press.
- [7] Flavio Calmon, Dennis Wei, Bhanukiran Vinzamuri, Karthikeyan Natesan Ramamurthy, and Kush R Varshney. 2017. Optimized Pre-Processing for Discrimination Prevention. In *Advances in Neural Information Processing Systems* (Long Beach, CA) (NIPS 2017), I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (Eds.). Curran Associates, Inc., 3992–4001. <http://papers.nips.cc/paper/6988-optimized-pre-processing-for-discrimination-prevention.pdf>
- [8] Victor Chernozhukov, Denis Chetverikov, Mert Demirer, Esther Duflo, Christian Hansen, Whitney Newey, and James Robins. 2018. Double/debiased machine learning for treatment and structural parameters. *The Econometrics Journal* 21, 1 (Feb. 2018), C1–C68. <https://doi.org/10.1111/ectj.12097>
- [9] Alexandra Chouldechova. 2017. Fair prediction with disparate impact: A study of bias in recidivism prediction instruments. *Big data* 5, 2 (2017), 153–163.
- [10] Alexandra Chouldechova, Diana Benavides-Prado, Oleksandr Fialko, and Rhema Vaithianathan. 2018. A case study of algorithm-assisted decision making in child maltreatment hotline screening decisions. In *Proceedings of the 1st Conference on Fairness, Accountability and Transparency (Proceedings of Machine Learning Research, Vol. 81)*, Sorelle A. Friedler and Christo Wilson (Eds.). PMLR, New York, NY, USA, 134–148. <http://proceedings.mlr.press/v81/chouldechova18a.html>
- [11] Amanda Coston, Alan Mishler, Edward H. Kennedy, and Alexandra Chouldechova. 2020. Counterfactual Risk Assessments, Evaluation, and Fairness (FAT* '20). Association for Computing Machinery, New York, NY, USA, 582–593. <https://doi.org/10.1145/3351095.3372851>
- [12] William Dieterich, Christina Mendoza, and Tim Brennan. 2016. *COMPAS Risk Scales: Demonstrating Accuracy Equity and Predictive Parity*. Technical Report.
- [13] Michele Donini, Luca Oneto, Shai Ben-David, John S Shawe-Taylor, and Massimiliano Pontil. 2018. Empirical Risk Minimization Under Fairness Constraints. In *Advances in Neural Information Processing Systems 31* (Montréal, Canada) (NeurIPS 2018), S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett (Eds.). Curran Associates, Inc., 2791–2801. <http://papers.nips.cc/paper/7544-empirical-risk-minimization-under-fairness-constraints.pdf>
- [14] Clark Glymour and Madelyn R. Glymour. 2014. Commentary: Race and sex are causes. *Epidemiology* 25, 4 (2014), 488–490. <https://doi.org/10.1097/EDE.0000000000000122>
- [15] Jinyong Hahn. 1998. On the Role of the Propensity Score in Efficient Semiparametric Estimation of Average Treatment Effects. *Econometrica* 66, 2 (1998), 315–331. <https://www.jstor.org/stable/2998560>
- [16] Moritz Hardt, Eric Price, Eric Price, and Nati Srebro. 2016. Equality of Opportunity in Supervised Learning. In *Advances in Neural Information Processing Systems 29* (NIPS 2016), D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett (Eds.). Curran Associates, Inc., 3315–3323. <http://papers.nips.cc/paper/6374-equality-of-opportunity-in-supervised-learning.pdf>
- [17] Paul W. Holland. 1986. Statistics and Causal Inference. *J. Amer. Statist. Assoc.* 81, 396 (1986), 968. <http://www.jstor.org/stable/2289069>
- [18] Lily Hu and Issa Kohler-Hausmann. 2020. What's Sex Got to Do with Machine Learning?. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency* (Barcelona, Spain) (FAT* '20). Association for Computing Machinery, New York, NY, USA, 513. <https://doi.org/10.1145/3351095.3375674>
- [19] Kosuke Imai and Zhichao Jiang. 2020. Principal Fairness for Human and Algorithmic Decision-Making. (2020). arXiv:2006.01770 [cs.CY] <https://arxiv.org/abs/2006.01770>
- [20] Melissa Jonnson. 2018. *The Influence of Risk Assessment Evidence on Judicial Sentencing Decisions*. Master's thesis. <http://summit.sfu.ca/item/18704>
- [21] Faisal Kamiran and Toon Calders. 2012. Data preprocessing techniques for classification without discrimination. *Knowledge and Information Systems* 33, 1 (2012), 1–33. <http://doi.org/10.1007/s10115-011-0463-8>
- [22] Michael Kearns, Aaron Roth, and Zhiwei Steven Wu. 2017. Meritocratic Fairness for Cross-Population Selection. In *Proceedings of the 34th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 70)*, Doina Precup and Yee Whye Teh (Eds.). PMLR, International Convention Centre, Sydney, Australia, 1828–1836. <http://proceedings.mlr.press/v70/kearns17a.html>
- [23] Edward H. Kennedy. 2016. Semiparametric Theory and Empirical Processes in Causal Inference. In *Statistical Causal Inference and Their Applications in Public Health Research*, Hua He, Pan Wu, and Ding-Geng (Din) Chen (Eds.). Springer, 141–167. https://doi.org/10.1007/978-3-319-41259-7_8
- [24] Niki Kilbertus, Mateo Rojas Carulla, Giambattista Parascandolo, Moritz Hardt, Dominik Janzing, and Bernhard Schölkopf. 2017. Avoiding Discrimination through Causal Reasoning. In *Advances in Neural Information Processing Systems 30* (Long Beach, CA) (NIPS 2017), I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (Eds.). Curran Associates, Inc., 656–666. <http://papers.nips.cc/paper/6668-avoiding-discrimination-through-causal-reasoning.pdf>
- [25] Michael P. Kim, Amirata Ghorbani, and James Zou. 2019. Multiaccuracy: Black-Box Post-Processing for Fairness in Classification. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society* (Honolulu, HI, USA) (AAIES '19). Association for Computing Machinery, New York, NY, USA, 247–254. <https://doi.org/10.1145/3306618.3314287>
- [26] Jon Kleinberg, Sendhil Mullainathan, and Manish Raghavan. 2017. Inherent Trade-Offs in the Fair Determination of Risk Scores. In *8th Innovations in Theoretical Computer Science Conference (ITCS 2017) (Leibniz International Proceedings in Informatics (LIPIcs), Vol. 67)*, Christos H. Papadimitriou (Ed.). Schloss Dagstuhl–Leibniz-Zentrum für Informatik, Dagstuhl, Germany, 43:1–43:23. <http://doi.org/10.4230/LIPIcs.ITCS.2017.0>
- [27] Matt J Kusner, Joshua Loftus, Chris Russell, and Ricardo Silva. 2017. Counterfactual Fairness. In *Advances in Neural Information Processing Systems 30* (NIPS 2017), I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (Eds.). Curran Associates, Inc., 4066–4076. <http://papers.nips.cc/paper/6995-counterfactual-fairness.pdf>
- [28] Jeff Larson and Julia Angwin. 2016. Technical Response to Northpointe. *ProPublica* (July 2016). <https://www.propublica.org/article/technical-response-to-northpointe>
- [29] Jeff Larson, Surya Mattu, Lauren Kirchner, and Julia Angwin. 2016. How We Analyzed the COMPAS Recidivism Algorithm. *ProPublica* (May 2016). <https://www.propublica.org/article/how-we-analyzed-the-compas-recidivism-algorithm>
- [30] Anthony W. Lowenkamp, Flores Kristin, and Bechtel Christopher T. 2016. False Positives, False Negatives, and False Analyses: A Rejoinder to Machine Bias: There's Software Used across the Country to Predict Future Criminals. And It's Biased against Blacks. *Federal Probation* 80, 2 (2016), 38–46. <https://www.uscourts.gov/federal-probation-journal/2016/09/false-positives-false-negatives-and-false-analyses-rejoinder>
- [31] Aditya Krishna Menon and Robert C Williamson. 2018. The cost of fairness in binary classification. In *Proceedings of the 1st Conference on Fairness, Accountability and Transparency (Proceedings of Machine Learning Research, Vol. 81)*, Sorelle A. Friedler and Christo Wilson (Eds.). PMLR, New York, NY, USA, 107–118. <http://proceedings.mlr.press/v81/menon18a.html>
- [32] Razieh Nabi, Daniel Malinsky, and Ilya Shpitser. 2019. Learning Optimal Fair Policies. In *Proceedings of the 36th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 97)*, Kamalika Chaudhuri and Ruslan Salakhutdinov (Eds.). PMLR, 4674–4682. <http://proceedings.mlr.press/v97/nabi19a.html>
- [33] Razieh Nabi and Ilya Shpitser. 2018. Fair Inference on Outcomes. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence* (New Orleans, LA). Association for the Advancement of Artificial Intelligence, 1931–1940. <https://aaai.org/ocs/index.php/AAAI/AAAI18/paper/view/16683>
- [34] Hari Krishnan Narasimhan. 2018. Learning with Complex Loss Functions and Constraints. In *Proceedings of the Twenty-First International Conference on Artificial Intelligence and Statistics (Proceedings of Machine Learning Research, Vol. 84)*, Amos Storkey and Fernando Perez-Cruz (Eds.). PMLR, 1646–1654. <http://proceedings.mlr.press/v84/narasimhan18a.html>
- [35] Jerzy Neyman. 1923. Justification of applications of the calculus of probabilities to the solutions of certain questions in agricultural experimentation. Excerpts English translation (Reprinted). *Statist. Sci.* 5 (1923), 463–472.
- [36] Ziad Obermeyer, Brian Powers, Christine Vogel, and Sendhil Mullainathan. 2019. Dissecting racial bias in an algorithm used to manage the health of populations. *Science* 366, 6464 (Oct. 2019), 447–453. <https://doi.org/10.1126/science.aax2342>
- [37] Lisa Rice and Deidre Swesnik. 2012. Discriminatory Effects of Credit Scoring on Communities of Color. *Suffolk University Law Review* 46 (2012), 935.
- [38] James Robins, Lingling Li, Eric Tchetgen, and Aad van der Vaart. 2008. Higher order influence functions and minimax estimation of nonlinear functionals. In *Probability and Statistics: Essays in Honor of David A. Freedman*, Deborah Nolan and Terry Speed (Eds.). Institute of Mathematical Statistics, Beachwood, Ohio, USA, 335–421. <https://doi.org/10.1214/193940307000000527>
- [39] Donald B Rubin. 2005. Causal Inference Using Potential Outcomes: Design, Modeling, Decisions. *J. Amer. Statist. Assoc.* 100, 469 (March 2005), 322–331. <https://doi.org/10.1198/016214504000001880>

- [40] Megan Stevenson. 2018. Assessing Risk Assessment in Action. *Minnesota Law Review* 103, 1 (2018), 83. <https://dx.doi.org/10.2139/ssrn.3016088>
- [41] Anastasios A. Tsiatis. 2006. *Semiparametric Theory and Missing Data*. Springer, New York, NY. <https://doi.org/10.1007/978-0-387-98135-2>
- [42] Mark J. van der Laan and James M. Robins. 2003. *Unified Methods for Censored Longitudinal Data and Causality*. Springer, New York, NY. <https://doi.org/10.1007/978-0-387-21700-0>
- [43] Aad van der Vaart. 2002. Semiparametric Statistics. In *Lectures on probability theory and statistics (Lecture notes in mathematics, 1781)*, Pierre Bernard (Ed.). Springer, Berlin. <https://doi.org/10.1007/b93152>
- [44] Tyler J. VanderWeele and Whitney R. Robinson. 2014. On the causal interpretation of race in regressions adjusting for confounding and mediating variables. *Epidemiology* 25, 4 (2014), 473–484. <https://doi.org/10.1097/EDE.0000000000000105>
- [45] Yixin Wang, Dhanya Sridhar, and David M. Blei. 2019. Equal Opportunity and Affirmative Action via Counterfactual Predictions. (2019). arXiv:1905.10870 <http://arxiv.org/abs/1905.10870>
- [46] Blake Woodworth, Suriya Gunasekar, Mesrob I. Ohannessian, and Nathan Srebro. 2017. Learning Non-Discriminatory Predictors. In *Proceedings of the 2017 Conference on Learning Theory (Proceedings of Machine Learning Research, Vol. 65)*, Satyen Kale and Ohad Shamir (Eds.). PMLR, Amsterdam, Netherlands, 1920–1953. <http://proceedings.mlr.press/v65/woodworth17a.html>
- [47] Muhammad Bilal Zafar, Isabel Valera, Manuel Gomez Rodriguez, and Krishna P. Gummadi. 2017. Fairness Beyond Disparate Treatment & Disparate Impact: Learning Classification without Disparate Mistreatment. In *Proceedings of the 26th International Conference on World Wide Web (Perth, Australia) (WWW '17)*. International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, CHE, 1171–1180. <https://doi.org/10.1145/3038912.3052660>
- [48] Junzhe Zhang and Elias Bareinboim. 2018. Fairness in Decision-Making – The Causal Explanation Formula. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence* (New Orleans, LA, 2018). Association for the Advancement of Artificial Intelligence, 2037–2045. <https://www.aaai.org/ocs/index.php/AAAI/AAAI18/paper/view/16949/15911>
- [49] Wenjing Zheng and van der Laan, Mark. 2010. Asymptotic Theory for Cross-validated Targeted Maximum Likelihood Estimation. *U.C. Berkeley Division of Biostatistics Working Paper Series Working Paper 273* (2010). <https://biostat.bepress.com/ucbbiostat/paper273/>

A PROOFS OF PROPOSITIONS

For convenience, we restate our three assumptions.

- 1. (Consistency) $Y = DY^1 + (1 - D)Y^0$
- 2. (Positivity) $\exists \delta \in (0, 1) : \mathbb{P}(\pi(A, X, S) \leq 1 - \delta) = 1$
- 3. (Ignorability) $Y^0 \perp\!\!\!\perp D \mid A, X, S$

Proof of Proposition 1 (Identification of error rates for input predictor S)

$$\begin{aligned}
 \text{cFPR}(S, a) &= \mathbb{P}(S = 1 \mid Y^0 = 0, A = a) \\
 &= \frac{\mathbb{P}(S = 1, Y^0 = 0, A = a)}{\mathbb{P}(Y^0 = 0, A = a)} \\
 &= \frac{\mathbb{E}[S(1 - Y^0)\mathbb{1}\{A = a\}]}{\mathbb{E}[(1 - Y^0)\mathbb{1}\{A = a\}]} \\
 &= \frac{\mathbb{E}[S(1 - \mathbb{E}[Y^0 \mid A, X, S, D = 0])\mathbb{1}\{A = a\}]}{\mathbb{E}[(1 - \mathbb{E}[Y^0 \mid A, X, S, D = 0])\mathbb{1}\{A = a\}]} \\
 &= \frac{\mathbb{E}[S(1 - \mu_0)\mathbb{1}\{A = a\}]}{\mathbb{E}[(1 - \mu_0)\mathbb{1}\{A = a\}]}
 \end{aligned}$$

$$\begin{aligned}
 \text{cFNR}(S, a) &= \mathbb{P}(S = 0 \mid Y^0 = 1, A = a) \\
 &= \frac{\mathbb{P}(S = 0, Y^0 = 1, A = a)}{\mathbb{P}(Y^0 = 1, A = a)} \\
 &= \frac{\mathbb{E}[(1 - S)Y^0\mathbb{1}\{A = a\}]}{\mathbb{E}[Y^0\mathbb{1}\{A = a\}]} \\
 &= \frac{\mathbb{E}[(1 - S)\mathbb{E}[Y^0 \mid A, X, S, D = 0]\mathbb{1}\{A = a\}]}{\mathbb{E}[\mathbb{E}[Y^0 \mid A, X, S, D = 0]\mathbb{1}\{A = a\}]} \\
 &= \frac{\mathbb{E}[(1 - S)\mu_0\mathbb{1}\{A = a\}]}{\mathbb{E}[\mu_0\mathbb{1}\{A = a\}]}
 \end{aligned}$$

The fourth equality in both derivations uses positivity and ignorability, and the fifth equality uses consistency.

Proof of Proposition 2 (Identification of the loss and fairness constraints)

Considering just the first component of the loss, we have:

$$\begin{aligned}
 (w^+) \mathbb{P}(S_\theta = 1, Y^0 = 0) &= (w^+) \mathbb{E}[S_\theta(1 - Y^0)] \\
 &= (w^+) \mathbb{E}[\mathbb{E}[S_\theta(1 - Y^0) \mid A, S, X]] \\
 &= (w^+) \mathbb{E} \{ \mathbb{E}[S_\theta \mid A, S](1 - \mathbb{E}[Y^0 \mid A, X, S]) \} \\
 &= (w^+) \mathbb{E} \{ \theta_{A,S}(1 - \mathbb{E}[Y^0 \mid A, X, S, D = 0]) \} \\
 &= (w^+) \mathbb{E} \{ \theta_{A,S}(1 - \mu_0) \}
 \end{aligned}$$

where the third equality uses that S_θ only depends on (A, S) , the fourth uses the definition of $\theta_{A,S}$ and ignorability, and the sixth uses consistency. Similar reasoning shows that $(w^-) \mathbb{P}(S = 0, Y^0 = 1) = (w^-) \mathbb{E} \{ (1 - \theta_{A,S})\mu_0 \}$. Combining these, we have

$$\begin{aligned}
 \mathcal{L}(S_\theta) &: = w^+ \mathbb{P}(S = 1, Y^0 = 0) + w^- \mathbb{P}(S = 0, Y^0 = 1) \\
 &= \mathbb{E}[\theta_{A,S}(w^+ - (w^+ + w^-)\mu_0)] + (w^-) \mathbb{E}[\mu_0] \\
 &= \theta^T \beta + (w^-) \mathbb{E}[\mu_0]
 \end{aligned}$$

We turn now to the fairness constraints. The error rates of the derived predictor S_θ depend on the error rates on the input predictor S as follows. Beginning with cFPR(S_θ, a), we have:

$$\begin{aligned}
 \mathbb{P}(S_\theta = 1 \mid Y^0 = 0, A = a) &= \\
 &= \sum_{s \in \{0,1\}} \mathbb{P}(S_\theta = 1 \mid Y^0 = 0, A = a, S = s) \mathbb{P}(S = s \mid Y^0 = 0, A = a) \\
 &= \sum_{s \in \{0,1\}} \mathbb{P}(S_\theta = 1 \mid A = a, S = s) \mathbb{P}(S = s \mid Y^0 = 0, A = a) \\
 &= \theta_{a,0}(1 - \text{cFPR}(S, a)) + \theta_{a,1} \text{cFPR}(S, a)
 \end{aligned}$$

where the first equality simply involves conditioning on S , and the second equality uses that $S_\theta \perp\!\!\!\perp Y^0 \mid A, S$. In other words, the false positive rate of S_θ depends only on θ and the false positive rate of the input predictor S . For the cFNR, by similar reasoning, we have:

$$\begin{aligned}
 \mathbb{P}(S_\theta = 0 \mid Y^0 = 1, A = a) &= \\
 &= 1 - \theta_{a,0}(\text{cFNR}(S, a)) + \theta_{a,1}(\text{cFNR}(S, a) - 1)
 \end{aligned}$$

The identification statements in the proposition follow by simply substituting in the expressions for cFPR(S, a), cFNR(S, a) from Proposition 1 and rearranging.

B PROOFS OF THEOREMS

We first introduce two lemmas used in the proofs of the theorems.

The first lemma gives sufficient conditions under which the optimal value of an estimated convex problem converges at a particular rate $f(n)$ to the optimal value of the target convex program. It is an adaptation of Theorem 3.5 in Shapiro (1991) that follows immediately from Theorems 2.1 and 3.4 in that same paper.

Lemma 1. (Shapiro, 1991) Let Θ be a compact subset of $\mathbb{R}^k$. Let $C(\Theta)$ denote the set of continuous real-valued functions on Θ , with $\mathcal{L} = C(\Theta) \times \dots \times C(\Theta)$ the r -dimensional Cartesian product. Let $\psi(\theta) = (\psi_0, \dots, \psi_r) \in \mathcal{L}$ be a vector of convex functions. Consider the quantity α^* defined as the solution to the following convex optimization program:

$$\begin{aligned} \alpha^* = \min_{\theta \in \Theta} \quad & \psi_0(\theta) \\ \text{subject to } & \psi_j(\theta) \leq 0, \quad j = 1, \dots, r \end{aligned}$$

Assume that Slater's condition holds, so that there is some $\theta \in \Theta$ for which the inequalities are satisfied and non-affine inequalities are strictly satisfied, i.e. $\psi_j(\theta) < 0$ if ψ_j is non-affine. Now consider a sequence of approximating programs, for $n = 1, 2, \dots$:

$$\begin{aligned} \hat{\alpha}_n = \min_{\theta \in \Theta} \quad & \hat{\psi}_{0n}(\theta) \\ \text{subject to } & \hat{\psi}_{jn}(\theta) \leq 0, \quad j = 1, \dots, r \end{aligned}$$

with $\hat{\psi}_n(\theta) := (\hat{\psi}_{0n}, \dots, \hat{\psi}_{rn}) \in \mathcal{L}$. Assume that $f(n)(\hat{\psi}_n - \psi)$ converges in distribution to a random element $W \in \mathcal{L}$ for some real-valued function $f(n)$. Then:

$$f(n)(\hat{\alpha}_n - \alpha_0) \rightsquigarrow L$$

for a particular random variable L . It follows that $\hat{\alpha}_n - \alpha_0 = O_{\mathbb{P}}(1/f(n))$.

Lemma 2. Let ξ, W be constant vectors and $\hat{\xi}_n, \hat{W}_n$ be random variables, with $\xi - \hat{\xi}_n = O_{\mathbb{P}}(1/f(n))$ for some real-valued $f(n)$. If, for all $M > 0$, $\mathbb{P}(\|W - \hat{W}_n\| > M) \leq \mathbb{P}(\|\xi - \hat{\xi}_n\| > CM)$ for some constant C , then $W - \hat{W}_n = O_{\mathbb{P}}(1/f(n))$.

PROOF. For any $\epsilon > 0$, there exists some $M_\epsilon > 0$ such that $\mathbb{P}(f(n)\|\xi - \hat{\xi}_n\| > M_\epsilon) < \epsilon$ for all n large enough. Set $M = M_\epsilon/C$. Then $\mathbb{P}(f(n)\|W - \hat{W}_n\| > M) < \epsilon$ for all n large enough. $\square$

B.1 Theorem 1 (loss gap)

We expand the loss by introducing the term $\hat{\beta}^T \hat{\theta}$, which is the quantity that is minimized in the course of computing $\hat{\theta}$. We proceed by splitting the loss into two terms and showing that each of those terms is $O_{\mathbb{P}}(1/f(n))$.

PROOF. The loss gap can be expanded as follows:

$$\begin{aligned} \mathcal{L}(S_{\hat{\theta}}) - \mathcal{L}(S_{\theta^*}) &= \beta^T \hat{\theta} - \beta^T \theta^* \\ &= \underbrace{(\beta^T \hat{\theta} - \hat{\beta}^T \hat{\theta})}_{(1)} + \underbrace{(\hat{\beta}^T \hat{\theta} - \beta^T \theta^*)}_{(2)} \end{aligned}$$

For term (1), we have

$$\begin{aligned} \hat{\theta}^T (\beta - \hat{\beta}) &\leq \|\hat{\theta}\| \|\beta - \hat{\beta}\| \\ &\leq 2\|\beta - \hat{\beta}\| \\ &= O_{\mathbb{P}}(1/f(n)) \end{aligned}$$

where the first line uses Cauchy-Schwarz, the second line follows from the fact that $\hat{\theta} \in [0, 1]^4$, and the third line follows by assumption. For term (2), we rely on Lemma 1. Note that we can write

$$\begin{aligned} \mathcal{L}(S_{\theta^*}) &= \min_{\theta \in \Theta} \psi_0(\theta) \\ &\text{subject to } \psi_j(\theta) \leq 0, \quad j = 1, \dots, r \\ \hat{\mathcal{L}}(S_{\hat{\theta}}) &= \min_{\theta \in \Theta} \hat{\psi}_0(\theta) \\ &\text{subject to } \hat{\psi}_j(\theta) \leq 0, \quad j = 1, \dots, r \end{aligned}$$

with $\Theta = [0, 1]^4$, and

$$\begin{aligned} \psi(\theta) &= (\mathcal{L}, \Delta^+ - \epsilon^+, -\Delta^+ - \epsilon^+, \Delta^- - \epsilon^-, -\Delta^- - \epsilon^-) \\ \hat{\psi}(\theta) &= (\hat{\mathcal{L}}, \hat{\Delta}^+ - \epsilon^+, -\hat{\Delta}^+ - \epsilon^+, \hat{\Delta}^- - \epsilon^-, -\hat{\Delta}^- - \epsilon^-) \end{aligned}$$

where for brevity we omit the argument S_{θ} to the loss $\mathcal{L}$ and the error rate differences Δ^+, Δ^- . Since these are linear programs, Slater's condition is satisfied. (We assume the LPs are feasible.) By assumption, each of the estimators in $\hat{\psi}(\theta)$ converges at rate $f(n)$, so $f(n)(\hat{\psi}(\theta) - \psi(\theta))$ converges to some (unknown) random variable. (We rule out pathological cases in which this does not happen.) Per Lemma 1, it follows that $\hat{\beta}^T \hat{\theta} - \beta^T \theta^* = \hat{\mathcal{L}}(S_{\hat{\theta}}) - \mathcal{L}(S_{\theta^*}) = O_{\mathbb{P}}(1/\sqrt{n})$.

The sum of the two terms in the loss gap is therefore also $O_{\mathbb{P}}(1/\sqrt{n})$. $\square$

B.2 Theorem 2 (excess unfairness)

The proof relies on Lemma 2 and the convergence of the estimated LP coefficient vectors $\hat{\beta}^+, \hat{\beta}^-$. When $\hat{\beta}^+, \hat{\beta}^-$ are close to β^+, β^- , the excess unfairness must be small for any $\theta \in \Theta = [0, 1]^4$, including of course $\hat{\theta}$.

PROOF. We have

$$\begin{aligned} &\mathbb{P}\left(\text{UF}^+(S_{\hat{\theta}}) > \delta \text{ or } \text{UF}^-(S_{\hat{\theta}}) > \delta\right) \\ &\leq \mathbb{P}\left(|\theta^T \beta^+| - |\theta^T \hat{\beta}^+| > \delta \text{ or } |\theta^T \beta^-| - |\theta^T \hat{\beta}^-| > \delta\right) \\ &\quad \text{for some } \theta \in [0, 1]^4 \\ &\leq \mathbb{P}\left(|\theta^T \beta^+ - \theta^T \hat{\beta}^+| > \delta \text{ or } |\theta^T \beta^- - \theta^T \hat{\beta}^-| > \delta\right) \\ &\quad \text{for some } \theta \in [0, 1]^4 \\ &\leq \mathbb{P}\left(\|\theta\| \cdot \|\hat{\beta}^+ - \beta^+\| > \delta \text{ or } \|\theta\| \cdot \|\hat{\beta}^- - \beta^-\| > \delta\right) \\ &\quad \text{for some } \theta \in [0, 1]^4 \end{aligned}$$

$$\begin{aligned}
&\leq \mathbb{P}\left(2\|\widehat{\beta}^+ - \beta^+\| > \delta \text{ or } 2\|\widehat{\beta}^- - \beta^-\| > \delta\right) \\
&\leq \mathbb{P}\left(2\|\widehat{\beta}^+ - \beta^+\| > \delta\right) + \mathbb{P}\left(2\|\widehat{\beta}^- - \beta^-\| > \delta\right) \\
&= \mathbb{P}\left(\|\widehat{\beta}^+ - \beta^+\| > \delta/2\right) + \mathbb{P}\left(\|\widehat{\beta}^- - \beta^-\| > \delta/2\right)
\end{aligned}$$

where the third inequality uses Cauchy-Schwartz, the fourth uses that $\theta \in [0, 1]^4 \implies \|\theta\| \leq 2$, and the fifth uses the union bound. (The norm here is the Euclidean norm.) The reasoning in the first inequality is as follows: if $\text{UF}^+(S_{\widehat{\theta}}) > \delta$, then $|\widehat{\theta}^T \beta^+| - |\widehat{\theta}^T \widehat{\beta}^+| > \delta$, since $\widehat{\theta}^T \widehat{\beta}^+ \leq \epsilon^+$ by construction. A necessary condition, then, is that $|\widehat{\theta}^T \beta^+| - |\widehat{\theta}^T \widehat{\beta}^+| > \delta$ for some $\theta \in [0, 1]^4$.

It follows from Lemma 2 that

$$\max\left\{\text{UF}^+(S_{\widehat{\theta}}), \text{UF}^-(S_{\widehat{\theta}})\right\} = O_{\mathbb{P}}(1/f(n))$$

□

C SAMPLE SPLITTING

A training sample, $\mathcal{D}_{\text{train}}$, is used to estimate $\widehat{\theta}$, while a separate sample $\mathcal{D}_{\text{test}}$ is used to estimate the risk and fairness properties of the derived predictor $S_{\widehat{\theta}}$ conditional on $\widehat{\theta}$. Within each sample, separate folds are used to estimate the nuisance parameters μ_0 and π and the target parameters.

The following schematic illustrates this procedure. k -fold cross fitting can be used within each sample to recover full sample size efficiency. For convenience, we assume that each of the four samples is of size n , though our results require only that each sample is $O(n)$.

D SIMULATIONS: DATA GENERATING PROCESS

The data generating process used in section 6.1.2 to illustrate Theorems 1 and 2 is as follows, for data $Z = (A, X, S, D, Y^0, Y^1, Y)$.

$$\mathbb{P}(A = 1) = 0.3$$

$$X | A \sim N(A * (1, -0.8, 4, 2)^T, I_4)$$

$$\mathbb{P}_{\text{pre}}(D = 1 | A, X) = \min\{0.975, \text{expit}((A, X)^T(0.2, -1, 1, -1, 1))\}$$

$$\mathbb{P}_{\text{post}}(D = 1 | A, X, S) = \min\{0.975, \text{expit}((A, X, S)^T(0.2, -1, 1, -1, 1, 1))\}$$

$$\mathbb{P}(Y^0 = 1 | A, X) = \text{expit}((A, X)^T(-5, 2, -3, 4, -5))$$

$$\mathbb{P}(Y^1 = 1 | A, X) = \text{expit}((A, X)^T(1, -2, 3, -4, 5))$$

$$Y = (1 - A)Y^0 + AY^1$$

where I_4 denotes the 4×4 identity matrix and N denotes a Gaussian distribution. The predictor $S(A, X)$ is trained using random forests. The pre-RAI decision making process doesn't depend on S ; the post-RAI process does.

For the simulations used in section 6.1.3 to illustrate fairness-performance tradeoffs, the distribution is identical except that $\mathbb{P}(Y^0 = 1 | A, X) = \text{expit}((A, X)^T(-4, 0.4, 0.6, 0.8, -1))$.

E ESTIMATING PREDICTIVE CHANGE

The predictive change $\mathbb{P}(S_{\widehat{\theta}} \neq S)$, which is referred to in section 6.2, represents the proportion of input predictions that the post-processed predictor $S_{\widehat{\theta}}$ flips. It can be estimated as follows:

$$\widehat{\mathbb{P}}(S_{\widehat{\theta}} \neq S) = \mathbb{P}_n\left\{\mathbb{P}\left(S_{\widehat{\theta}} \neq S | A, S\right)\right\} =$$

$$\mathbb{P}_n\left\{\sum_{a \in \{0,1\}} [\theta_{a,0} \mathbb{1}\{A = a, S = 0\} + (1 - \theta_{a,1}) \mathbb{1}\{A = a, S = 1\}]\right\}$$

Since this is a sample average, confidence intervals can be derived via the central limit theorem.

F NOTATION

Input data	
$Z = (A, X, D, S, Y) \sim \mathbb{P}$	Sensitive feature A , covariates X , decision (treatment, intervention) D , input predictor S , outcome Y
Derived predictor	
$S_\theta \sim \text{Bern}(\theta_{A,S})$ $\theta_{a,s} = \mathbb{P}(S_\theta = 1 \mid A = a, S = s)$ $\theta_{A,S} = \sum_{a,s \in \{0,1\}} \mathbb{1}\{A = a, S = s\} \theta_{a,s}$ $\theta = (\theta_{0,0}, \theta_{0,1}, \theta_{1,0}, \theta_{1,1})^T$ $\tilde{\theta} = (0, 1, 0, 1)$	Predictor derived from S Conditional probability that defines S_θ RV that takes value $\theta_{a,s}$ with probability $\mathbb{P}(A = a, S = s)$ Optimization parameter The value such that $S_{\tilde{\theta}} = S$
Nuisance parameters	
$\pi = \pi(A, X, S) = \mathbb{P}(D = 1 \mid A, X, S)$ $\mu_0 = \mu_0(A, X, S, D) = \mathbb{E}[Y \mid A, X, S, D = 0]$ $\phi = \frac{1-D}{1-\pi}(Y - \mu_0) + \mu_0$	Propensity score for the decision Outcome regression Uncentered influence function for $E[Y^0]$
Loss parameters	
w^+, w^- $\beta_{a,s} = \mathbb{E}[\mathbb{1}\{A = a, S = s\}(w^+ - (w^+ + w^-)\mu_0)]$ $\beta = (\beta_{0,0}, \beta_{0,1}, \beta_{1,0}, \beta_{1,1})^T$ $\mathcal{L}(S_\theta) = w^+ \mathbb{P}(S_\theta = 1, Y^0 = 0) + w^- \mathbb{P}(S_\theta = 0, Y^0 = 1)$ $\Gamma(S_\theta)$	Weights on the false positive and false negative rates A coefficient in the loss, for $a, s \in \{0, 1\}$ Vector of loss coefficients Loss of S_θ , equivalent to $\theta^T \beta + w^- \mathbb{E}[\mu_0]$ loss change $\mathcal{L}(S_\theta) - \mathcal{L}(S)$
Fairness parameters	
$\text{cFPR}(S_\theta, a) = \mathbb{P}(S_\theta = 1 \mid Y^0 = 0, A = a)$ $\text{cFNR}(S_\theta, a) = \mathbb{P}(S_\theta = 0 \mid Y^0 = 1, A = a)$ $\beta^+ = (1 - \text{cFPR}(S, 0), \text{cFPR}(S, 0), \text{cFPR}(S, 1) - 1, -\text{cFPR}(S, 1))$ $\beta^- = (-\text{cFNR}(S, 0), \text{cFNR}(S, 0) - 1, \text{cFNR}(S, 1), 1 - \text{cFNR}(S, 1))$ $\Delta^+(S_\theta) = \theta^T \beta^+ = \text{cFPR}(S_\theta, 0) - \text{cFPR}(S_\theta, 1)$ $\Delta^-(S_\theta) = \theta^T \beta^- = \text{cFNR}(S_\theta, 0) - \text{cFNR}(S_\theta, 1)$ ϵ^+, ϵ^- $\text{UF}^+(S_\theta) = \max(\Delta^+(S_\theta) - \epsilon^+, 0)$ $\text{UF}^-(S_\theta) = \max(\Delta^-(S_\theta) - \epsilon^-, 0)$	Counterfactual FPR for S_θ for group a Counterfactual FNR for S_θ for group a Coefficients defining the fairness constraints Error rate differences of the predictor S_θ in the cFPR Error rate differences of the predictor S_θ in the cFNR Fairness constraints on Δ^+ and Δ^- Excess unfairness in the cFPR Excess unfairness in the cFNR
Optimal fair derived predictor	
$\arg \min_\theta \mathcal{L}(S_\theta) \text{ s.t. } \Delta^+(S_\theta) \leq \epsilon^+, \Delta^-(S_\theta) \leq \epsilon^-$	Parameter defining the optimal fair derived predictor S_{θ^*}