

Pragmatic Fairness: Developing Policies with Outcome Disparity Control

Limor Gultchin

University of Oxford, The Alan Turing Institute, Google DeepMind

LIMOR.GULTCHIN@GMAIL.COM

Siyuan Guo

University of Cambridge, Max Planck Institute for Intelligent Systems

Alan Malek

Google DeepMind

Silvia Chiappa

Google DeepMind

Ricardo Silva

University College London

Editors: Francesco Locatello and Vanessa Didelez

Abstract

We introduce a causal framework for designing optimal policies that satisfy classes of fairness constraints. We take a pragmatic approach asking what we can do with an action space available from historical data, with no further experimentation and novel actions immediately available. We propose two different fairness constraints: a “moderation breaking” constraint which aims at reducing disparity in outcome levels across sensitive attributes to the extent the provided action space permits; and an “equal benefit” constraint which aims at distributing gain from the new and maximized policy equally across sensitive attribute levels, and thus at keeping pre-existing preferential treatment in place or avoiding the introduction of new disparity. We introduce practical methods for implementing the constraints and illustrate their uses on experiments with semi-synthetic models.

1. Introduction

The fairness of decisions made by machine learning models involving underprivileged groups has seen increasing attention and scrutiny by the academic community and beyond. A growing body of literature has been looking at the unfavourable treatment that might arise from historical biases present in the data, data collection practices, or the limits of modelling choices and techniques. Within this field of study, the vast majority of works has considered the problem of designing *fair prediction systems*, i.e. systems whose outcomes satisfy certain properties with respect to membership in a sensitive group (Verma and Rubin, 2018; Barocas et al., 2019; Mehrabi et al., 2019; Wachter et al., 2020; Mitchell et al., 2021; Pessach and Shmueli, 2022). In contrast, relatively little attention has been given to the problem of designing *fair optimal policies* (Joseph et al., 2016, 2018; Gillen et al., 2019; Kusner et al., 2019; Nabi et al., 2019; Chohlas-Wood et al., 2021). In this case, the goal is to design a decision making system that specifies how to select *actions* that maximize a *downstream outcome* of interest subject to fairness constraints.

The consideration of an outcome downstream of the action allows a more flexible and powerful approach to fair decision making, as it enables to optimize future outcomes rather than matching historical ones, and also to enforce fairness constraints on the effects of the actions rather than




Figure 1: (a): Causal graph with associated distribution $p(A, S, X, Y; \sigma_A) = p(Y|A, S, X)p(A|S, X; \sigma_A)p(S, X)$ describing our fair optimal policy setting. (b): A schema of the fair prediction system setting, where we do not control the true outcome but a predictor of it. A long term distribution shift for Y in the population may take place by future applications of the new predictor, but no short term, individual level effects are expected to be significant.

on the actions themselves. Enforcing fairness constraints on the effects of the actions in decision problems that correspond to allocations of resources or goods may help reduce the risk of unfair delayed impact (Dwork et al., 2012; Liu et al., 2018; D’Amour et al., 2020).

In this work, we propose a causal framework for designing fair optimal policies that is inspired by the public health literature (Jackson and VanderWeele, 2018; Jackson, 2018, 2020). We assume the causal graph in figure 1a, where action A depends on sensitive attribute S and covariates X ; and where S and X , which can be associated, potentially directly influence the outcome Y . We define the policy with the conditional distribution $p(A | S, X; \sigma_A)$ parameterized by σ_A —represented in the graph as a *regime indicator* (Dawid, 2007; Correa and Bareinboim, 2020). The aim is to learn a parametrization that maximizes Y in expectation while controlling its dependence on S . This graph describes common real-world settings in which the action can only indirectly control for the dependence of Y on S , and therefore to a level that depends on the available action space. In addition, it enables to learn the optimal policy from available historical data collected using a baseline policy $p(A | S, X; \sigma_A = \emptyset)$, overcoming issues such as, e.g., cost or ethical constraints in randomizing actions, which are common in real-world applications where fairness is of relevance. To highlight the difference between our fair optimal policy framework and the fair prediction system framework, we contrast it to the causal graph in figure 1b, where the action is the definition of a predictor \hat{Y} that is designed to approximate some unknown real outcome Y . We describe the goal as learning a parametrization such that $p(\hat{Y} | S, X; \sigma_Y)$ matches the distribution of past outcomes $p(Y | S, X; \sigma_Y = \emptyset)$ while controlling for dependence of the outcome distribution on S .

We consider two constraints for controlling disparity in the downstream outcomes with respect to S , which may be applicable for different use cases, depending on context and available actions: (i) a moderation breaking constraint, which aims at actively reducing disparity to the extent permitted by the action space; and (ii) an equal benefit constraint, which aims at equalizing the disparity of a new policy with that of the baseline policy in order to maintain the preferential treatment present in the baseline policy or to conservatively avoid introducing new disparity. We demonstrate the performance of our framework on two real-world based settings.

Symbol	Explanation
σ_A	Regime indicator for conditional distribution
A	Action influenced by S and X
S	Sensitive attribute
X	Covariates influencing both A and Y
Y	Outcome
Y_{σ_A}	Potential outcome under the learned policy
Y_{\emptyset}	Potential outcome under the baseline policy
$p(A S, X; \emptyset)$	Baseline policy
$p(A S, X; \sigma_A)$	Parameterized learned policy
$\mu_Y(a, s, x)$	Expected outcome given a, s , and x
$f(s, x), g(a, s, x), h(a, x)$	Components of μ_Y for ModBrk constraint
ϵ	Slack variable for fairness constraint
$\mu_{\sigma_A}^Y$	Expected outcome under policy σ_A
$g_{\sigma_A}(s, x), h_{\sigma_A}(x)$	Adjusted components of $\mu_Y(\sigma_A)$
$\mathcal{N}(\cdot, \cdot)$	Gaussian distribution notation
$\Phi(\cdot)$	Cumulative distribution function for Gaussian
$E[\cdot]$	Expectation operator
$\rho(s, x)$	Correlation coefficient for joint Gaussian assumptions
$F(\cdot)$	Cumulative distribution function
μ_A, V_A	Mean and variance in Gaussian policies
μ_Y, V_Y	Mean and variance of outcomes

Table 1: Symbols and notations in this paper.

2. Outcome Disparity Controlled Policy

We assume the setting represented by the causal graph in figure 1a, where the sensitive attribute S corresponds to characteristics of an individual—such as race, gender, disabilities, sexual or political orientation—which we wish to protect against some measure of unfairness. We focus on discrete S , but the proposed methods can be adapted to settings with continuous S .

This causal graph encodes the following statistical independencies assumptions:

$$(a) \sigma_A \perp\!\!\!\perp Y \mid \{A, S, X\}$$

$$(b) \sigma_A \perp\!\!\!\perp \{S, X\}.$$

These assumptions are not restrictive, as they are satisfied in any setting in which the partial ordering among nodes is given by $\{S \cup X, \sigma_A, A, Y\}$. Assumption (a) enables us to learn the optimal policy based on historical data collected using a *baseline policy* $p(A | S, X; \sigma_A = \emptyset)$ (observational-data regime), rather than by taking actions (interventional-data regime). Such a baseline policy represents the action allocation that was in place during the collection of the data pre-optimization. Assumption (b) allows us to compute the constraints with the proposed estimation methods below. We denote with Y_{σ_A} the *potential outcome* random variable, which represents the outcome resulting from taking actions according to $p(A | S, X; \sigma_A)$ and has distribution equal to $p(y; \sigma_A) := \int_{a,s,x} p(y | a, s, x) p(a, s, x; \sigma_A)$. In the remainder of the paper, we indicate the baseline policy and potential outcome with $p(A | S, X; \emptyset)$ and Y_{\emptyset} respectively.

The goal of the decision maker is to learn a parametrization σ_A that maximizes the expectation of the potential outcome, $\mathbb{E}[Y_{\sigma_A}]$, while also controlling the disparity in Y_{σ_A} across S in the following ways:

- (i) through a *moderation breaking* (ModBrk) constraint the aims at removing dependence of $\mathbb{E}[Y_{\sigma_A}]$ on S as much as possible via what can be controlled, namely the allocation of A as determined by $p(A|S, X; \sigma_A)$;
- (ii) through an *equal benefit* (EqB) constraint that requires the distribution of $Y_{\sigma_A} - Y_{\emptyset}$ to be approximately equal across different values of S , ensuring that gain from the new policy is distributed equally.

2.1. Disparity Control via ModBrk Constraint

Without loss of generality¹, $\mu^Y(a, s, x) := \mathbb{E}[Y | a, s, x]$ can be written as

$$\mu^Y(a, s, x) = f(s, x) + g(a, s, x) + h(a, x), \quad (1)$$

where $h(a, x)$ is defined to be a function of components of (A, X) which are marginally independent of S . In particular, $\mathbb{E}[h(A, X) | S = s; \sigma_A]$ and $\mathbb{E}[h(A, X) | X = x, S = s; \sigma_A]$ are not functions of S .

The above leads to the following decomposition

$$\begin{aligned} \mu_{\sigma_A}^Y(s, x) &:= \mathbb{E}[Y_{\sigma_A} | s, x] = \int_a \mu^Y(a, s, x) p(a | s, x; \sigma_A) \\ \mu_{\sigma_A}^Y(s, x) &= f(s, x) + g_{\sigma_A}(s, x) + h_{\sigma_A}(x), \end{aligned}$$

where $g_{\sigma_A}(s, x) := \mathbb{E}[g(A, s, x) | s, x; \sigma_A]$, and $h_{\sigma_A}(x) := \mathbb{E}[h(A, x) | s, x; \sigma_A]$. This decomposition contains (i) a component $f(s, x)$ that cannot be affected by σ_A , but which can contribute to disparity; (ii) a component $h_{\sigma_A}(x)$ that can be adjusted by σ_A to increase expected outcomes, but which is not affected by S ; and (iii) a component $g_{\sigma_A}(s, x)$ by which the choice of σ_A can influence differences that are *moderated* by S .

This decomposition is *not* unique. However, regardless of the exact form of the decomposition, it follows that for any two levels s and \bar{s} of S , we have

$$\mathbb{E}[Y_{\sigma_A} | S = s] - \mathbb{E}[Y_{\sigma_A} | S = \bar{s}] :=$$

$$\mathbb{E}[g_{\sigma_A}(s, X) | S = s] - \mathbb{E}[g_{\sigma_A}(\bar{s}, X) | S = \bar{s}] + \mathbb{E}[f(s, X) | S = s] - \mathbb{E}[f(\bar{s}, X) | S = \bar{s}].$$

If we want to enforce an upper bound on $|\mathbb{E}[Y_{\sigma_A} | S = s] - \mathbb{E}[Y_{\sigma_A} | S = \bar{s}]|$ as an operational definition of disparity of σ_A as moderated by S , it suffices to upper bound

$$|\mathbb{E}[g_{\sigma_A}(s, X) | S = s] - \mathbb{E}[g_{\sigma_A}(\bar{s}, X) | S = \bar{s}] + \mathbb{E}[f(s, X) | S = s] - \mathbb{E}[f(\bar{s}, X) | S = \bar{s}]|.$$

Considering that $\mathbb{E}[f(s, X) | S = s]$ is not a function of σ_A (since $(S, X) \perp\!\!\!\perp \sigma_A$) and solely a function of s , we will consider them to be nuisance parameters and instead enforce an upper bound

$$|\mathbb{E}[g_{\sigma_A}(s, X) | S = s] - \mathbb{E}[g_{\sigma_A}(\bar{s}, X) | S = \bar{s}]| \leq \epsilon^{1/2}$$

1. This can be seen by setting $f(s, x) = h(a, x) = 0$.

uniformly on all possible combinations (s, \bar{s}) . Taking the square of both sides above to ensure differentiability, we get to the following mathematical program

$$\arg \max_{\sigma_A} \mathbb{E}[Y_{\sigma_A}] \quad \text{s.t.} \quad (\mathbb{E}[g_{\sigma_A}(s, X) | S = s] - \mathbb{E}[g_{\sigma_A}(\bar{s}, X) | S = \bar{s}])^2 \leq \epsilon, \forall s, \bar{s}. \quad (2)$$

The slack ϵ is chosen based on domain requirements and the feasibility of the problem: central to the setup is that we work with a given space of policies, which is constrained by real-world phenomena and, in general, can only do so much to reduce disparity. In practice, a range of ϵ values should be tried, and the corresponding disparity constrains measured, as well as their impact of the expected outcome.

Parceling out f and h puts the focus on the *interaction* between σ_A and S , which is the only component that can be controlled in order to reduce disparity in expectation. In our experiments, we explicitly parameterize Eq. (1) as the ensemble of three neural networks, each representing f , g and h , which learns this conditional expectation by mean squared error minimization. For simplicity we do *not* enforce any identifiability, which was not required for the derivation above, relying on standard regularization to guide the smoothness of g_{σ_A} to be no less than what suffices to reconstruct $\mathbb{E}[Y | a, s, x]$ after accounting for terms in which S and σ_A do not interact.

Motivation & Suitability. ModBrk aims at removing S 's influence on the outcome. As such, it is appropriate when disparity is illegitimate. This is analogous to the appropriateness of the demographic parity constraint in fair prediction systems.

However, ModBrk controls disparity via the separate actions and therefore to the extent permitted by the action space. Some insights into ModBrk can be gained by noticing that we are operating in the (estimated) true process that follows a particular σ_A . We marginalize the intermediate process between (A, S, X) and the outcome Y , but implicitly assume that actions change mediating events, such as M in figure 2a. The extent by which we can

mitigate unfairness is a property of the real-world space of policies. For instance, we cannot interfere in the direct dependence between S and Y except for the moderating effect that M might have under a new policy. This is in contrast to previous work, such as [Nabi et al. \(2019\)](#), which solves a planning problem in a “projection” of the real-world process where unfair information has been removed by some criteria.

An example in which the ModBrk constraint would be suitable is that of a company looking to change its outreach campaign A to maximize job applications Y while mitigating current imbalances in their demographics S and level of experience X (figure 2b, [green labels](#)). The company cannot control factors such as cultural preference among applicants for industry sectors, but can induce modifications by focusing recruiting efforts in events organized by minority groups in relevant conferences and by choosing recruiting strategies that do not interact with group membership.




Figure 2: (a): Causal graph providing further intuition for the ModBrk constraint. (b): Example use cases for the suggested constrains. See full description in the text.

Likewise, if we have policies to increase a measure of wealth or other economical indicators across subgroups of S , we may not wish to do so while keeping disparities uncontrolled.

We emphatically do *not* claim ModBrk is an universal principle. For instance, if Y is a particular medical outcome, we may want to have as much of that maximized regardless of differences in subgroup status.

2.2. Disparity Control via EqB Constraint

We formalize the requirement that the distribution of $Y_{\sigma_A} - Y_{\emptyset}$ must be approximately equal across different values of S using the cumulative distribution function (cdf) \mathcal{F} , leading to the constrained objective

$$\arg \max_{\sigma_A} \mathbb{E}[Y_{\sigma_A}] \quad \text{s.t.} \quad \mathcal{F}(Y_{\sigma_A} - Y_{\emptyset} | S = s) = \mathcal{F}(Y_{\sigma_A} - Y_{\emptyset} | S = \bar{s}), \forall s, \bar{s}. \quad (3)$$

The EqB constraint focuses on $Y_{\sigma_A} - Y_{\emptyset}$ because it is proposing a notion of fairness which looks to equally distribute gains from a new policy compared to an initial condition, represented by Y_{\emptyset} , or in fact keep preferential treatment present in the original policy. That is part of the pragmatic nature of our proposal.

In general, the distribution of $Y_{\sigma_A} - Y_{\emptyset}$ is not identifiable, due to the fact that the decision maker has access only to realization of Y under the baseline policy. This problem is similar to the “fundamental problem of causal inference” in the context of hard interventions, for which bounding approaches have been suggested (Pearl, 2000; Wu et al., 2019; Miratrix et al., 2018). Similarly, we propose matching upper bounds and lower bounds of the cdf.

Motivation & Suitability. EqB aims at not increasing disparity compared to the baseline policy. The EqB constraint would be appropriate in two scenarios. The first scenario is when we would like to keep the disparity present in the baseline policy when moving to the new one, since it includes legitimate disparity or desirable preferential treatment. This is a similar motivation to the equalized odds constraint in fair prediction systems. This scenario would arise, for example, when wishing to design an income taxation system to optimize disposable income or savings levels (figure 2b, purple labels). Disparities in Y across gender or racial groups S may exist, as well as possibly related individual characteristics X , like socio-economic status, employment type and housing situation. The decision maker should take into account and maintain desirable preferential treatments present in the baseline policy, for example, for working moms.

The second scenario is when we would like to not introduce greater disparity in a new policy, acknowledging the limitations of the action space. This scenario would arise, e.g., when wishing to perform a software update of a medical or well-being app that includes different designs, with users belonging to different age groups. The decision maker might want to ensure that the new policy does not change existing levels of difference in health or wellbeing across groups, which relates to the usage of the app. The app use, as well as the downstream outcome Y itself, has pre-existing relations to age groups S as well as to other associated user characteristics X , but the decision maker wants the update rollout to not exacerbate the already existing gap in outcome under the baseline policy. The decision maker may not want to forcefully reduce disparities in outcome in such a case – it is likely that younger age groups would be more responsive and happier with a software update, and have higher levels of the downstream outcome Y overall; the choice of policy can only do so much to change that.

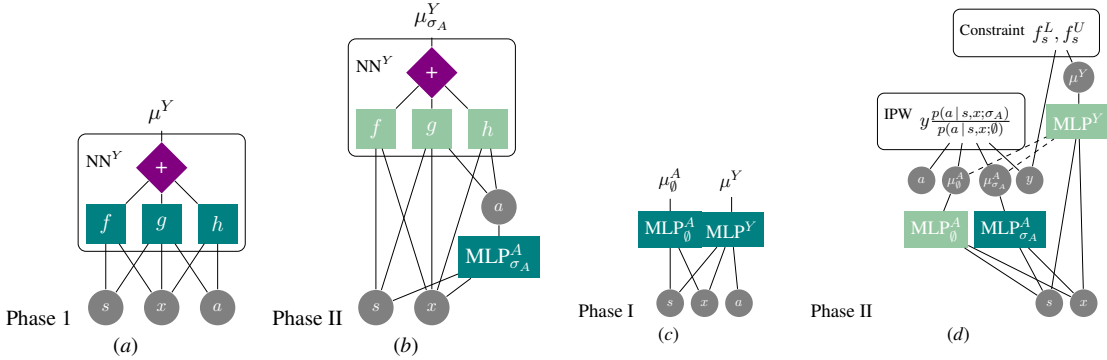


Figure 3: (a, b) Training phases for ModBrk. (c, d) Training phases for EqB. Gray circles: inputs. Teal blocks: parameter layers. Light green blocks: fixed parameters. Purple diamonds: additive gates. Dashed edge: alternating inputs.

3. Relation to Prior Work

While more attention has been given to the problem of designing fair prediction systems, some works have considered the problem of designing fair optimal policies. Like us, [Nabi et al. \(2019\)](#) use observational data, but in order to learn a policy as if particular path-specific effects between S and A and S and Y were completely deactivated. They do not place any constraints on the distribution of Y_{σ_A} given S and X , require complex counterfactual computations, and aim to achieve a notion of fair policy which targets specific causal subpaths. Critically, they rely on a manipulation of S , and do not have an equivalent to our pragmatic view, asking what can be done with a set of available actions. [Chohlas-Wood et al. \(2021\)](#) consider a more general utility function than ours as the optimization objective but focuses on enforcing a notion similar to demographic parity on the choice of the policies, rather than considering a fairness notion on the outcomes. The most similar work to ours is the one from [Kusner et al. \(2019\)](#), still in the observational data regime. Crucially, this work is different in its motivation: it aims to extend the work in [Kusner et al. \(2017\)](#) to the policy setting, and thus relies on manipulation of S . Further, it consists mostly of budget treatment allocations and interference problems. It does not have our pragmatic view, and does not consider parameterized policy spaces that take into account a covariate vector X . Due to the difference in their disparity definitions, these methods would not be directly comparable to ours. There is also a growing literature on fair policy optimization in online settings, but those deal with sequential decisions and interventional data, which are fundamentally different from ours (see Appendix A).

4. Method

In this section, we describe how the ModBrk and EqB constrained objectives (2) and (3) can be estimated from an observational dataset $\mathcal{D} = \{a^i, s^i, x^i, y^i\}_{i=1}^N$, $(a^i, s^i, x^i, y^i) \sim p(A, S, X, Y; \sigma_\theta)$, and introduce a method for optimizing σ_A using neural networks for each constraint. In both cases, we enforce the constraints via an augmented Lagrangian approach, casting them as inequality constraints controlled by a slack value ϵ (see Appendix B). Different choices of ϵ lead to different trade-offs between utility and constraints, similarly to varying Lagrange multiplier values.

4.1. ModBrk Constraint

For the ModBrk constrained objective (2), we learn a deterministic policy using an MLP neural network $\text{MLP}_{\sigma_A}^A$, i.e. $p(A|s, x; \sigma_A) = \delta_{A=\text{MLP}_{\sigma_A}^A(s, x)}$, where δ denotes the delta function and $\text{MLP}_{\sigma_A}^A(s, x)$ the output of $\text{MLP}_{\sigma_A}^A$ given input s, x . This gives $\mu_{\sigma_A}^Y(s, x) = \mathbb{E}[Y_{\sigma_A}|s, x] = \int_a \mu^Y(a, s, x)p(a|s, x; \sigma_A) = \mu^Y(\text{MLP}_{\sigma_A}^A(s, x), s, x)$.

We model $\mu^Y(a, s, x) = f(s, x) + g(a, s, x) + h(a, x)$ using a structured neural network NN^Y that separates into the three components f, g, h reflecting the decomposition of $\mu^Y(a, s, x)$. We learn the parameters of $\text{MLP}_{\sigma_A}^A$ and NN^Y in two phases, as outlined in figure 3a. In Phase I, we estimate the parameters of NN^Y from \mathcal{D} . In Phase II, we learn the parameters of $\text{MLP}_{\sigma_A}^A$ by optimizing objective (2) with $\mathbb{E}[Y_{\sigma_A}] \approx \frac{1}{N} \sum_{i=1}^N \mathbb{E}[Y_{\sigma_A} | s^i, x^i] = \frac{1}{N} \sum_{i=1}^N \mu^Y(\text{MLP}_{\sigma_A}^A(s^i, x^i), s^i, x^i)$, using the NN^Y trained in Phase I. If all variables are discrete, after estimating $\mu^Y(a, s, x)$ and $p(s, x)$, computing (2) can be cast as an LP (see Appendix C). Consistency is then a standard result which follows immediately. Notice that this formulation can also be extended to continuous S by defining the constraint via partial derivative, $\mathbb{E} \left[\left| \frac{\partial g_{\sigma_A}(s, X)}{\partial s} \right| \middle| S = s \right] \leq \epsilon$.

Action clipping. We suggest ensuring overlap of $p(a|s, x; \sigma_A)$ with $p(a|s, x; \emptyset)$ by adaptively constraining the output of $\text{MLP}_{\sigma_A}^A$ to be within an interval that resembles the one observed under the baseline policy. We consider two options: (a) matching the minimal and maximal value a seen for each s, x combination (binning continuous elements) and (b): extend that interval according to the difference between the minimal and maximal a value seen with each s, x . In the experiments, we opted for constraining the output of $\text{MLP}_{\sigma_A}^A$ to be in the interval

$$\left[\min_{A, S, X} -\eta \text{gap}_{A, S, X}, \max_{A, S, X} +\eta \text{gap}_{A, S, X} \right],$$

where $\text{gap}_{A, S, X} = \max_{A, S, X} - \min_{A, S, X}$. We enforced this interval with a shifted sigmoid function in the last layer of $\text{MLP}_{\sigma_A}^A$. We set $\eta = 1$ to allow some extrapolation and increase in utility.

4.2. EqB Constraint

For the EqB constrained objective (3), we require knowledge of the baseline policy as well as creating upper and lower bounds on the cdf $\mathcal{F}(Y_{\sigma_A} - Y_{\emptyset} | S = s), \forall s$. $Y_{\sigma_A} - Y_{\emptyset}$ is a counterfactual quantity, which is not generally identifiable. In fact, we do not have access to the joint distribution of Y_{σ_A} and Y_{\emptyset} as we never observed both variables simultaneously for the same individual. However, we show next that if one is willing to make a parametric assumption, then $Y_{\sigma_A} - Y_{\emptyset}$ is partially identifiable and amenable to bounding. We assume that the baseline policy is Gaussian and only consider Gaussian policies, with equal homogeneous variance, i.e. $p(A|s, x; \emptyset) = \mathcal{N}(\mu_{\emptyset}^A(s, x), V^A)$, $p(A|s, x; \sigma_A) = \mathcal{N}(\mu_{\sigma_A}^A(s, x), V^A)$. In addition, we assume that Y_{σ_A} and Y_{\emptyset} are jointly Gaussian conditioned on S, X , with marginal means $\mu_{\sigma_A}^Y(s, x)$, $\mu_{\emptyset}^Y(s, x)$ and equal homogenous marginal variance V^Y .

The assumption on Y_{σ_A} and Y_{\emptyset} enables us to bound the population cdf by maximizing/minimizing it with respect to the unknown correlation coefficient $\rho(s, x) \in [-1, 1]$, where $\rho(s, x) := \frac{\text{Cov}(Y_{\sigma_A}, Y_{\emptyset} | s, x)}{V^Y}$. Let Φ denote the cdf of the standard Gaussian. We can write

$$\mathbb{P}(Y_{\sigma_A} - Y_{\emptyset} \leq z | x, s) = \Phi \left(\frac{z - \mu_{\sigma_A}^Y(s, x) + \mu_{\emptyset}^Y(s, x)}{\sqrt{2V^Y(1 - \rho(s, x))}} \right).$$

Defining $\mu(s, x) := \mu_{\sigma_A}^Y(s, x) - \mu_{\emptyset}^Y(s, x)$ and $f_{s,x}^z(\rho) := \Phi\left(\frac{z - \mu(s, x)}{\sqrt{2V^Y(1-\rho)}}\right)$, we can show that, for any s, x , the following bounds are the tightest: (i) $f_{s,x}^z(-1) \leq \mathbb{P}(Y_{\sigma_A} - Y_{\emptyset} \leq z | x, s) \leq f_{s,x}^z(1)$, for $z - \mu(s, x) > 0$; (ii) $f_{s,x}^z(1) \leq \mathbb{P}(Y_{\sigma_A} - Y_{\emptyset} \leq z | x, s) \leq f_{s,x}^z(-1)$, for $z - \mu(s, x) < 0$. Using N_s to indicate the number of elements in \mathcal{D} with $s^i = s$, we obtain global lower and upper bounds estimates for $\mathbb{P}(Y_{\sigma_A} - Y_{\emptyset} \leq z | s)$ as $F_s^L(z) = \frac{1}{N_s} \sum_{i:s^i=s} f_{s,x^i}^z(-\text{sign}(z - \mu(s, x^i)))$ and $F_s^U(z) = \frac{1}{N_s} \sum_{i:s^i=s} f_{s,x^i}^z(\text{sign}(z - \mu(s, x^i)))$.

We operationalize the constraint by minimizing the mean squared error (MSE) of the bounds differences, i.e. our final objective is

$$\arg \max_{\sigma_A} \mathbb{E}[Y_{\sigma_A}] \quad \text{s.t.} \quad \sum_{z \in S_z} \left(\|F_s^L(z) - F_{\bar{s}}^L(z)\|_2^2 + \|F_s^U(z) - F_{\bar{s}}^U(z)\|_2^2 \right) \leq \epsilon, \quad (4)$$

$\forall s, \bar{s}$, where S_z is a grid of values.

We propose to estimate $\mathbb{E}[Y_{\sigma_A}]$ with the inverse probability weighting (IPW) estimator. This is an alternative estimation approach to the one used for `ModBrk`, to demonstrate the modularity of our approach and constraints.

$$\mathbb{E}[Y_{\sigma_A}] = \int_{a,s,x,y} y \frac{p(a | s, x; \sigma_A)}{p(a | s, x; \emptyset)} p(a, s, x, y; \emptyset) \approx \frac{1}{N} \sum_{i=1}^N y^i \frac{p(a^i | s^i, x^i; \sigma_A)}{p(a^i | s^i, x^i; \emptyset)}.$$

We model $\mu_{\emptyset}^A(s, x)$, $\mu_{\sigma_A}^A(s, x)$ and $\mu^Y(a, s, x) = \mathbb{E}[Y | a, s, x]$ using MLP neural networks MLP_{\emptyset}^A , $\text{MLP}_{\sigma_A}^A$, and MLP^Y respectively. We learn the parameters of these networks in two phases as outlined in figure 3b. In Phase I, we learn the parameters of MLP_{\emptyset}^A and MLP^Y from \mathcal{D} with MSE losses between predicted and observed actions and outcome. We obtain a MAP estimate of $\mu_{\emptyset}^Y(s, x) = \int_a \mu^Y(a, s, x) p(a | s, x, \emptyset)$ as $\hat{\mu}_{\emptyset}^Y(s, x) = \int_a \mu^Y(a, s, x) \delta_{A=\mu_{\emptyset}^A(s, x)}$. We estimate V^A and V^Y through averaging the MSE of target and predicted mean output from MLP^A and MLP^Y respectively². In Phase II, we learn the parameters of $\text{MLP}_{\sigma_A}^A$ that maximize objective (4) using the MLP_{\emptyset}^A and MLP^Y trained in Phase I.

Misspecification and parametric restrictiveness. Notice that, as the Gaussianity assumption is placed on the conditional joint distribution of $(Y_{\sigma_A}, Y_{\emptyset})$ given S and X rather than on the marginal distribution. It is thus reasonable to invoke the central limit theorem and assume that, given enough samples, the residuals of $Y_{\emptyset} | s, x$ and $Y_{\sigma_A} | s, x$ are normally distributed, since one can view the residuals as aggregating various minor factors that explain the remaining variability after taking into account a strong signal. Also notice that our proposed approach can be easily extended to heteroskedastic Gaussian regression models which are very flexible and can accommodate many real-world datasets. Finally, a nonparametric approach for the constraint computation would be possible e.g. via Fréchet bound (see Appendix D), but that may come at a computational cost. To illustrate the effects of misspecification on our proposed assumption, we present below a formal result showing that our constraint estimation error is bounded from above by how the true joint density deviates from the Gaussian density (see Appendix D for a proof and discussion).

2. We assume V^Y to be homoskedastic.

Proposition 1 For fixed s, x , let $Q = \mathbb{P}(Y_{\sigma_A} - Y_{\emptyset} \leq z \mid s, x)$ and our estimator

$$\hat{Q} = \Phi \left(\frac{z - \mu_{\sigma_A}^Y(s, x) + \mu_{\emptyset}^Y(s, x)}{\sqrt{2V^Y(1 - \rho(s, x))}} \right).$$

Let $f_{s,x}$ denote the joint density function of $Y_{\sigma_A}, Y_{\emptyset} \mid s, x$ and $\phi_{s,x}$ the joint Gaussian density function with mean $[\mu_{\sigma_A}^Y(s, x), \mu_{\emptyset}^Y(s, x)]$ and shared variance V with correlation $\rho(s, x)$. Then $|Q - \hat{Q}| \leq \|f_{s,x} - \phi_{s,x}\|_1$, where $\|\cdot\|_1$ denotes the L_1 norm.

5. Experiments

We evaluated the `ModBrk` and `EqB` constraint methods³ on the New York City Public School District (NYCSchools) dataset compiled in [Kusner et al. \(2019\)](#), which we augmented to include actions; and further tested `ModBrk` on the Infant Health and Development Program (IHDP) dataset, specifically the real-data example examining dosage effects described in Section 6.2 of [Hill \(2011\)](#) (we could not use this dataset for `EqB` due to its reliance on parametric assumptions). Below, we briefly discuss the datasets and defer further details to Appendix E.

Action-augmented NYCSchools dataset. We adopted the same sensitive attribute and covariates as [Kusner et al. \(2019\)](#), and augmented the dataset with generated actions and outcomes. We created continuous actions corresponding to funding level decisions as $A = (w_{SX}^T SX)^2 + \max(0, w_X^T X) + \mathcal{N}(0.5, 0.4)$, where E is the original percent of students taking the SAT/ACT exams (pre-college entry) in the dataset. The proposed form of Y is discussed in Appendix E.

IHDP dataset. The IHDP dataset describes a program that targeted low-birth-weight, premature infants, providing them with intensive high-quality child care and home visits from a trained provider. As continuous action A we used the self-selected number of participation days in the program. The outcome Y corresponds to child score attainment in cognitive tests at age three. We considered mother’s race (white vs. non-white) as sensitive attribute S . We reinterpreted the original setting as a resource allocation problems as follows: rather than as a self-decision, we view the number of days in treatment as external, e.g., by assigning different individuals to different lengths of participation. In this case, the `ModBrk` constraint goal is to break the moderation of the allocation of days in program by the group membership, such that the resource allocation policy is not responsible for the difference in scores attained by different groups.

5.1. Results

We evaluated `ModBrk` and `EqB` by comparing these methods with: (1) optimizing the policy with no disparity constraint (Unconstrained), (2) optimizing the policy without using S (Drop S), and (3) using the baseline policy ($\sigma_A = \emptyset$). As discussed in Section 2, previous work on fair policy or action allocation differs from our setting and objective, and therefore cannot be meaningfully compared. For `ModBrk`, we also compare against different levels of constant actions (`ConstA`). We do not include it for `EqB` since the IPW estimator is not suitable for δ -distributions. The range of ϵ values in the plots was determined such that the minimal ϵ slack values correspond to the smallest constraint value achievable with our optimization strategy (this is the lowest possible constraint value achieved when setting $\epsilon = 0$); the maximal ϵ slack value is the one that would closely match

3. The code reproducing our results is available at <https://github.com/limorigu/PragmaticFairness>.

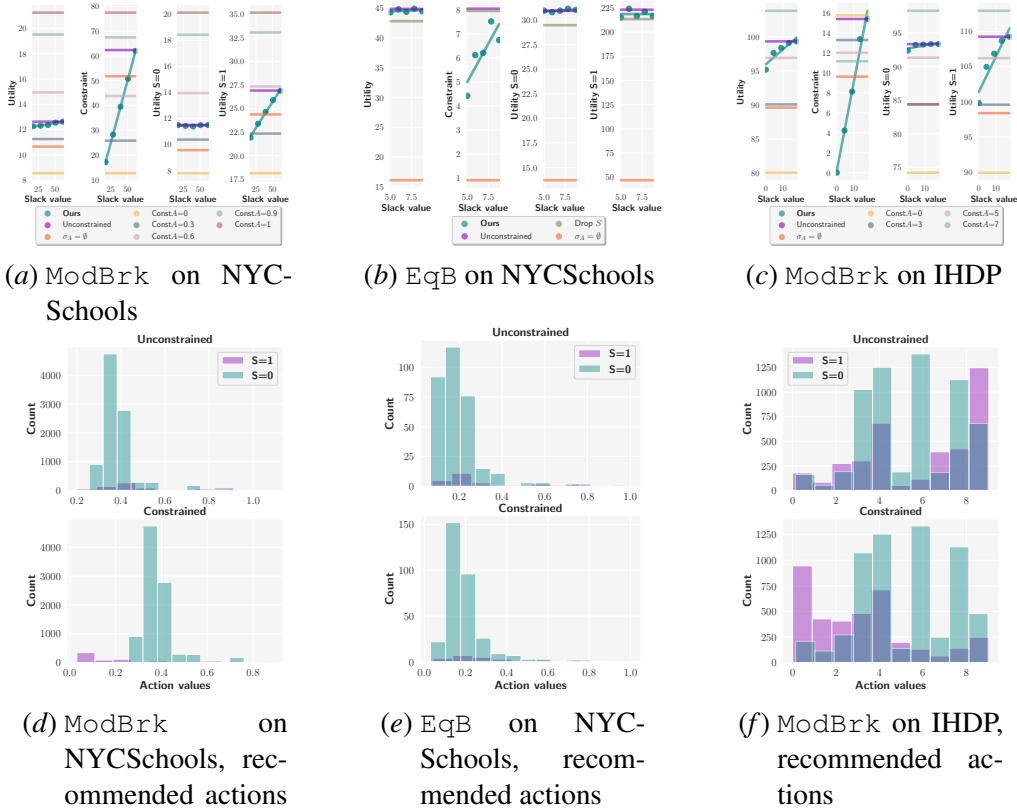


Figure 4: (Top) influence of the slack value on obtained values of the objective function and fairness constraint violations. (Bottom) recommended actions by group for unconstrained optimization and for constrained optimization, where the slack on fairness violation set to 0 or the highest level shown in the top plot, respectively.

the constraint violation under unconstrained optimization. In practice, we propose to explore the frontier resulting from choices of slack values and select the most acceptable trade-off for the user.

ModBrk Constraint Method. The results for ModBrk on the NYCSchools and IHDP datasets are presented in Fig. 4(a), 4(d) and in Figs. 4(c), 4(f) respectively⁴. For both datasets, as we increase the tolerance on fairness violations, in the form of higher slack value ϵ in (2), we see a major increase in constraint value, while allowing some increase in utility $\mathbb{E}[Y_{\sigma_A}]$. Observing the utility values broken down by group membership ($\mathbb{E}[Y_{\sigma_A} | S = 1]$ vs. $\mathbb{E}[Y_{\sigma_A} | S = 0]$), we see most of this increase of utility is driven by the privileged group, $S = 1$ (a majority-white student body in the NYCSchools dataset, and white mothers for the IHDP dataset). This is to be expected given that we are trying to minimize interactions involving $S = 1$ and A , i.e. where the membership in the privileged group inflates utility values, via higher g values. These higher g values for the $S = 1$ group also translate into higher utility values for $S = 1$ according the baseline policy, as can be

4. Dropping S had no influence here on utility and constraint values, thus we removed this baseline to avoid confusion due to overlap and subsequent mismatch of colors between figure and legend.

seen from the orange line, indicating baseline policy actions in the corresponding figures. Notice that the choice of an appropriate slack ϵ here is a matter of trade-off based on user preferences of utility vs. constraint value, and will depend on the specific dataset. We can gain deeper insight into the working of our approach by inspecting the histogram of recommended actions at the end of the policy model training. For both datasets, we see that applying the constraint with a small ϵ value results in mapping more $S = 1$ members to lower action values compared to the $S = 0$ group, indicating that to enforce the constraint more tightly means allocating lower value actions to the privileged group, while giving out as many high actions to $S = 0$ as possible under the action clipping setting described previously. Notice that in both settings we succeed in increasing the utility compared to the baseline policy. Recall that we employ action clipping to ensure some overlap with the baseline actions for estimation purposes. This means that we are bound to some extent to the interval of actions seen for each s, x in the baseline policy (this trade-off can be explored via choice of η value, see Section 4.1). One could also increase η to achieve higher utility at the cost of increased total action allocation (i.e. higher “budget”) and sacrifice in coverage. This is also why we cannot achieve utility values that are as high as the highest Const_A in the figures; we ask how best to distribute actions within the effective “budget”.

EqB Constraint Method. We present results for EqB on the NYC Schools dataset in Figs. 4(b) and 4(e). As we increase the tolerance on fairness violations - in the form of higher slack value ϵ in (4) - we observe an increase in constraint value with almost no change in overall utility $\mathbb{E}[Y_{\sigma_A}]$. Observing the utility values broken down by group membership ($\mathbb{E}[Y_{\sigma_A} | S = 1]$ vs. $\mathbb{E}[Y_{\sigma_A} | S = 0]$), we see also no significant change in utility coming from either group. This result indicates that our method learns a policy that assigns actions ensuring equal benefit without sacrificing the utility of either group. Note that we also see an unusually high estimate of utility for $S = 1$. This could be explained by the IPW estimator not extrapolating well to unseen data, as the $S = 1$ group only consists of 7% of an already small dataset. On the right hand side of Fig. 4(b) we observe similar recommended action histograms for the unconstrained (top right) and constrained (bottom right) cases. This shows that our method decreases disparity with a similar “budget”. This is possible because we are not attempting to reduce disparity present under the baseline policy. The similarity in action histograms is partially due to the IPW objective which aligns recommended actions with observed ones to achieve higher weighting. To validate our approach, we compute the ground truth constraint value through counterfactual realization of the policy model’s predicted action mean in the data generating process and compute distributional difference of $Y_{\sigma_A} - Y_{\emptyset} | S$ for the different sensitive attribute groups. Notice also that our method succeeds in increasing the utility compared to the baseline policy. Although the baseline policy has the lowest constraint, this is expected as we are comparing distributions of $Y_{\sigma_A} - Y_{\emptyset} | S$, where $\sigma_A = \emptyset$. Dropping S has a slight increase in constraint and a decrease in utility compared to unconstrained setting, as our policy model performs better through taking into account S to break the indirect association between S and Y .

6. Conclusion

We introduced a causal framework to learn fair policies given access to observational data and an action space. Taking a pragmatic view, we asked what is the best utility that can be achieved with the provided action space, while controlling two notions of disparity: one focusing on mitigating a possible moderation effect involving group membership and the policy, and the other focusing on ensuring equal benefit with respect to a baseline policy. We see this work as a first conceptual

contribution in defining pragmatic fair impact policies, and envisage various possible future directions, including extending the proposed methods beyond binary sensitive attributes, to a multi-stage policy setting, to handle unmeasured confounding and to online optimization.

Limitations discussion. *Conceptual.* Our pragmatic approach relies on data that we have about mechanisms *in this world*, including possible inequities our actions in question cannot affect. We also do so within a stationary framework that currently does not account for feedback, although we see it as an important next step. This comes with caveats. First, whether or not it is desirable to control for disparities of outcomes across levels of the sensitive attribute is problem-dependent. For instance, if the action space consists of solely two options, to give a medical treatment or not, it may be unclear why we should take into consideration group differences in recovery: other things being equal, we just want to maximize the number of lives saved. In contrast, if Y is a relative measure, e.g., of wealth, pure wealth maximization may be judged to be harmful if disparities among groups in S are exacerbated. In this case, we may settle for a scenario with less aggregated wealth if disparities are controlled. Such value judgements are *not* to be decided algorithmically. Our goal is to provide a formalization of disparity control *if* it is judged to be desirable, and to provide an estimate of *whether* different levels of control can be achieved under an acceptable loss of total expected outcome, *given* an action space that is, again, a property of the real-world. We simply provide a framework to examine what is possible. However, if our data reflects biased mechanisms existing in the world that is outside of reach for our actions, we will not be able to change those, and they will be reflected in estimates of what is possible. This is in contrast to methods that aim to model alternative fair worlds.

Technical. The parametric assumption we make for the operationalization of the EqB constraint is one that could be avoided and extended for greater generalization (e.g., via Fréchet bounds, see Appendix D.1). However, in this work we opted for a parametric assumption to focus on a general conceptual introduction. For the ModBrk constraint we propose a simple MLP estimation of the decomposition in Eq. 1. One could envision a more elaborate formulation that will avoid a possible failure mode, where g simply subsumes f and h . We did not observe empirically that this happens in practice. However, one could avoid such possible issue by including additional regularization or residual connections between the components of the MLP.

ACKNOWLEDGMENTS

The authors would like to thank Anian Ruoss and Jessica Schrouff for their helpful feedback on the manuscript, as well as the Alan Turing Institute for providing access to computational resources. RS was partially funded by the EPSRC grant EP/W024330/1.

REFERENCES

- Abien Fred Agarap. Deep learning using rectified linear units (relu). *arXiv preprint arXiv:1803.08375*, 2018.
- Solon Barocas, Moritz Hardt, and Arvind Narayanan. *Fairness and Machine Learning*. fairml-book.org, 2019. <http://www.fairmlbook.org>.
- Alex Chohlas-Wood, Madison Coots, Emma Brunskill, and Sharad Goel. Learning to be fair: A consequentialist approach to equitable decision-making. *CoRR*, abs/2109.08792, 2021.

- Juan Correa and Elias Bareinboim. A calculus for stochastic interventions: Causal effect identification and surrogate experiments. In *AAAI Conference on Artificial Intelligence*, pages 10093–10100, 2020.
- Alexander D’Amour, Hansa Srinivasan, James Atwood, Pallavi Baljekar, D Sculley, and Yoni Halpern. Fairness is not static: deeper understanding of long term fairness via simulation studies. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, pages 525–534, 2020.
- Alexander Philip Dawid. Fundamentals of statistical causality. Technical report, University of Sheffield, 2007.
- Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard S. Zemel. Fairness through awareness. In *Innovations in Theoretical Computer Science Conference*, pages 214 – 226, 2012.
- Stephen Gillen, Christopher Jung, Michael Kearns, and Aaron Roth. Online learning with an unknown fairness metric. In *Advances in Neural Information Processing Systems*, 2019.
- Jennifer L. Hill. Bayesian nonparametric modeling for causal inference. *Journal of Computational and Graphical Statistics*, 20(1):217–240, 2011.
- Wen Huang, Lu Zhang, and Xintao Wu. Achieving counterfactual fairness for causal bandit, 2021.
- John W. Jackson. On the interpretation of path-specific effects in health disparities research. *Epidemiology*, 29(4):517–520, 2018.
- John W. Jackson. Meaningful causal decompositions in health equity research: definition, identification, and estimation through a weighting framework, 2020.
- John W. Jackson and Tyler J. VanderWeele. Decomposition analysis to identify intervention targets for reducing disparities. *Epidemiology (Cambridge, Mass.)*, 29(6):825, 2018.
- Matthew Joseph, Michael Kearns, Jamie Morgenstern, and Aaron Roth. Fairness in learning: Classic and contextual bandits. *arXiv preprint arXiv:1605.07139*, 2016.
- Matthew Joseph, Michael Kearns, Jamie Morgenstern, Seth Neel, and Aaron Roth. Meritocratic fairness for infinite and contextual bandits. In *AAAI/ACM Conference on AI, Ethics, and Society*, pages 158–163, 2018.
- Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization, 2017.
- Matt Kusner, Chris Russell, Joshua Loftus, and Ricardo Silva. Making decisions that reduce discriminatory impacts. In *International Conference on Machine Learning*, pages 3591–3600, 2019.
- Matt J. Kusner, Joshua Loftus, Chris Russell, and Ricardo Silva. Counterfactual fairness. In *Advances in Neural Information Processing Systems*, pages 4069–4079, 2017.
- Lydia T. Liu, Sarah Dean, Esther Rolf, Max Simchowitz, and Moritz Hardt. Delayed impact of fair machine learning. In *International Conference on Machine Learning*, pages 3150–3158, 2018.

- G. D. Makarov. Estimates for the Distribution Function of a Sum of Two Random Variables When the Marginal Distributions are Fixed. *Theory of Probability & Its Applications*, 26(4):803–806, 1982.
- Ninareh Mehrabi, Fred Morstatter, Nripsuta Saxena, Kristina Lerman, and Aram Galstyan. A survey on bias and fairness in machine learning, 2019. URL <https://arxiv.org/abs/1908.09635>.
- Luke Miratrix, Jane Furey, Avi Feller, Todd Grindal, and Lindsay C Page. Bounding, an accessible method for estimating principal causal effects, examined and explained. *Journal of Research on Educational Effectiveness*, 11(1):133–162, 2018.
- Shira Mitchell, Eric Potash, Solon Barocas, Alexander D’Amour, and Kristian Lum. Algorithmic fairness: Choices, assumptions, and definitions. *Annual Review of Statistics and Its Application*, 8(1):141–163, 2021. doi: 10.1146/annurev-statistics-042720-125902. URL <https://doi.org/10.1146/annurev-statistics-042720-125902>.
- Razieh Nabi, Daniel Malinsky, and Ilya Shpitser. Learning optimal fair policies. In *International Conference on Machine Learning*, pages 4674–4682, 2019.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems*, pages 8024–8035. 2019.
- Vishakha Patil, Ganesh Ghalme, Vineet Nair, and Y. Narahari. Achieving fairness in the stochastic multi-armed bandit problem, 2020.
- Judea Pearl. *Causality: Models, Reasoning, and Inference*. Cambridge University Press, 2000.
- Dana Pessach and Erez Shmueli. A review on fairness in machine learning. *ACM Comput. Surv.*, 55(3), feb 2022. ISSN 0360-0300. doi: 10.1145/3494672. URL <https://doi.org/10.1145/3494672>.
- Candice Schumann, Zhi Lang, Nicholas Mattei, and John P. Dickerson. Group fairness in bandit arm selection, 2022.
- Sahil Verma and Julia Rubin. Fairness definitions explained. In *2018 IEEE/ACM International Workshop on Software Fairness (Fairware)*, pages 1–7. IEEE, 2018.
- Sandra Wachter, Brent Mittelstadt, and Chris Russell. Bias preservation in machine learning: the legality of fairness metrics under eu non-discrimination law. *W. Va. L. Rev.*, 123:735, 2020.
- Yongkai Wu, Lu Zhang, and Xintao Wu. Counterfactual fairness: Unidentification, bound and algorithm. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence*, 2019.

Appendix A. Further discussion of related works

Considering interventional data in the online setting, [Joseph et al. \(2016, 2018\)](#) derive algorithms that have “meritocratic fairness,” requiring that, at every round, no arm with lower expected reward is preferred to one with a higher expected reward. This goal can be seen as an individual-level fairness metric, but it is not at odds with a policy that maximizes utility; instead, its main concern is to offer more conservative exploration.

Recently, fairness in bandit settings has received a lot of additional attention ([Huang et al., 2021](#); [Schumann et al., 2022](#); [Patil et al., 2020](#)). We see interventional data and online sequential settings as an interesting future direction, especially appropriate for use cases with low individual impact, such as software and internet applications that interact with users naturally and frequently, at low participation cost. However, we focus on observational-data regimes which we see as often more appropriate, especially when the collection of interventional data could be harmful, or when it would be unfair for individuals to essentially pay the price for exploration and calibration of learning agents.

Appendix B. Augmented Lagrangian penalties

We can write the constrained problem as $\min_x \max_{\lambda \geq 0} f(x) + \lambda(c(x) - b)$ (possibly with a flipped sign, depending on the optimization problem specifics). f is our policy optimization objective, c stands for the computed constraint, and b is the level we wish to hold it at. The max returns $f(x)$ when x satisfies the constraints (as the maximum is obtained at $\lambda = 0$), and ∞ otherwise (as the maximum is at $\lambda = \infty$). There could be a difficulty in optimizing the form above directly, as λ jumps from 0 to ∞ when passing through the constraint boundary. To fix this we add a penalty on making large changes to λ , obtaining $\min_x \max_{\lambda \geq 0} f(x) + \lambda(c(x) - b) - 1/(2\mu\|\lambda - \lambda'\|_2)$, where λ' are the Lagrange multipliers from the previous iteration and μ is a penalty term iteratively increased.

The inner max can be solved in closed form as

$$\lambda = \begin{cases} 0 & \text{if } k(x) \leq -\frac{\lambda'}{\mu} \\ \lambda' + \mu k(x) & \text{otherwise} \end{cases} \quad (5)$$

where $k(x) := (c(x) - b)$, representing the inequality constraint. Plugging the above into a complete optimization problem,

$$\begin{aligned} & \min_x f(x) + \phi(k(x), \lambda', \mu) \\ \text{s.t. } & \phi(k(x), \lambda', \mu) = \begin{cases} -\frac{\lambda'^2}{2\mu} & \text{if } k(x) \leq -\frac{\lambda'}{\mu} \\ \lambda'k(x) + \frac{\mu k(x)^2}{2} & \text{otherwise} \end{cases} \end{aligned} \quad (6)$$

This suggests a straight-forward algorithm: at each step, minimize the above, update the value of λ via [Eq. 5](#), and repeat until convergence.

Appendix C. The ModBrk constraint as an LP

In this section we show how the optimization problem in Eq. 2 can be cast as a Linear Program (LP) for the common case where the variables are discrete. We decouple the estimation problem and assume that $\mu^Y(a, s, x)$ and $p(s, x)$ are known or appropriately estimated.

First, consider the case when A is binary; then, the policy may be identified by $\mu_{\sigma_A}^A(s, x) = p(A = 1|S = s, X = x; \sigma_A)$. For convenience, define $\delta(s, x) = \mu^Y(1, s, x) - \mu^Y(0, s, x)$. Then, we may evaluate

$$\begin{aligned} \mathbb{E}[Y_\pi] &= \sum_{s,x} p(s, x) (\mu^Y(1, s, x)\mu_{\sigma_A}^A(s, x) + \mu^Y(0, s, x)(1 - \mu_{\sigma_A}^A(s, x))) = \\ &\quad \sum_{s,x} p(s, x)(\mu^Y(0, s, x) + \delta(s, x)\mu_{\sigma_A}^A(s, x)), \end{aligned}$$

which simplifies the optimization problem to

$$\begin{aligned} \max_{\sigma_A} \quad & \sum_{s,x} p(s, x)\delta(s, x)\mu_{\sigma_A}^A(s, x) \\ \text{s.t.} \quad & \sum_x (p(s, x)\delta(s, x)\mu_{\sigma_A}^A(s, x) - p(s', x)\delta(s', x)\mu_{\sigma_A}^A(s', x)) \leq \sqrt{\epsilon} \quad \forall s, s' < s, \text{ and} \\ & \sum_x (p(s, x)\delta(s, x)\mu_{\sigma_A}^A(s, x) - p(s', x)\delta(s', x)\mu_{\sigma_A}^A(s', x)) \geq -\sqrt{\epsilon} \quad \forall s, s' < s, \end{aligned}$$

which is a linear program in σ_A with a linear objective and linear constraints.

To generalize to the non-binary case, we define, by a slight abuse of notation, $\mu_{\sigma_A}^A(a, s, x) = p(A = a|S = s, X = x; \sigma_A)$ and evaluate the objective as

$$\mathbb{E}[Y_\pi] = \sum_{s,x,a} p(s, x)\mu^Y(a, s, x)\mu_{\sigma_A}^A(a, s, x),$$

leading to the linear program

$$\begin{aligned} \max_{\mu_{\sigma_A}^A(a,s,x)} \quad & \sum_{s,x,a} p(s, x)\mu^Y(a, s, x)\mu_{\sigma_A}^A(a, s, x) \\ \text{s.t.} \quad & \sum_{a,x} p(s, x) (\mu^Y(a, s, x)\mu_{\sigma_A}^A(a, s, x) - p(s, x)\mu^Y(a, s, x)\mu_{\sigma_A}^A(a, s, x)) \\ & \quad - \sum_{a,x} p(s', x) (\mu^Y(a, s', x)\mu_{\sigma_A}^A(a, s', x) - p(s, x)\mu^Y(a, s', x)\mu_{\sigma_A}^A(a, s', x)) \leq \sqrt{\epsilon} \quad \forall s, s' < s, \\ & \sum_{a,x} p(s, x) (\mu^Y(a, s, x)\mu_{\sigma_A}^A(a, s, x) - p(s, x)\mu^Y(a, s, x)\mu_{\sigma_A}^A(a, s, x)) \\ & \quad - \sum_{a,x} p(s', x) (\mu^Y(a, s', x)\mu_{\sigma_A}^A(a, s', x) - p(s, x)\mu^Y(a, s', x)\mu_{\sigma_A}^A(a, s', x)) \geq -\sqrt{\epsilon} \quad \forall s, s' < s, \\ & \sum_a \mu_{\sigma_A}^A(a, s, x) = 1 \quad \forall s, x, \quad \text{and} \quad \mu_{\sigma_A}^A(a, s, x) \geq 0 \quad \forall a, s, x. \end{aligned}$$

The last two constraints ensure that the policy is a valid conditional probability distribution.

Appendix D. The EqB constraint formulation

We first formulate the EqB constraint under a Gaussian assumption primarily for simplicity of exposition. We derive a conservative bound in section 4.2 to guarantee that $F_s^L(z) \leq F_s^{true}(z) \leq F_s^U(z)$, where $F_s^{true}(z) = \mathbb{P}(Y_{\sigma_A} - Y_{\emptyset} \leq z | s)$. Note that no tighter bound exists without further assumptions. This is a result that follows directly from the properties of a Gaussian distribution: that is, for a bivariate Gaussian (X, Y) with fixed marginals, one gets the smallest/largest value of the cdf $\mathbb{P}(X - Y \leq z)$ among all possible correlation coefficients when their correlation is -1 or 1 (depending on z and the marginal means). This can be directly verified in the equation following the second paragraph in Section 4.2 by simple differentiation. Thus we know the global bound, $F_s^U(z) - F_s^L(z)$.

Relying on such-tight-as-possible bounds, equating the upper and lower bounds across groups will indeed ensure similar constraints, and is the best we can do if we do not want to deal with the quality of approximation itself, but rather enforce a similarity between approximations across groups. In other words, the goal is to make the distribution of $D \equiv Y_{\sigma_A} - Y_{\emptyset}$ as independent as possible of the values of S , meaning any pair s and \bar{s} should imply the same conditional distribution for D (up to an agreeable tolerance level ϵ). Ideally, if the conditional distribution of D was identified, the constraint would be

$$|\mathcal{P}(D \leq z | s) - \mathcal{P}(D \leq z | \bar{s})| \leq \epsilon$$

As these distributions are *not* identified, we apply the principle of aiming at *not having any evidence against the lack of equality between these distributions*. This means making *the set of all possible $\mathcal{P}(D \leq z | s)$ and $\mathcal{P}(D \leq z | \bar{s})$ the same for all z* . As the smallest of such sets are given by intervals, this is equivalent to making the upper bounds the same and the lower bounds the same, *not* mixing upper and lower bounds:

$$\begin{aligned} |\mathcal{P}^U(D \leq z | s) - \mathcal{P}^U(D \leq z | \bar{s})| &\leq \epsilon \\ |\mathcal{P}^L(D \leq z | s) - \mathcal{P}^L(D \leq z | \bar{s})| &\leq \epsilon \end{aligned}$$

Finally, in order to reduce the number of possible Lagrange multipliers, we use the norm in Eq. 4 as an alternative to enforcing all of the constraints separately.

For extensions beyond the normality assumption, there are very standard worst-case tight bounds based on standard copula theory. We describe those in the following.

D.1. Fréchet bound (Makarov, 1982)

Consider a random vector (X, Y) with continuous marginal CDF $F_X(x) = \mathbb{P}(X \leq x)$, $F_Y(y) = \mathbb{P}(Y \leq y)$. The copula of (X, Y) is

$$C(u, v) = \mathbb{P}(X \leq F_X^{-1}(u), Y \leq F_Y^{-1}(v))$$

The Fréchet-Hoeffding Theorem states that

$$\max\{u + v - 1, 0\} \leq C(u, v) \leq \min\{u, v\}$$

Thus we can bound the unknown joint distribution with estimable marginal distributions. However, we note that the benefit of this nonparametric approach comes at a computational cost, since we would need to estimate the CDF of marginal distributions. Additionally, we may end up with loose bounds, for example, when marginal probabilities take values around 0.5.

D.2. Proof of model misspecification

Lemma 1 *The difference distribution $W := X - Y$ for two (possibly dependent) random variables has density f_W taking the form*

$$f_W(w) = \int_{-\infty}^{\infty} f_{X,Y}(x, x-w) dx \quad (7)$$

where $f_{X,Y}(x, y)$ is the joint density function of X, Y .

Proof We can write the cumulative density function of the difference $X - Y$ as the following:

$$F_W(w) := \mathbb{P}(W \leq w) = \mathbb{P}(X - Y \leq w) = \int_{-\infty}^{\infty} \int_{-\infty}^{w+y} f_{X,Y}(x, y) dx dy$$

Note the density function is the derivative of the CDF. Thus, we have:

$$\begin{aligned} f_W(w) &= \frac{d}{dw} F_W(w) = \frac{d}{dw} \int_{-\infty}^{\infty} \int_{-\infty}^{w+y} f_{X,Y}(x, y) dx dy \\ &= \int_{-\infty}^{\infty} \frac{d}{dw} \int_{-\infty}^w f_{X,Y}(t+y, y) dt dy \quad (\text{change of variable } x = t+y) \\ &= \int_{-\infty}^{\infty} f_{X,Y}(w+y, y) dy \\ &= \int_{-\infty}^{\infty} f_{X,Y}(x, x-w) dx \quad (\text{change of variable } x = w+y) \end{aligned}$$

■

Lemma 2 *Let $W := X - Y$ and $W' := X' - Y'$ where $f_{X,Y}$ is the joint density function of X, Y and $\phi_{X,Y}$ is the joint density function of X', Y' . Then $\forall w$,*

$$|\mathbb{P}(W \leq w) - \mathbb{P}(W' \leq w)| \leq \|f_{X,Y} - \phi_{X,Y}\|_1 \quad (8)$$

Proof Let f_W be the density function of W and ϕ_W be the density function of W' .

$$\begin{aligned} |\mathbb{P}(W \leq w) - \mathbb{P}(W' \leq w)| &= \left| \int_{-\infty}^w f_W(z) - \phi_W(z) dz \right| \\ &= \left| \int_{-\infty}^w \int_{-\infty}^{\infty} f_{X,Y}(x, x-z) - \phi_{X,Y}(x, x-z) dx dz \right| \\ &= \left| \int_{-\infty}^{\infty} \int_{x-w}^{\infty} f_{X,Y}(x, y) - \phi_{X,Y}(x, y) dy dx \right| \\ &\leq \int_{-\infty}^{\infty} \int_{x-w}^{\infty} |f_{X,Y}(x, y) - \phi_{X,Y}(x, y)| dy dx \\ &\leq \|f_{X,Y} - \phi_{X,Y}\|_1, \forall w \end{aligned}$$

■

Proposition 1 (Model Misspecification) For fixed s, x , let $Q = \mathbb{P}(Y_{\sigma_A} - Y_{\emptyset} \leq z | s, x)$ and our estimator $\hat{Q} = \Phi\left(\frac{z - \mu_{\sigma_A}^Y(s, x) + \mu_{\emptyset}^Y(s, x)}{\sqrt{2V^Y(1 - \rho(s, x))}}\right)$. Let $f_{s, x}$ denote the joint density function of $Y_{\sigma_A}, Y_{\emptyset} | s, x$ and $\phi_{s, x}$ denote the joint Gaussian density function with mean $[\mu_{\sigma_A}^Y(s, x), \mu_{\emptyset}^Y(s, x)]$ and shared variance V with correlation $\rho(s, x)$. Then

$$\left| Q - \hat{Q} \right| \leq \|f_{s, x} - \phi_{s, x}\|_1, \quad (9)$$

where $\|\cdot\|_1$ denotes the L_1 norm.

Informally speaking, the accuracy of our estimator is bounded by the L_1 norm between our Gaussian approximated joint density function and the real underlying joint density function of observed treatment effect and counterfactual treatment effect conditioned on individual circumstance (s, x) .

Proof We can think of $\phi_{s, x}$ as a jointly normal approximation to the true underlying joint density $f_{s, x}$ of observed treatment effect and counterfactual treatment effect conditioned on individual circumstance (s, x) . Then by Lemma 2, we have

$$\left| \mathbb{P}(Y_{\sigma_A} - Y_{\emptyset} \leq z | s, x) - \Phi\left(\frac{z - \mu_{\sigma_A}^Y(s, x) + \mu_{\emptyset}^Y(s, x)}{\sqrt{2V^Y(1 - \rho(s, x))}}\right) \right| \leq \|f_{s, x} - \phi_{s, x}\|_1.$$

■

Appendix E. Additional dataset details

To the best of our knowledge, both datasets were made public under CC0 license.

Action-augmented NYC Schools dataset. As sensitive attribute S , we consider the majority race of students in the school (white vs. non-white), while as covariates X we consider: (1) if the school offers advanced placement or international baccalaureate classes, (2) whether the school offers calculus courses, and (3) the number of full-time counselors employed (fractional values indicate part-time work). We then create continuous actions corresponding to funding level decisions as $A = (w_{SX}^T SX)^2 + \max(0, w_X^T X) + \mathcal{N}(0.5, \sigma)$, where w_{SX} , w_X are sampled uniformly in $[0, 1]$, and $\sigma = 0.4$ (to allow variability, but not encourage negative values, given that the mean is set at 0.5). We then scale A to lie in $[0, 1]$. We generate the outcome Y that represents the percent of students who take the college entrance exams in the US (SAT/ACT), as $Y = 20E + \beta^T [S, X, SX, A, XA, SA, SAX] + \gamma^T [X, A, AX] + \mathcal{N}(1, \sigma)$, where E is the original percent of students taking the SAT/ACT exams (pre-college entry) appearing in the dataset. We multiply E by 20 to bring it within the same scale as the other components in the structural equation. β and γ are sampled from a normal distribution, with higher mean values for coefficients affecting terms involving A , to ensure the outcome is noticeably responsive to changes in A . In the original dataset, E is the outcome and therefore can be interpreted as a function of S and X .

The form of Y was picked as a simple demonstration of one that would include the original Y 's in the original dataset, as well as interaction terms across X, S, A , and with a component that does depend on S versus one that does not. The goal was to upweight the original Y such that the magnitude of the contribution of the different parts are on the same scale, and the noise was chosen to hold the Gaussian assumption we introduce, as well as retaining a meaningful signal to noise

ratio when trying to fit a model for the conditional $\mathbb{E}[Y|S, X, A]$. We plotted the data to ensure that the Gaussianity assumption holds and the observed residuals to ensure that the model is reasonably fitted.

The original dataset includes 345 data points. For the `ModBrk` constraint, we use a bootstrapped version (by re-sampling S, X, E with replacement), increasing the number of data points to 20,000.

IHDP dataset. The IHDP dataset consists of data from a program that targeted low-birth-weight, premature infants, providing them with intensive high-quality child care and home visits from a trained provider. The continuous action A is the self-selected number of participation days in the program. The outcome Y is the child score attainment in cognitive tests at age three. As sensitive attribute S , we consider mother’s race (white vs. non-white). After removing duplicate encoding and NaN values, we are left with 31 covariates, 22 categorical and 9 continuous. We use a bootstrapped version (by re-sampling S, X with replacement, and obtaining corresponding A, Y predictions from two black-box models trained on the original S, Y in the dataset), increasing the number of samples to 20,000.

We re-interpret the original setting into a resource allocation problems as follows: rather than as a self-decision, we suggest to think of the number of days in treatment as external, e.g., voucher allocation or by assigning different individuals to different lengths of participation. In this case, the `ModBrk` constraint’s goal is to break the moderation of the allocation of days in program by the group membership (white vs. non-white mothers), such that the resource allocation policy is not responsible for the difference in test scores attained by different groups.

Appendix F. Additional experimental details

All experiments were implemented using the Pytorch library (Paszke et al., 2019). The `EqB` experiments were run on a standard personal Mac, using CPUs. The `ModBrk` experiments were run on a single Microsoft Azure NVIDIA Tesla P100 GPU server (NC6s v2 instance). We provide the corresponding environment setups in the accompanying code. After performing hyper-parameter tuning, we ran the experiments appearing in the manuscripts with the following configurations:

F.1. `EqB` NYCSchools

Our stage I model consists of two MLPs, where MLP_{θ}^A learns μ_{θ}^A and MLP^Y learns μ^Y as described in the main text. Both MLPs share the same structure, one hidden linear layer and one ReLU layer (Agarap, 2018), with the hidden dimension set to 128 (for each component). We train the stage I model for 500 epochs for MLP^Y and 1,000 epochs for MLP_{θ}^A , with a learning rate set at 0.01 and 0.0001 respectively, and the optimizer is Adam (Kingma and Ba, 2017). We verify the training of our stage I models by checking R^2 score involving predictions and the target. We also check the predictions of the trained model MLP^Y for constant actions against the original model’s.

The stage II model consists of one linear layer and one ReLU layer, with the hidden dimension set to 10 (for each component). In order to speed up its convergence, we standardize the output of the stage II model with constant values. We train the stage II model for 1,000 epochs, with learning rate 0.01.

F.2. ModBrk NYCSchools

Our stage I model is an MLP composed of additive f, g, h components, as described in the main text. For this setting, we use three linear layers, and two ReLU non-linearities (Agarap, 2018), with the hidden dimension set to 256 (for each component). We train the stage I model for 3,000 epochs, with a learning rate set at 0.005, and the optimizer is Adam (Kingma and Ba, 2017). We verify the training of our stage I model by checking explained variance and R^2 scores on a held-out set. We also check the predictions of the trained model for constant actions against the original model's.

The stage II model consists of 3 linear layers and 2 ReLU nonlinearities, with a single batchnorm layer following the first input layer. We also include a shifted Sigmoid at the end of the network to enforce our action clipping, as described in the main text. For this setting, the hidden dimension is set to 64, the learning rate is 0.001, and we use the Adam optimizer (Kingma and Ba, 2017). We use scheduling for smoother training, and for both, we use a single batch of size 20,000.

F.3. ModBrk IHDP

We use similar model architectures to the ones above for this dataset, with the following modifications. For the stage I model, we use only 2 linear layers and 1 ReLU non-linearity (Agarap, 2018), with the hidden dimensions set at 512. This model is trained for 1,000 epochs. We compare constant action predictions compared to those of a standard black box model trained for $\mathbb{E}[Y|S, X, A]$. The learning rate used for stage I is 0.001, with a small weight decay of 0.1 for regularization.

For stage II, we use 50 hidden dimensions and a learning rate of 0.0005. We train for 3,000 epochs. See complete configuration specifications in the files `configs/NYCSchools_EqB_configs.yml`, `configs/NYCSchools_ModBrk_configs.yml`, `configs/IHDP_ModBrk_configs.yml` in the accompanying code.