

A Statistical Test for Probabilistic Fairness

Bahar Taskesen

Ecole Polytechnique Fédérale de Lausanne
Switzerland
bahar.taskesen@epfl.ch

Daniel Kuhn

Ecole Polytechnique Fédérale de Lausanne
Switzerland
daniel.kuhn@epfl.ch

Jose Blanchet

jose.blanchet@stanford.edu
Stanford University
USA

Viet Anh Nguyen

Stanford University
USA
viet-anh.nguyen@stanford.edu

ABSTRACT

Algorithms are now routinely used to make consequential decisions that affect human lives. Examples include college admissions, medical interventions or law enforcement. While algorithms empower us to harness all information hidden in vast amounts of data, they may inadvertently amplify existing biases in the available datasets. This concern has sparked increasing interest in fair machine learning, which aims to quantify and mitigate algorithmic discrimination. Indeed, machine learning models should undergo intensive tests to detect algorithmic biases before being deployed at scale. In this paper, we use ideas from the theory of optimal transport to propose a statistical hypothesis test for detecting unfair classifiers. Leveraging the geometry of the feature space, the test statistic quantifies the distance of the empirical distribution supported on the test samples to the manifold of distributions that render a pre-trained classifier fair. We develop a rigorous hypothesis testing mechanism for assessing the probabilistic fairness of any pre-trained logistic classifier, and we show both theoretically as well as empirically that the proposed test is asymptotically correct. In addition, the proposed framework offers interpretability by identifying the most favorable perturbation of the data so that the given classifier becomes fair.

CCS CONCEPTS

• **Applied computing** → *IT governance; Law*; • **Social and professional topics** → *Race and ethnicity; Geographic characteristics; Sexual orientation; Gender; Age*; • **Theory of computation** → *Mathematical optimization*.

KEYWORDS

fairness, algorithmic bias, equal opportunity, equalized odds, Wasserstein distance

ACM Reference Format:

Bahar Taskesen, Jose Blanchet, Daniel Kuhn, and Viet Anh Nguyen. 2021. A Statistical Test for Probabilistic Fairness. In *ACM Conference on Fairness*,

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

FACCT '21, March 1–10, 2021, Virtual Event, Canada

© 2021 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-8309-7/21/03...\$15.00

<https://doi.org/10.1145/3442188.3445927>

Accountability, and Transparency (FACCT '21), March 1–10, 2021, Virtual Event, Canada. ACM, New York, NY, USA, 18 pages. <https://doi.org/10.1145/3442188.3445927>

1 INTRODUCTION

The past decade witnessed data and algorithms becoming an integrative part of the human society. Recent technological advances are now allowing us to collect and store an astronomical amount of unstructured data, and the unprecedented computing power is enabling us to convert these data into decisional insights. Nowadays, machine learning algorithms can uncover complex patterns in the data to produce an exceptional performance that can match, or even surpass, that of humans. These algorithms, as a consequence, are proliferating in every corner of our lives, from suggesting us the next vacation destination to helping us create digital paintings and melodies. Machine learning algorithms are also gradually assisting humans in consequential decisions such as deciding whether a student is admitted to college, picking which medical treatment to be prescribed to a patient, and determining whether a person is convicted. Arguably, these decisions impact radically many people's lives, together with the future of their loved ones.

Algorithms are conceived and function following strict rules of logic and algebra; it is hence natural to expect that machine learning algorithms deliver objective predictions and recommendations. Unfortunately, in-depth investigations reveal the excruciating reality that state-of-the-art algorithmic assistance is far from being free of biases. For example, a predictive algorithm widely used in the United States criminal justice system is more likely to *misclassify* African-American offenders into the group of high recidivism risk compared to white-Americans [12, 46]. The artificial intelligence tool developed by Amazon also learned to penalize gender-related keywords such as “women’s” in the profile screening process, and thus may prefer to recommend hiring male candidates for software development and technical positions [17]. Further, Google’s ad-targeting algorithm displayed advertisements for higher-paying executive jobs more often to men than to women [18].

There are several possible explanations for why cold, soulless algorithms may trigger biased recommendations. First, the data used to train machine learning algorithms may already encrypt human biases manifested in the data collection process. These biases arise as the result of a suboptimal design of experiments, or from historically biased human decisions that accumulate over centuries. Machine-learned algorithms, which are apt to detect underlying patterns from data, will unintentionally learn and maintain these

existing biases [9, 43]. For example, secretary or primary school teacher are professions which are predominantly taken by women, thus, natural language processing systems are inclined to associate female attributes to these jobs. Second, training a machine learning algorithm typically involves minimizing the prediction error which privileges the majority populations over the minority groups. Clinical trials, for instance, typically involve very few participants from the minority groups such as indigenous people, and thus medical interventions recommended by the algorithms may not align perfectly to the characteristics and interests of patients from the minority groups. Finally, even when the sensitive attributes are not used in the training phase, strong correlations between the sensitive attributes and the remaining variables in the dataset may be exploited to generate unjust actions. For example, the sensitive attribute of race can be easily inferred with high accuracy based on common non-sensitive attributes such as the travel history of passengers or the grocery shopping records of customers.

The pressing needs to redress undesirable algorithmic biases have propelled the rising field of fair machine learning¹. A building pillar of this field involves the verification task: given a machine learning algorithm, we are interested in verifying if this algorithm satisfies a chosen criterion of fairness. This task is performed in two steps: first, we choose an appropriate notion of fairness, then the second step invokes a computational procedure, which may or may not involve data, to decide if the chosen fairness criterion is fulfilled. A plethora of criteria for fair machine learning were proposed in the literature, many of them are motivated by philosophical or sociological ideologies or legal constraints. For example, anti-discrimination laws may prohibit making decisions based on sensitive attributes such as age, gender, race or sexual orientation. Thus, a naïve strategy, called fairness through unawareness, involves removing all sensitive attributes from the training data. However, this strategy seldom guarantees any fairness due to the inter-correlation issues [27, 30], and thus potentially fails to generate inclusive outcomes [2, 6, 36, 41]. Other notions of fairness aim to either promote individual fairness [21], prevent disparate treatment [70] or avoid disparate mistreatment [23, 71] of the algorithms. Towards similar goals, notions of *group* fairness focus on reducing the difference of favorable outcomes proportions among different sensitive groups. Examples of group fairness notions include disparate impact [70], demographic parity (statistical parity) [10, 21], equality of opportunity [31] and equalized odds [31]. The notion of counterfactual fairness [27] was also suggested as a measure of causal fairness. Despite the abundance of available notions, there is unfortunately no general consensus on the most suitable measure to serve as the industry standard. Moreover, except in trivial cases, it is not possible for a machine learning algorithm to simultaneously satisfy multiple notions of fairness [5, 37]. Therefore, the choice of the fairness notion is likely to remain more an art than a science.

This paper focuses not on the normative approach to choosing an ideal notion of machine learning fairness. We endeavor in this paper to shed more light on the computational procedure to complement the verification task. Concretely, we position ourselves in the classification setting, which is arguably the most popular task in machine learning. Moreover, we will focus on notions of group

fairness, and we employ the framework of statistical hypothesis test instead of algorithmic test.

Contributions. Our paper makes two concrete contributions to the problem of fairness testing of machine learning’s classifiers.

- (1) We propose the Wasserstein projection framework to perform statistical hypothesis test of group fairness for classification algorithms. We derive in details the computation of the test statistic and the limiting distribution when fairness is measured using the probabilistic equality of opportunity and probabilistic equalized odds criteria.
- (2) We demonstrate that the Wasserstein projection hypothesis testing paradigm is asymptotically correct and can exploit additional information on the geometry of the feature space. Moreover, we also show that this paradigm promotes transparency and interpretability through the analysis of the most favorable distributions.

The remaining of the paper is structured as follows. In Section 2, we introduce the general problem of statistical hypothesis test of classification fairness, and depict the current landscape of fairness testing in the literature. Section 3 details our Wasserstein projection approach to this problem. Sections 4 and 5 apply the proposed framework to test if a pre-trained logistic classifier satisfies the fairness notion of probabilistic equal opportunity and probabilistic equalized odds, respectively. Numerical experiments are presented in Section 6 to empirically validate the correctness and demonstrate the power of our proposed paradigm. Section 7 concludes the paper with outlooks on the broader impact of our Wasserstein projection hypothesis testing approach.

All technical proofs are relegated to the Appendix.

2 STATISTICAL TESTING FRAMEWORK FOR FAIRNESS AND LITERATURE REVIEW

We consider throughout this paper a generic binary classification setting. Let $\mathcal{X} = \mathbb{R}^d$ and $\mathcal{Y} = \{0, 1\}$ be the space of feature inputs and label outputs of interest. We assume that there is a single sensitive attribute corresponding to each data point and its space is denoted by $\mathcal{A} = \{0, 1\}$. A probabilistic classifier is represented by a function $h(\cdot) : \mathcal{X} \rightarrow [0, 1]$ that outputs for each given sample $x \in \mathcal{X}$ the probability that x belongs to the positive class. The deterministic classifier predicts class 1 if $h(x) \geq \tau$ and class 0 otherwise, where $\tau \in [0, 1]$ is a classification threshold. Note that the function h depends only on the feature X , but not on the sensitive attribute A , thus predicting Y using h satisfies fairness through unawareness.

The central goal of this paper is to provide a statistical test to detect if a classifier h fails to satisfy a prescribed notion of machine learning fairness. A statistical hypothesis test can be cast with the null hypothesis being

$$\mathcal{H}_0: \text{the classifier } h \text{ is fair,}$$

against the alternative hypothesis being

$$\mathcal{H}_1: \text{the classifier } h \text{ is not fair.}$$

In this paper, we focus on statistical notions of *group* fairness, which are usually defined using conditional probabilities. A prevalent notion of fairness in machine learning is the criterion of equality of

¹Comprehensive surveys on fair machine learning can be found in [5, 13, 14, 44].

opportunity², which requires that the true positive rate are equal between subgroups.

Definition 2.1 (Equal opportunity [31]). A classifier $h(\cdot) : \mathcal{X} \rightarrow [0, 1]$ satisfies the equal opportunity criterion relative to \mathbb{Q} if

$$\mathbb{Q}(h(X) \geq \tau | A = 1, Y = 1) = \mathbb{Q}(h(X) \geq \tau | A = 0, Y = 1),$$

where τ is the classification threshold.

Another popular criterion of machine learning fairness is the equalized odds, which is more stringent than the equality of opportunity: it requires that the positive outcome is conditionally independent of the sensitive attributes given the true label.

Definition 2.2 (Equalized odds [31]). A classifier $h(\cdot) : \mathcal{X} \rightarrow [0, 1]$ satisfies the equalized odds criterion relative to \mathbb{Q} if

$$\mathbb{Q}(h(X) \geq \tau | A = 1, Y = y) = \mathbb{Q}(h(X) \geq \tau | A = 0, Y = y) \quad \forall y \in \mathcal{Y},$$

where τ is the classification threshold.

Notice that the criteria of fairness presented in Definitions 2.1 and 2.2 are dependent on the distribution \mathbb{Q} : a classifier h can be fair relative to a distribution \mathbb{Q}_1 , but it may become unfair with respect to another distribution $\mathbb{Q}_2 \neq \mathbb{Q}_1$. If we denote by \mathbb{P} the true population distribution that governs the random vector (X, A, Y) , then it is imperative and reasonable to test for group fairness with respect to \mathbb{P} . For example, to test for the equality of opportunity, we can reformulate a two-sample equal conditional mean test of the null hypothesis

$$\mathcal{H}_0 : \mathbb{E}_{\mathbb{P}}[\mathbb{1}_{h(X) \geq \tau} | A = 1, Y = 1] = \mathbb{E}_{\mathbb{P}}[\mathbb{1}_{h(X) \geq \tau} | A = 0, Y = 1],$$

and one can potentially employ a Welch's t -test with proper adjustment for the randomness of the sample size. Unfortunately, deriving the test becomes complicated when the null hypothesis involves an equality of multi-dimensional quantities, which arises in the case of equalized odds, due to the complication of the covariance terms. Variations of the permutation tests were also proposed to detect discriminatory behaviour of machine learning algorithms following the same formulation of the one-dimensional two-sample equality of conditional mean test [19, 66]. However, these permutation tests follow a black-box mechanism and are unable to be generalized to multi-dimensional tests. Tests based on group fairness notions can also be accomplished using an algorithmic approach as in [19, 29, 35, 57].

From a broader perspective, deriving tests for fairness is an active area of research, and many testing procedures have been recently proposed to test for individual fairness [34, 68], for counterfactual fairness [6, 27] and diverse other criteria [3, 66, 67].

Literature related to optimal transport. Optimal transport is a long-standing field that dates back to the seminal work of Gaspard Monge [45]. In the past few years, it has attracted significant attention in the machine learning and computer science communities thanks to the availability of fast approximation algorithms [4, 8, 16, 20, 28]. Optimal transport is particularly successful in various learning tasks, notably generative mixture models [38, 49], image processing [1, 24, 39, 50, 63], computer vision and graphics [51, 52, 56, 61, 62], clustering [32], dimensionality reduction [11, 25, 55, 58, 59], domain adaptation [15, 47], signal

²We use two terms "equality of opportunity" and "equal opportunity" interchangeably.

processing [65] and data-driven distributionally robust optimization [7, 26, 40, 72]. Recent comprehensive survey on optimal transport and its applications can be found in [38, 53].

In the context of fair classification, ideas from optimal transport have been used to construct fair logistic classifier [64], to detect classifiers that does not obey group fairness notions, or to ensure fairness by pre-processing [29], to learn a fair subspace embedding that promotes fair classification [69], to test individual fairness [68], or to construct a counterfactual test [6].

3 WASSERSTEIN PROJECTION FRAMEWORK FOR STATISTICAL TEST OF FAIRNESS

We hereby provide a fresh alternative to the testing problem of machine learning fairness. On that purpose, for a given classifier h , we define abstractly the following set of distributions

$$\mathcal{F}_h = \{\mathbb{Q} \in \mathcal{P} : \text{the classifier } h \text{ is fair relative to } \mathbb{Q}\}, \quad (1)$$

where \mathcal{P} denotes the space of all distributions on $\mathcal{X} \times \mathcal{A} \times \mathcal{Y}$. Intuitively, the set \mathcal{F}_h contains all probability distributions under which the classifier h satisfies the prescribed notion of fairness. It is trivial to see that if \mathcal{F}_h contains the true data-generating distribution \mathbb{P} , then the classifier h is fair relative to \mathbb{P} . Thus, we can reinterpret the hypothesis test of fairness using the hypotheses

$$\mathcal{H}_0 : \mathbb{P} \in \mathcal{F}_h, \quad \mathcal{H}_1 : \mathbb{P} \notin \mathcal{F}_h.$$

Testing the inclusion of \mathbb{P} in \mathcal{F}_h is convenient if \mathcal{P} is endowed with a distance. In this paper, we equip \mathcal{P} with the Wasserstein distance.

Definition 3.1 (Wasserstein distance). The type-2 Wasserstein distance between two probability distributions \mathbb{Q} and \mathbb{Q}' supported on Ξ is defined as

$$\mathbb{W}(\mathbb{Q}', \mathbb{Q}) = \min_{\pi \in \Pi(\mathbb{Q}', \mathbb{Q})} \sqrt{\mathbb{E}_{\pi}[c(\xi', \xi)^2]},$$

where the set $\Pi(\mathbb{Q}', \mathbb{Q})$ contains all joint distributions of the random vectors $\xi' \in \Xi$ and $\xi \in \Xi$ under which ξ' and ξ have marginal distributions \mathbb{Q}' and \mathbb{Q} , respectively, and $c : \Xi \times \Xi \rightarrow [0, \infty]$ constitutes a lower semi-continuous ground metric.

The type-2 Wasserstein distance³ is a special instance of the optimal transport. The squared Wasserstein distance between \mathbb{Q}' and \mathbb{Q} can be interpreted as the cost of moving the distribution \mathbb{Q}' to \mathbb{Q} , where $c(\xi', \xi)$ is the cost of moving a unit mass from ξ' to ξ . Being a distance on \mathcal{P} , \mathbb{W} is symmetric, non-negative and vanishes to zero if $\mathbb{Q}' = \mathbb{Q}$. The Wasserstein distance is hence an attractive measure to identify if \mathbb{P} belongs to \mathcal{F}_h . Using this insight, the hypothesis test for fairness has the equivalent representation

$$\mathcal{H}_0 : \inf_{\mathbb{Q} \in \mathcal{F}_h} \mathbb{W}(\mathbb{P}, \mathbb{Q}) = 0, \quad \mathcal{H}_1 : \inf_{\mathbb{Q} \in \mathcal{F}_h} \mathbb{W}(\mathbb{P}, \mathbb{Q}) > 0.$$

Even though \mathbb{P} remains elusive to our knowledge, we are given access to a set of i.i.d test samples $\{(\hat{x}_i, \hat{a}_i, \hat{y}_i)\}_{i=1}^N$ generated from the true distribution \mathbb{P} . Thus we can rely on the empirical value

$$\inf_{\mathbb{Q} \in \mathcal{F}_h} \mathbb{W}(\hat{\mathbb{P}}^N, \mathbb{Q}),$$

which is the distance from the empirical distribution supported on the samples $\hat{\mathbb{P}}^N = \sum_{i=1}^N \delta_{(\hat{x}_i, \hat{a}_i, \hat{y}_i)}$ to the set \mathcal{F}_h . To perform the test, it is sufficient to study the limiting distribution of the test statistic using proper scaling under the null hypothesis \mathcal{H}_0 . The outcome of

³From this point, we omit the term "type-2" for brevity.

the test is determined by comparing the test statistic to the quantile value of the limiting distribution at a chosen level of significant $\alpha \in (0, 1)$.

Advantages. The Wasserstein projection framework to hypothesis testing that we described above offers several advantages over the existing methods.

- (1) Geometric flexibility: The definition of the Wasserstein distance implies that there exists a joint ground metric c on the space of the features, the sensitive attribute and the label. If the modelers or the regulators possess any structural information on an appropriate metric on $\Xi = \mathcal{X} \times \mathcal{A} \times \mathcal{Y}$, then this information can be exploited in the testing procedure. Thus, the Wasserstein projection framework equips the users with an additional freedom to inject prior geometric information into the statistical test.
- (2) Multivariate generalizability: Certain notions of fairness, such as equalized odds, are prescribed using multiple equalities of conditional expectations. The Wasserstein projection framework encapsulates these equalities simultaneously in the definition of the set \mathcal{F}_h , and provides a *joint* test of these equalities without the hassle of decoupling and testing individual equalities as being done in the currently literature.
- (3) Interpretability: If we denote by \mathbb{Q}^* the projection of the empirical distribution $\hat{\mathbb{P}}^N$ onto the set of distributions \mathcal{F}_h , i.e.,

$$\mathbb{Q}^* = \arg \min_{\mathbb{Q} \in \mathcal{F}_h} \mathbb{W}(\hat{\mathbb{P}}^N, \mathbb{Q}),$$

then \mathbb{Q}^* encodes the minimal perturbation to the empirical samples so that the classifier h becomes fair. The distribution \mathbb{Q}^* is thus termed the most favorable distribution, and examining \mathbb{Q}^* can reveal the underlying mechanism and explain the outcome of the hypothesis test. The accessibility to \mathbb{Q}^* showcases the expressiveness of the Wasserstein projection framework.

Whilst theoretically sound and attractive, there are three potential difficulties with the Wasserstein projection approach to statistical test of fairness. First, to project $\hat{\mathbb{P}}^N$ onto the set \mathcal{F}_h , we need to solve an infinite-dimensional optimization problem, which is inherently difficult. Second, for many notions of machine learning fairness such as the equality of opportunity and the equalized odds, the corresponding set \mathcal{F}_h in (1) is usually prescribed using *nonlinear* constraints. For example, if we consider the equal opportunity criterion in Definition 2.1, then the set \mathcal{F}_h can be re-expressed using a fractional function of the probability measure as

$$\mathcal{F}_h = \left\{ \mathbb{Q} \in \mathcal{P} \text{ such that } \frac{\mathbb{Q}(h(X) \geq \tau, A = 1, Y = 1)}{\mathbb{Q}(A = 1, Y = 1)} = \frac{\mathbb{Q}(h(X) \geq \tau, A = 0, Y = 1)}{\mathbb{Q}(A = 0, Y = 1)} \right\}.$$

Apart from involving nonlinear constraints, it is easy to verify that the set \mathcal{F}_h is also non-convex, which amplifies the difficulty of computing the projection onto \mathcal{F}_h . Finally, the limiting distribution of the test statistic is difficult to analyze due to the discontinuity of the probability function at the set $\{x \in \mathcal{X} : h(x) = \tau\}$. The asymptotic analysis with this discontinuity is of a combinatorial nature, and is significantly more problematic than the asymptotic analysis of smooth quantities.

While these difficulties may be overcome via various ways, in this paper we choose the following combination of remedies. First,

we will use a relaxed notion of fairness termed *probabilistic fairness*, which was originally introduced in [54]. Second, when computing the Wasserstein distances between distributions on $\mathcal{X} \times \mathcal{A} \times \mathcal{Y}$, we use

$$c((x', a', y'), (x, a, y)) = \|x - x'\| + \infty|a - a'| + \infty|y - y'| \quad (2)$$

as the ground metric, where $\|\cdot\|$ is a norm on \mathbb{R}^d . This case corresponds to having an absolute trust in the label and in the sensitive attribute of the training samples. This absolute trust restriction is common in the literature of fair machine learning [64, 68].

We now briefly discuss the advantage of using the ground metric of the form (2). Denote by $p \in \mathbb{R}_{++}^{|\mathcal{A}| \times |\mathcal{Y}|}$ the array of the true marginals of (A, Y) , in particular, $p_{ay} = \mathbb{P}(A = a, Y = y)$ for all $a \in \mathcal{A}$ and $y \in \mathcal{Y}$. Further, let $\hat{p}^N \in \mathbb{R}_{++}^{|\mathcal{A}| \times |\mathcal{Y}|}$ be the array of the empirical marginals of (A, Y) under the empirical measure $\hat{\mathbb{P}}^N$, that is, $\hat{p}_{ay}^N = \hat{\mathbb{P}}^N(A = a, Y = y)$ for all $a \in \mathcal{A}$ and $y \in \mathcal{Y}$. Throughout this paper, we assume that the empirical marginals are proper, that is, $\hat{p}_{ay}^N \in (0, 1)$ for any $(a, y) \in \mathcal{A} \times \mathcal{Y}$. We define temporarily the simplex set $\Delta := \{\bar{p} \in \mathbb{R}_{++}^{|\mathcal{A}| \times |\mathcal{Y}|} : \sum_{a \in \mathcal{A}, y \in \mathcal{Y}} \bar{p}_{ay} = 1\}$. Subsequently, for any marginals $\bar{p} \in \Delta$, we define the marginally-constrained set of distributions

$$\mathcal{F}_h(\bar{p}) \triangleq \left\{ \mathbb{Q} \in \mathcal{P} : \begin{array}{l} h \text{ is fair relative to } \mathbb{Q} \\ \mathbb{Q}(A = a, Y = y) = \bar{p}_{ay} \quad \forall (a, y) \in \mathcal{A} \times \mathcal{Y} \end{array} \right\}.$$

Using these notations, one can readily verify that $\mathcal{F}_h = \cup_{\bar{p} \in \Delta} \mathcal{F}_h(\bar{p})$. Moreover, the next result asserts that in order to compute the projection of $\hat{\mathbb{P}}^N$ onto \mathcal{F}_h , it suffices to project onto the marginally-constrained set $\mathcal{F}_h(\hat{p}^N)$.

Lemma 3.2 (Projection with marginal restrictions). Suppose that the ground metric is chosen as in (2). If a measure $\mathbb{Q} \in \mathcal{F}_h$ satisfies $\mathbb{W}(\hat{\mathbb{P}}^N, \mathbb{Q}) < \infty$, then $\mathbb{Q} \in \mathcal{F}_h(\hat{p}^N)$.

A useful consequence of Lemma 3.2 is that

$$\inf_{\mathbb{Q} \in \mathcal{F}_h} \mathbb{W}(\hat{\mathbb{P}}^N, \mathbb{Q}) = \inf_{\mathbb{Q} \in \mathcal{F}_h(\hat{p}^N)} \mathbb{W}(\hat{\mathbb{P}}^N, \mathbb{Q}), \quad (3)$$

where the feasible set of the problem on the right-hand side is the marginally-constrained set $\mathcal{F}_h(\hat{p}^N)$ using the empirical marginals \hat{p}^N . For two notions of probabilistic fairness that we will explore in this paper, projecting $\hat{\mathbb{P}}^N$ onto $\mathcal{F}_h(\hat{p}^N)$ is arguably easier than onto \mathcal{F}_h . Thus, this choice of ground metric improves the tractability when computing the test statistic.

Third, and finally, we will focus on the logistic regression setting, which is one of the most popular classification methods [33]. In this setting, the conditional probability $\mathbb{P}[Y = 1 | X = x]$ is modelled by the sigmoid function $h_\beta(x) = (1 + \exp(-\beta^\top x))^{-1}$, where $\beta \in \mathbb{R}^d$ is the regression parameter. Moreover, a classifier with $\beta = 0$, is trivially fair. Thus, it suffices to consider $\beta \neq 0$.

Notations. We use $\|\cdot\|_*$ to denote the dual norm of $\|\cdot\|$. For any integer N , we define $[N] := \{1, 2, \dots, N\}$. Given N test samples $(\hat{x}_i, \hat{a}_i, \hat{y}_i)_{i=1}^N$, we use $\mathcal{I}_y \triangleq \{i \in [N] : \hat{y}_i = y\}$ to denote the index set of observations with label y . The parameters λ_i are defined as

$$\forall i \in [N] : \lambda_i = \begin{cases} (\hat{p}_{11}^N)^{-1} & \text{if } (\hat{a}_i, \hat{y}_i) = (1, 1), \\ -(\hat{p}_{01}^N)^{-1} & \text{if } (\hat{a}_i, \hat{y}_i) = (0, 1), \\ (\hat{p}_{10}^N)^{-1} & \text{if } (\hat{a}_i, \hat{y}_i) = (1, 0), \\ -(\hat{p}_{00}^N)^{-1} & \text{if } (\hat{a}_i, \hat{y}_i) = (0, 0). \end{cases} \quad (4)$$

4 TESTING FAIRNESS FOR PROBABILISTIC EQUAL OPPORTUNITY CRITERION

In this section, we use the ingredients introduced in the previous section to concretely construct a statistical test for the fairness of a logistic classifier h_β . Specifically, we will employ the probabilistic equal opportunity criterion which was originally proposed in [54].

Definition 4.1 (Probabilistic equal opportunity criterion [54]). A logistic classifier $h_\beta : \mathcal{X} \rightarrow [0, 1]$ satisfies the probabilistic equalized opportunity criteria relative to a distribution \mathbb{Q} if

$$\mathbb{E}_{\mathbb{Q}}[h_\beta(X)|A = 1, Y = 1] = \mathbb{E}_{\mathbb{Q}}[h_\beta(X)|A = 0, Y = 1].$$

The probabilistic equal opportunity criterion, which serves as a surrogate for the equal opportunity criterion in Definition 2.1, depends on the smooth and bounded sigmoid function h_β but is independent of the classification threshold τ . Motivated by [42], we empirically illustrate in Figure 1 that the probabilistic surrogate provides a good approximation of the equal opportunity criterion. Figure 1a plots the absolute difference of the classification probabilities $|\mathbb{P}(h(X) \geq \frac{1}{2}|A = 1, Y = 1) - \mathbb{P}(h(X) \geq \frac{1}{2}|A = 0, Y = 1)|$, while Figure 1b plots the absolute difference of the sigmoid expectations $|\mathbb{E}_{\mathbb{P}}[h(X)|A = 1, Y = 1] - \mathbb{E}_{\mathbb{P}}[h(X)|A = 0, Y = 1]|$. One may observe that the regions of β so that the absolute differences fall close to zero are similar in both plots. This implies that a logistic classifier h_β which is equal opportunity fair is also likely to be *probabilistic* equal opportunity fair, and vice versa.




Figure 1: Comparison of fairness notions for $d = 2$ and $h_\beta(x) = 1/(1 + \exp(\frac{1}{3} - \beta_1 x_1 - \beta_2 x_2))$.

We use the superscript “opp” to emphasize that fairness is measured using the probabilistic equal opportunity criterion. Consequently, the set of distributions $\mathcal{F}_{h_\beta}^{\text{opp}}$ that makes the logistic classifier h_β fair is

$$\mathcal{F}_{h_\beta}^{\text{opp}} = \left\{ \mathbb{Q} \in \mathcal{P} \text{ such that : } \mathbb{E}_{\mathbb{Q}}[h_\beta(X)|A = 1, Y = 1] = \mathbb{E}_{\mathbb{Q}}[h_\beta(X)|A = 0, Y = 1] \right\}.$$

The statistical hypothesis test to verify whether the classifier h_β is fair is formulated with the null and alternative hypotheses

$$\mathcal{H}_0^{\text{opp}} : \mathbb{P} \in \mathcal{F}_{h_\beta}^{\text{opp}}, \quad \mathcal{H}_1^{\text{opp}} : \mathbb{P} \notin \mathcal{F}_{h_\beta}^{\text{opp}}.$$

The remainder of this section unfolds as follows. In Section 4.1, we delineate the computation of the projection of $\hat{\mathbb{P}}^N$ onto $\mathcal{F}_{h_\beta}^{\text{opp}}$. Section 4.2 studies the limiting distribution of the test statistic, while Section 4.3 examines the most favorable distribution.

4.1 Wasserstein Projection

Lemma 3.2 suggests that it is sufficient to consider the projection onto the marginally-constrained set $\mathcal{F}_{h_\beta}^{\text{opp}}(\hat{\mathbb{P}}^N)$, where $\hat{\mathbb{P}}^N$ is the empirical marginals of the empirical distribution $\hat{\mathbb{P}}^N$. In particular, $\mathcal{F}_{h_\beta}^{\text{opp}}(\hat{\mathbb{P}}^N)$ is

$$\mathcal{F}_{h_\beta}^{\text{opp}}(\hat{\mathbb{P}}^N) = \left\{ \mathbb{Q} \in \mathcal{P} \text{ such that : } \begin{aligned} & (\hat{\mathbb{P}}_{11}^N)^{-1} \mathbb{E}_{\mathbb{Q}}[h_\beta(X) \mathbb{1}_{(1,1)}(A, Y)] = (\hat{\mathbb{P}}_{01}^N)^{-1} \mathbb{E}_{\mathbb{Q}}[h_\beta(X) \mathbb{1}_{(0,1)}(A, Y)] \\ & \mathbb{Q}(A = a, Y = y) = \hat{\mathbb{P}}_{ay}^N \quad \forall (a, y) \in \mathcal{A} \times \mathcal{Y} \end{aligned} \right\},$$

where the equality follows from the law of conditional expectation. Notice that the set $\mathcal{F}_{h_\beta}^{\text{opp}}(\hat{\mathbb{P}}^N)$ is prescribed using *linear* constraints of \mathbb{Q} , and thus it is more amenable to optimization than the set $\mathcal{F}_{h_\beta}^{\text{opp}}$. It is also more convenient to work with the *squared* distance function \mathcal{R} whose input is the empirical distribution $\hat{\mathbb{P}}^N$ and its corresponding vector of empirical marginals $\hat{\mathbb{P}}^N$ by

$$\mathcal{R}^{\text{opp}}(\hat{\mathbb{P}}^N, \hat{\mathbb{P}}^N) := \begin{cases} \inf & \mathbb{W}(\mathbb{Q}, \hat{\mathbb{P}}^N)^2 \\ \text{s.t.} & \mathbb{E}_{\mathbb{Q}}[h_\beta(X) ((\hat{\mathbb{P}}_{11}^N)^{-1} \mathbb{1}_{(1,1)}(A, Y) - (\hat{\mathbb{P}}_{01}^N)^{-1} \mathbb{1}_{(0,1)}(A, Y))] = 0 \\ & \mathbb{E}_{\mathbb{Q}}[\mathbb{1}_{(a,y)}(A, Y)] = \hat{\mathbb{P}}_{ay}^N \quad \forall (a, y) \in \mathcal{A} \times \mathcal{Y}. \end{cases}$$

Notice that the constraints of the above infimum problem are linear in the measure \mathbb{Q} , but the functions inside the expectation operators are possibly *nonlinear* functions of $\hat{\mathbb{P}}^N$. Using the equivalent characterization (3), the following relation holds

$$\inf_{\mathbb{Q} \in \mathcal{F}_{h_\beta}^{\text{opp}}} \mathbb{W}(\hat{\mathbb{P}}^N, \mathbb{Q}) = \inf_{\mathbb{Q} \in \mathcal{F}_{h_\beta}^{\text{opp}}(\hat{\mathbb{P}}^N)} \mathbb{W}(\hat{\mathbb{P}}^N, \mathbb{Q}) = \sqrt{\mathcal{R}^{\text{opp}}(\hat{\mathbb{P}}^N, \hat{\mathbb{P}}^N)}.$$

We now proceed to show how computing the projection can be reduced to solving a finite-dimensional optimization problem.

Proposition 4.2 (Dual reformulation). The squared projection distance $\mathcal{R}^{\text{opp}}(\hat{\mathbb{P}}^N, \hat{\mathbb{P}}^N)$ equals to the optimal value of the following finite-dimensional optimization problem

$$\sup_{\gamma \in \mathbb{R}} \frac{1}{N} \sum_{i \in \mathcal{I}_1} \inf_{x_i \in \mathcal{X}} \{ \|x_i - \hat{x}_i\|^2 + \gamma \lambda_i h_\beta(x_i) \}. \quad (5)$$

While Proposition 4.2 asserts that computing the *squared* projection distance $\mathcal{R}^{\text{opp}}(\hat{\mathbb{P}}^N, \hat{\mathbb{P}}^N)$ is equivalent to solving a finite-dimensional problem, unfortunately, this saddle point problem is in general difficult. Indeed, because h_β is non-convex, even finding the optimal inner solution x_i^* for a fixed value of the outer variable $\gamma \in \mathbb{R}$ is generally NP-hard [48]. The situation can be partially alleviated if $\|\cdot\|$ is an Euclidean norm on \mathbb{R}^d .

Lemma 4.3 (Univariate reduction). Suppose that $\|\cdot\|$ is the Euclidean norm on \mathbb{R}^d , we have

$$\mathcal{R}^{\text{opp}}(\hat{\mathbb{P}}^N, \hat{\mathbb{P}}^N) = \sup_{\gamma \in \mathbb{R}} \frac{1}{N} \sum_{i \in \mathcal{I}_1} \min_{k_i \in [0, \frac{1}{8}]} \gamma^2 \lambda_i^2 \|\beta\|_2^2 k_i^2 + \frac{\gamma \lambda_i}{1 + \exp(\gamma \lambda_i \|\beta\|_2^2 k_i - \beta^\top \hat{x}_i)}. \quad (6)$$

The proof of Lemma 4.3 follows trivially from application of Lemma B.1 to reformulate the inner infimum problems for each

$i \in \mathcal{I}_1$. Lemma 4.3 offers a significant reduction in the computational complexity to solve the inner subproblems of (5). Instead of optimizing over d -dimensional vector x_i , the representation in Lemma 4.3 suggests that it suffices to search over a 1-dimensional space for k_i . While the objective function is still non-convex in k_i , we can perform a grid search over a compact interval to find the optimal solution for k_i to high precision. The grid search operations can also be parallelized across the index i thanks to the independent structure of the inner problems. Furthermore, the objective function of the supremum problem is a point-wise minimum of linear, thus concave, functions of γ . Hence, the outer problem is a concave maximization problem in γ , which can be solved using a golden section search algorithm.

4.2 Limiting Distribution

We now characterize the limit properties of $\mathcal{R}^{\text{OPP}}(\hat{\mathbb{P}}^N, \hat{p}^N)$. The next theorem asserts that the limiting distribution is of the chi-square type.

Theorem 4.4 (Limiting distribution – Probabilistic equal opportunity). Suppose that $(\hat{x}_i, \hat{a}_i, \hat{y}_i)$ are i.i.d. samples from \mathbb{P} . Under the null hypothesis $\mathcal{H}_0^{\text{OPP}}$, we have

$$N \times \mathcal{R}^{\text{OPP}}(\hat{\mathbb{P}}^N, \hat{p}^N) \xrightarrow{d} \theta \chi_1^2,$$

where χ_1^2 is a chi-square distribution with 1 degree of freedom,

$$\theta = \left(\mathbb{E}_{\mathbb{P}} \left[\left\| \nabla h_{\beta}(X) \left(\frac{\mathbb{1}_{(1,1)}(A, Y)}{p_{11}} - \frac{\mathbb{1}_{(0,1)}(A, Y)}{p_{01}} \right) \right\|_*^2 \right] \right)^{-1} \frac{\sigma_1^2}{p_{01}^2 p_{11}^2}$$

with $\sigma_1^2 = \text{Cov}(Z_1)$, and Z_1 is the random variable

$$\begin{aligned} Z_1 &= h_{\beta}(X) \left(p_{01} \mathbb{1}_{(1,1)}(A, Y) - p_{11} \mathbb{1}_{(0,1)}(A, Y) \right) \\ &\quad + \mathbb{1}_{(0,1)}(A, Y) \mathbb{E}_{\mathbb{P}}[\mathbb{1}_{(1,1)}(A, Y) h_{\beta}(X)] \\ &\quad - \mathbb{1}_{(1,1)}(A, Y) \mathbb{E}_{\mathbb{P}}[\mathbb{1}_{(0,1)}(A, Y) h_{\beta}(X)]. \end{aligned}$$

Construction of the hypothesis test. Based on the result of Theorem 4.4, the statistical hypothesis test proceeds as follows. Let $\eta_{1-\alpha}^{\text{OPP}}$ denote the $(1-\alpha) \times 100\%$ quantile of $\theta \chi_1^2$, where $\alpha \in (0, 1)$ is the predetermined significance level. By Theorem 4.4, the statistical decision has the form

$$\text{Reject } \mathcal{H}_0^{\text{OPP}} \text{ if } \hat{s}_N^{\text{OPP}} > \eta_{1-\alpha}^{\text{OPP}},$$

with

$$\hat{s}_N^{\text{OPP}} = N \times \mathcal{R}^{\text{OPP}}(\hat{\mathbb{P}}^N, \hat{p}^N).$$

The limiting distribution $\theta \chi_1^2$ is nonpivotal because θ depends on the true distribution \mathbb{P} . Luckily, because the quantile function of $\theta \chi_1^2$ is continuous in θ , if $\hat{\theta}_N$ is a consistent estimator of θ then it is also valid to use the quantile of $\hat{\theta}_N \chi_1^2$ for the purpose of testing. We thus proceed to discuss a consistent estimator $\hat{\theta}_N$ constructed from the available data. First, notice that \hat{p}_{01}^N and \hat{p}_{11}^N are consistent estimators for p_{01} and p_{11} . Similarly, the law of large numbers asserts that the denominator term in the definition of θ can be estimated

by the sample average

$$\begin{aligned} &\mathbb{E}_{\mathbb{P}} \left[\left\| \nabla h_{\beta}(X) \left(\frac{\mathbb{1}_{(1,1)}(A, Y)}{p_{11}} - \frac{\mathbb{1}_{(0,1)}(A, Y)}{p_{01}} \right) \right\|_*^2 \right] \\ &\approx \hat{T}^N = \frac{\|\beta\|_*^2}{N} \sum_{i=1}^N h_{\beta}(\hat{x}_i)^2 (1 - h_{\beta}(\hat{x}_i))^2 \left(\frac{\mathbb{1}_{(1,1)}(\hat{a}_i, \hat{y}_i)}{(\hat{p}_{11}^N)^2} + \frac{\mathbb{1}_{(0,1)}(\hat{a}_i, \hat{y}_i)}{(\hat{p}_{01}^N)^2} \right). \end{aligned}$$

Under the null hypothesis $\mathcal{H}_0^{\text{OPP}}$, Z_1 has mean 0. The sample average estimate of σ_1^2 is $\hat{\sigma}_1^N \approx (\hat{\sigma}^N)^2$ with

$$\begin{aligned} (\hat{\sigma}_1^N)^2 &= \frac{1}{N} \sum_{i=1}^N \left[h_{\beta}(\hat{x}_i) \left(p_{01} \mathbb{1}_{(1,1)}(\hat{a}_i, \hat{y}_i) - p_{11} \mathbb{1}_{(0,1)}(A, Y) \right) \right. \\ &\quad + \mathbb{1}_{(0,1)}(\hat{a}_i, \hat{y}_i) \left(\sum_{j=1}^N \mathbb{1}_{(1,1)}(\hat{a}_j, \hat{y}_j) h_{\beta}(\hat{x}_j) \right) \\ &\quad \left. - \mathbb{1}_{(1,1)}(\hat{a}_i, \hat{y}_i) \left(\sum_{j=1}^N \mathbb{1}_{(0,1)}(\hat{a}_j, \hat{y}_j) h_{\beta}(\hat{x}_j) \right) \right]^2. \end{aligned} \quad (7)$$

Using a nested arguments involving the continuous mapping theorem and Slutsky's theorem, the estimator

$$\hat{\theta}^N = \frac{(\hat{\sigma}_1^N)^2}{\hat{T}^N (\hat{p}_{01}^N)^2 (\hat{p}_{11}^N)^2}$$

is consistent for θ . Let the corresponding $(1-\alpha) \times 100\%$ quantile of the random variable $\hat{\theta}^N \chi_1^2$ be $\hat{\eta}_{1-\alpha}^{\text{OPP}}$. The statistical test decision using the plug-in consistent estimate becomes

$$\text{Reject } \mathcal{H}_0^{\text{OPP}} \text{ if } \hat{s}_N^{\text{OPP}} > \hat{\eta}_{1-\alpha}^{\text{OPP}}.$$

4.3 Most Favorable Distributions

We now discuss the construction of the most favorable distribution \mathbb{Q}^* , the projection of the empirical distribution $\hat{\mathbb{P}}^N$ onto the set $\mathcal{F}_{h_{\beta}}^{\text{OPP}}$. Intuitively, \mathbb{Q}^* is the distribution closest to $\hat{\mathbb{P}}^N$ that makes h_{β} a fair classifier under the equal opportunity criterion. If $\|\cdot\|$ is the Euclidean norm, the information about \mathbb{Q}^* can be recovered from the optimal solution of problem (6) by the result of the following lemma.

Lemma 4.5 (Most favorable distribution). Suppose that $\|\cdot\|$ is the Euclidean norm. Let γ^* be the optimal solution of problem (6), and for any $i \in \mathcal{I}_1$, let k_i^* be a solution of the inner minimization of (6) with respect to γ^* . Then the most favorable distribution $\mathbb{Q}^* = \arg \min_{\mathbb{Q} \in \mathcal{F}_{h_{\beta}}^{\text{OPP}}} \mathbb{W}(\hat{\mathbb{P}}^N, \mathbb{Q})$ is a discrete distribution of the form

$$\mathbb{Q}^* = \frac{1}{N} \left(\sum_{i \in \mathcal{I}_0} \delta_{(\hat{x}_i, \hat{a}_i, \hat{y}_i)} + \sum_{i \in \mathcal{I}_1} \delta_{(\hat{x}_i - k_i^* \gamma^* \lambda_i \beta, \hat{a}_i, \hat{y}_i)} \right).$$

By using the result of Lemma 4.3, it is easy to verify that \mathbb{Q}^* satisfies $\mathbb{W}(\mathbb{Q}^*, \hat{\mathbb{P}}^N)^2 = \mathcal{R}^{\text{OPP}}(\hat{\mathbb{P}}^N, \hat{p}^N)$. Moreover, one can also show that $\mathbb{Q}^* \in \mathcal{F}_{h_{\beta}}^{\text{OPP}}$. These two observations imply that \mathbb{Q}^* is the projection of $\hat{\mathbb{P}}^N$ onto $\mathcal{F}_{h_{\beta}}^{\text{OPP}}$. The detailed proof is omitted.

Lemma 4.5 suggests that in order to obtain the most favorable distribution, it suffices to perturb only the data points with positive label. This is intuitively rational because the notion of probabilistic equality of opportunity only depends on the positive label, and thus the perturbation with a minimal energy requirement should only

move sample points with $\hat{y}_i = 1$. When the underlying geometry is the Euclidean norm, the optimal perturbation of the point \hat{x}_i is to move it along a line dictated by β with a scaling factor $k_i^* \gamma^* \lambda_i$. Notice that λ_i defined in (4) are of opposite signs between samples of different sensitive attributes, which implies that it is optimal to perturb \hat{x}_i in opposite directions dependent on whether $\hat{a}_i = 0$ or $\hat{a}_i = 1$. This is, again, rational because moving points in opposite direction brings the clusters of points closer to the others, which reduces the discrepancy in the expected value of $h_\beta(X)$ between subgroups.

As a final remark, we note that \mathbb{Q}^* is not necessarily unique. This is because of the non-convexity of the inner problem over k_i in (6), which leads to the non-uniqueness of the optimal solution k_i^* (see Appendix B and Figure 5).

5 TESTING FAIRNESS FOR PROBABILISTIC EQUALIZED ODDS CRITERION

In this section, we extend the Wasserstein projection framework to the statistical test of probabilistic equalized odds for a pre-trained logistic classifier.

Definition 5.1 (Probabilistic equalized odds criterion [54]). A logistic classifier $h_\beta(\cdot) : \mathcal{X} \rightarrow [0, 1]$ satisfies the probabilistic equalized odds criteria relative to \mathbb{Q} if

$$\mathbb{E}_{\mathbb{Q}}[h_\beta(X)|A = 1, Y = y] = \mathbb{E}_{\mathbb{Q}}[h_\beta(X)|A = 0, Y = y] \quad \forall y \in \mathcal{Y}.$$

The notion of probabilistic equalized odds requires that the conditional expectation of h_β to be independent of A for any label subgroup, thus it is more stringent than the probabilistic equal opportunity studied in the previous section. We use the superscript “odd” in this section to emphasize on this specific notion of fairness. The definition of the probabilistic equalized odds prescribes the following set of distributions

$$\mathcal{F}_{h_\beta}^{\text{odd}} = \left\{ \begin{array}{l} \mathbb{Q} \in \mathcal{P} \text{ such that :} \\ \mathbb{E}_{\mathbb{Q}}[h_\beta(X)|A = 1, Y = 1] = \mathbb{E}_{\mathbb{Q}}[h_\beta(X)|A = 0, Y = 1] \\ \mathbb{E}_{\mathbb{Q}}[h_\beta(X)|A = 1, Y = 0] = \mathbb{E}_{\mathbb{Q}}[h_\beta(X)|A = 0, Y = 0] \end{array} \right\}.$$

Correspondingly, the Wasserstein projection hypothesis test for probabilistic equalized odds can be formulated as

$$\mathcal{H}_0^{\text{odd}} : \mathbb{P} \in \mathcal{F}_{h_\beta}^{\text{odd}}, \quad \mathcal{H}_1^{\text{odd}} : \mathbb{P} \notin \mathcal{F}_{h_\beta}^{\text{odd}}.$$

In the sequence, we study the projection onto the manifold $\mathcal{F}_{h_\beta}^{\text{odd}}$ in Section 5.1. Section 5.2 examines the asymptotic behaviour of the test statistic, and we close this section by studying the most favorable distribution \mathbb{Q}^* in Section 5.3.

5.1 Wasserstein Projection

Following a similar strategy as in Section 4, we define the set

$$\mathcal{F}_{h_\beta}^{\text{odd}}(\hat{p}^N) = \left\{ \begin{array}{l} \mathbb{Q} \in \mathcal{P} \text{ such that :} \\ (\hat{p}_{11}^N)^{-1} \mathbb{E}_{\mathbb{Q}}[h_\beta(X) \mathbb{1}_{(1,1)}(A, Y)] = (\hat{p}_{01}^N)^{-1} \mathbb{E}_{\mathbb{Q}}[h_\beta(X) \mathbb{1}_{(0,1)}(A, Y)] \\ (\hat{p}_{10}^N)^{-1} \mathbb{E}_{\mathbb{Q}}[h_\beta(X) \mathbb{1}_{(1,0)}(A, Y)] = (\hat{p}_{00}^N)^{-1} \mathbb{E}_{\mathbb{Q}}[h_\beta(X) \mathbb{1}_{(0,0)}(A, Y)] \\ \mathbb{Q}(A = a, Y = y) = \hat{p}_{ay}^N \quad \forall (a, y) \in \mathcal{A} \times \mathcal{Y} \end{array} \right\},$$

and the squared distance function

$$\mathcal{R}^{\text{odd}}(\hat{\mathbb{P}}^N, \hat{p}^N) = \begin{cases} \inf \mathbb{W}(\mathbb{Q}, \hat{\mathbb{P}}^N)^2 \\ \text{s.t. } \mathbb{E}_{\mathbb{Q}}[h_\beta(X)((\hat{p}_{11}^N)^{-1} \mathbb{1}_{(1,1)}(A, Y) - (\hat{p}_{01}^N)^{-1} \mathbb{1}_{(0,1)}(A, Y))] = 0 \\ \mathbb{E}_{\mathbb{Q}}[h_\beta(X)((\hat{p}_{10}^N)^{-1} \mathbb{1}_{(1,0)}(A, Y) - (\hat{p}_{00}^N)^{-1} \mathbb{1}_{(0,0)}(A, Y))] = 0 \\ \mathbb{E}_{\mathbb{Q}}[\mathbb{1}_{(a,y)}(A, Y)] = \hat{p}_{ay}^N \quad \forall (a, y) \in \mathcal{A} \times \mathcal{Y}. \end{cases}$$

The equivalent relation (3) suggests that the projection onto the set of distributions $\mathcal{F}_{h_\beta}^{\text{odd}}$ satisfies

$$\inf_{\mathbb{Q} \in \mathcal{F}_{h_\beta}^{\text{odd}}} \mathbb{W}(\hat{\mathbb{P}}^N, \mathbb{Q}) = \inf_{\mathbb{Q} \in \mathcal{F}_{h_\beta}^{\text{odd}}(\hat{p}^N)} \mathbb{W}(\hat{\mathbb{P}}^N, \mathbb{Q}) = \sqrt{\mathcal{R}^{\text{odd}}(\hat{\mathbb{P}}^N, \hat{p}^N)}.$$

The squared distance $\mathcal{R}^{\text{odd}}(\hat{\mathbb{P}}^N, \hat{p}^N)$ can be computed by solving the saddle point problem in the following proposition.

Proposition 5.2 (Dual reformulation). The squared projection distance $\mathcal{R}^{\text{odd}}(\hat{\mathbb{P}}^N, \hat{p}^N)$ equals to the optimal value of the following finite-dimensional optimization problem

$$\sup_{\gamma \in \mathbb{R}, \zeta \in \mathbb{R}} \frac{1}{N} \sum_{i=1}^N \inf_{x_i \in \mathcal{X}} \{ \|x_i - \hat{x}_i\|^2 + (\gamma \lambda_i \mathbb{1}_1(\hat{y}_i) + \zeta \lambda_i \mathbb{1}_0(\hat{y}_i)) h_\beta(x_i) \}. \quad (8)$$

To complete this section, we now discuss an efficient way to compute $\mathcal{R}^{\text{odd}}(\hat{\mathbb{P}}^N, \hat{p}^N)$. The next lemma reveals that computing $\mathcal{R}^{\text{odd}}(\hat{\mathbb{P}}^N, \hat{p}^N)$ can be decomposed into two subproblems of similar structure.

Lemma 5.3 (Univariate reduction). We have

$$\mathcal{R}^{\text{odd}}(\hat{\mathbb{P}}^N, \hat{p}^N) = \mathcal{R}^{\text{opp}}(\hat{\mathbb{P}}^N, \hat{p}^N) + U_N,$$

where U_N is computed as

$$U_N = \sup_{\zeta \in \mathbb{R}} \frac{1}{N} \sum_{i \in \mathcal{I}_0} \inf_{x_i \in \mathcal{X}} \{ \|x_i - \hat{x}_i\|^2 + \zeta \lambda_i h_\beta(x_i) \}.$$

Furthermore, if $\|\cdot\|$ is the Euclidean norm on \mathbb{R}^d , then

$$U_N = \sup_{\zeta \in \mathbb{R}} \frac{1}{N} \left\{ \sum_{i \in \mathcal{I}_0} \min_{k_i \in [0, \frac{1}{\beta}]} \zeta^2 \lambda_i^2 \|\beta\|_2^2 k_i^2 + \frac{\zeta \lambda_i}{1 + \exp(\zeta \lambda_i \|\beta\|_2^2 k_i - \beta^\top \hat{x}_i)} \right\}. \quad (9)$$

Notice that problem (9) has a similar structure to problem (6); the mere difference is that the summation in the objective function of (9) runs over the index set $\mathcal{I}_0 = \{i \in [N] : \hat{y}_i = 0\}$ instead of \mathcal{I}_1 in (6). Solving for U_N thus incurs the same computational complexity as, and can also be performed in parallel with, computing $\mathcal{R}^{\text{opp}}(\hat{\mathbb{P}}^N, \hat{p}^N)$.

5.2 Limiting Distribution

The next result asserts that the squared projection distance \mathcal{R}^{odd} has the $O(N^{-1})$ convergence rate.

Theorem 5.4 (Limiting distribution – Probabilistic equalized odds). Suppose that $(\hat{x}_i, \hat{a}_i, \hat{y}_i)$ are i.i.d. samples from \mathbb{P} . Under the null hypothesis $\mathcal{H}_0^{\text{odd}}$, we have

$$N \times \mathcal{R}^{\text{odd}}(\hat{\mathbb{P}}^N, \hat{p}^N) \xrightarrow{d} \sup_{\gamma, \zeta} \left\{ \gamma H_1 + \zeta H_0 + \mathbb{E}_{\mathbb{P}} \left[\left\| \begin{pmatrix} \gamma \\ \zeta \end{pmatrix}^\top \begin{pmatrix} p_{11}^{-1} \mathbb{1}_{(1,1)}(A, Y) - p_{01}^{-1} \mathbb{1}_{(0,1)}(A, Y) \\ p_{10}^{-1} \mathbb{1}_{(1,0)}(A, Y) - p_{00}^{-1} \mathbb{1}_{(0,0)}(A, Y) \end{pmatrix} \nabla h_{\beta}(X) \right\|_*^2 \right] \right\},$$

where $\nabla h_{\beta}(X) = h_{\beta}(X)(1 - h_{\beta}(X)\beta)$, and $H_y = \mathcal{N}(0, \sigma_y^2)/(p_{1y}p_{0y})$ with $\sigma_y^2 = \text{Cov}(Z_y)$, and Z_y are random variables

$$\begin{aligned} Z_y &= h_{\beta}(X) \left(p_{0y} \mathbb{1}_{(1,y)}(A, Y) - p_{1y} \mathbb{1}_{(0,y)}(A, Y) \right) \\ &\quad + \mathbb{1}_{(0,y)}(A, Y) \mathbb{E}_{\mathbb{P}}[\mathbb{1}_{(1,y)}(A, Y) h_{\beta}(X)] \\ &\quad - \mathbb{1}_{(1,y)}(A, Y) \mathbb{E}_{\mathbb{P}}[\mathbb{1}_{(0,y)}(A, Y) h_{\beta}(X)]. \end{aligned}$$

Construction of the hypothesis test. Contrary to the explicit chi-square limiting distribution for the probabilistic equal opportunity fairness in Theorem 4.4, the limiting distribution for the probabilistic equalized odds fairness is not available in closed form. Nevertheless, the limiting distribution in this case can be obtained by sampling H_0 and H_1 and solving a collection of optimization problems for each sample. Notice that the objective function of the supremum problem presented in Theorem 5.4 is continuous in H_1 and H_0 , one thus can define

$$\hat{H}_y = \mathcal{N}(0, \hat{\sigma}_y^2)/(\hat{p}_{1y}^N \hat{p}_{0y}^N),$$

where $\hat{\sigma}_y^2$ is the sample average estimate of σ_y^2 , which can be computed using an equation similar to (7). The limiting distribution can be computed by solving the optimization problem with plug-in values

$$\sup_{\gamma, \zeta} \left\{ \gamma \hat{H}_1 + \zeta \hat{H}_0 + \mathbb{E}_{\hat{\mathbb{P}}^N} \left[\left\| \begin{pmatrix} \gamma \\ \zeta \end{pmatrix}^\top \begin{pmatrix} (\hat{p}_{11}^N)^{-1} \mathbb{1}_{(1,1)}(A, Y) - (\hat{p}_{01}^N)^{-1} \mathbb{1}_{(0,1)}(A, Y) \\ (\hat{p}_{10}^N)^{-1} \mathbb{1}_{(1,0)}(A, Y) - (\hat{p}_{00}^N)^{-1} \mathbb{1}_{(0,0)}(A, Y) \end{pmatrix} \nabla h_{\beta}(X) \right\|_*^2 \right] \right\}.$$

Notice that the expectation is taken over the empirical distribution $\hat{\mathbb{P}}^N$, and can be written as a finite sum. The last optimization problem can be solved efficiently using quadratic programming for any realization of \hat{H}_1 and \hat{H}_0 . The objective values can be collected to compute the $(1 - \alpha) \times 100\%$ -quantile estimate $\hat{\eta}_{1-\alpha}^{\text{odd}}$ of the limiting distribution. The statistical test decision using the plug-in estimate becomes

$$\text{Reject } \mathcal{H}_0^{\text{odd}} \text{ if } \hat{s}_N^{\text{odd}} > \hat{\eta}_{1-\alpha}^{\text{odd}},$$

where $\hat{s}_N^{\text{odd}} = N \times \mathcal{R}^{\text{odd}}(\hat{\mathbb{P}}^N, \hat{p}^N)$.

5.3 Most Favorable Distributions

If the feature space \mathcal{X} is endowed with an Euclidean norm, then the most favorable distribution \mathbb{Q}^* , defined in this section as the projection of $\hat{\mathbb{P}}^N$ onto $\mathcal{F}_{h_{\beta}}^{\text{odd}}$, can be constructed by exploiting Lemma 5.3.

Lemma 5.5 (Most favorable distribution). Suppose that $\|\cdot\|$ is the Euclidean norm. Let γ^* and ζ^* be the optimal solution of

problems (6) and (9), respectively. For any $i \in \mathcal{I}_1$, let k_i^* be the solution of the inner minimization of (6) with respect to γ^* , and for any $i \in \mathcal{I}_0$, let k_i^* be a solution of the inner minimization of (9) with respect to ζ^* . Then the most favorable distribution $\mathbb{Q}^* = \arg \min_{\mathbb{Q} \in \mathcal{F}_{h_{\beta}}^{\text{odd}}} \mathbb{W}(\hat{\mathbb{P}}^N, \mathbb{Q})$ is a discrete distribution of the form

$$\mathbb{Q}^* = \frac{1}{N} \left(\sum_{i \in \mathcal{I}_0} \delta_{(\hat{x}_i - k_i^* \zeta^* \lambda_i \beta, \hat{a}_i, \hat{y}_i)} + \sum_{i \in \mathcal{I}_1} \delta_{(\hat{x}_i - k_i^* \gamma^* \lambda_i \beta, \hat{a}_i, \hat{y}_i)} \right).$$

The proof of Lemma 5.5 follows from verifying that $\mathbb{Q}^* \in \mathcal{F}_{h_{\beta}}^{\text{odd}}$ and that $\mathbb{W}(\mathbb{Q}^*, \hat{\mathbb{P}}^N)^2 = \mathcal{R}^{\text{odd}}(\hat{\mathbb{P}}^N, \hat{p}^N)$ using Lemma 5.3, the detailed proof is omitted. For probabilistic equalized odds, the most favorable distribution \mathbb{Q}^* alters the locations of both $i \in \mathcal{I}_0$ and $i \in \mathcal{I}_1$. The directions of perturbation are dependent on λ_i , which is determined using (4). Notice that λ_i carry opposite signs corresponding to whether $\hat{a}_i = 0$ or $\hat{a}_i = 1$, thus the perturbations will move \hat{x}_i in opposite directions based on the value of the sensitive attribute \hat{a}_i .

6 NUMERICAL EXPERIMENT

All experiments are run on an Intel Xeon based cluster composed of 287 compute nodes each with 2 Skylake processors running at 2.3 GHz with 18 cores each. We only use 2 nodes of this cluster and all optimization problems are implemented in Python version 3.7.3. In all experiments, we use the 2-norm to measure distances in the feature space. Moreover, we focus on the hypothesis test of probabilistic equal opportunity, and thus the Wasserstein projection, the limiting distribution and the most favorable distribution follow from the results presented in Section 4.

6.1 Validation of the Hypothesis Test

We now demonstrate that our proposed Wasserstein projection framework for statistical test of fairness is a valid, or asymptotically correct, test. We consider a binary classification setting in which \mathcal{X} is 2-dimensional feature space. The true distribution \mathbb{P} has true marginal values p_{ay} being

$$p_{11} = 0.2, p_{01} = 0.1, p_{10} = 0.3, p_{00} = 0.4.$$

Moreover, conditioning on (A, Y) , the feature X follows a Gaussian distribution of the form

$$\begin{aligned} X|A=1, Y=1 &\sim \mathcal{N}([6, 0], [3.5, 0; 0, 5]), \\ X|A=0, Y=1 &\sim \mathcal{N}([-2, 0], [5, 0; 0, 5]), \\ X|A=1, Y=0 &\sim \mathcal{N}([6, 0], [3.5, 0; 0, 5]), \\ X|A=0, Y=0 &\sim \mathcal{N}([-4, 0], [5, 0; 0, 5]). \end{aligned}$$

The true distribution \mathbb{P} is thus a mixture of Gaussian, and under this specification, a simple algebraic calculation indicates that a logistic classifier with $\beta = (0, 1)^\top$ is fair with respect to the probabilistic equal opportunity criterion in Definition 4.1. We thus focus on verifying fairness for this specific classifier. In the first experiment, we empirically validate Theorem 4.4. To this end, we generate $N \in \{100, 500\}$ i.i.d. samples from \mathbb{P} to be used as the test data, and then calculate the squared projection distance $\mathcal{R}^{\text{OPP}}(\hat{\mathbb{P}}^N, \hat{p}^N)$ using Proposition 4.2. The process is repeated 2,000 times to obtain an empirical estimate of the distribution of $N \times \mathcal{R}^{\text{OPP}}(\hat{\mathbb{P}}^N, \hat{p}^N)$.




Figure 2: Empirical distribution of $N \times \mathcal{R}^{\text{opp}}(\hat{\mathbb{P}}^N, \hat{p}^N)$ taken over 2,000 replications (histogram) versus the limiting distribution $\theta\chi_1^2$ (blue curve) with different sample sizes N . Fig. 2a-2b are density plots, Fig. 2c-2d are cumulative distribution plots.

We also generate another set of one million i.i.d. samples from \mathbb{P} to estimate the limiting distribution $\theta\chi_1^2$. Figure 2 shows that the empirical distribution of $N \times \mathcal{R}^{\text{opp}}(\hat{\mathbb{P}}^N, \hat{p}^N)$ converges to the limiting distribution $\theta\chi_1^2$ as N increases.

The second set of experiments aims to show that our proposed Wasserstein projection hypothesis test is asymptotically valid. We generate $N \in \{100, 500, 1000\}$ i.i.d. samples from \mathbb{P} and calculate the test statistic $N \times \mathcal{R}^{\text{opp}}(\hat{\mathbb{P}}^N, \hat{p}^N)$. The same data is used to estimate $\hat{\theta}^N$ and compute the $(1 - \alpha) \times 100\%$ -quantile of $\hat{\theta}^N \chi_1^2$ to perform the quantile based test as laid out in Section 4.2. We repeat this procedure for 2,000 replications to keep track of the rejection projection at different significant values of $\alpha \in \{0.5, 0.3, 0.1, 0.05, 0.01\}$. Table 1 summarizes the rejection probabilities of Wasserstein projection tests for equal opportunity criterion under the null hypothesis $\mathcal{H}_0^{\text{opp}}$. We can observe that at sample size $N > 100$, the rejection probability is close to the desired level α , which empirically validates our testing procedure.

$N = 100$	$N = 500$	$N = 1000$	α
0.511	0.4905	0.5	0.50
0.282	0.2895	0.299	0.30
0.048	0.0895	0.093	0.10
0.007	0.0425	0.0405	0.05
0.0	0.0065	0.005	0.01

Table 1: Comparison of the null rejection probabilities of probabilistic equal opportunity tests with different significance levels α and test sample sizes N .

6.2 Most Favorable Distribution Analysis

In this section, we visualize the most favorable distribution \mathbb{Q}^* from Lemma 4.5 for a vanilla logistic regression classifier with weight $\beta = (0.4, 0.12)^\top$. We simply generate 28 samples with equal subgroup proportions to form the empirical distribution $\hat{\mathbb{P}}^N$. To find the support of \mathbb{Q}^* , we solve problem (6), whose optimizer dictates the transportation plan of each sample \hat{x}_i . Figure 3 visualizes the original test samples that forms $\hat{\mathbb{P}}^N$, along with the most favorable distribution \mathbb{Q}^* . Green lines in the figure represent how samples are perturbed. As we are testing for the probabilistic notion of equal opportunity, only the samples with positive label $\hat{y}_i = 1$




Figure 3: Visualization of the most favorable distribution \mathbb{Q}^* for a logistic classifier with weight $\beta = (0.4, 0.12)^\top$. The black arrow indicates the vector β . Colors represent class, while symbolic shapes encode the sensitive values. The green lines show the transport plan of the empirical test samples from their original positions (indicated with transparent colors) to their ultimate destinations (with non-transparent colors).

presented in blue are perturbed in order to obtain \mathbb{Q}^* . Furthermore, we observe that the positively-labeled test samples are transported along the axis directed by β (black arrow). Moreover, the samples with different sensitive attributes, represented by different shapes, move in opposite direction so that they get closer to each other, which reduces the discrepancy in the expected value of $h_\beta(X)$ between the relevant subgroups.

6.3 The COMPAS Dataset

COMPAS (Correctional Offender Management Profiling for Alternative Sanctions)⁴ is a commercial tool used by judges and parole officers for scoring criminal defendant’s likelihood of recidivism. The COMPAS dataset is used by the COMPAS algorithm to compute the risk score of reoffending for defendants, and also contains the criminal records within 2 years after the decision. The dataset consists of 6,172 samples with 10 attributes including gender, age category, race, etc. We concentrate on the subset of the data with violent recidivism, and we use race (African-American and Caucasian) as the sensitive attribute. We split 70% of the COMPAS data to train a Tikhonov-regularized logistic classifier, with the tuning penalty parameter λ chosen in the range from 0 to 100 with 50

⁴<https://www.propublica.org/datastore/dataset/compas-recidivism-risk-score-data-and-analysis>

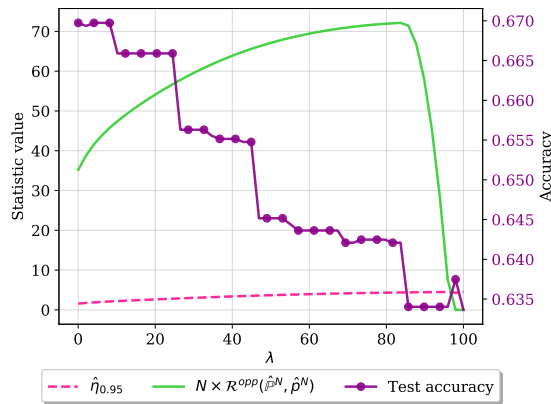


Figure 4: Test statistic and accuracy of Tikhonov regularized logistic regression on test data with rejection threshold $\hat{\eta}_{0.95}$.

equi-distant points. The remaining 30% of the data is used as the test samples for auditing.

Figure 4 demonstrates the relation between the accuracy and the degree of fairness with respect to the regularization parameter λ . Strong regularization penalty (high values of λ) results in small values of the test statistic, but the classifier has low test accuracy. On the contrary, weak penalization leads to undesirable fairness level but higher prediction accuracy. The pink dashed line in Figure 4 shows the rejection threshold of the Wasserstein projection test at significance level $\alpha = 0.05$ for varying value of the regularization parameter λ . We can observe that the Wasserstein projection test recommends a rejection of the null hypothesis $\mathcal{H}_0^{\text{OPP}}$ for a wide range of λ . Only at λ sufficiently large that the test fails to reject the null hypothesis.

7 CONCLUDING REMARKS AND BROADER IMPACT

In this paper, we propose a statistical hypothesis test for group fairness of classification algorithms based on the theory of optimal transport. Our test statistic relies on computing the projection distance from the empirical distribution supported on the test samples to the manifold of distributions that renders the classifier fair. When the notion of fairness is chosen to be either the probabilistic equal opportunity or the probabilistic equalized odds, we show that the projection can be computed efficiently. We provide the limiting distribution of the test statistic and show that our Wasserstein projection test is asymptotically correct. Our proposed test also offers the flexibility to incorporate the geometric information of the feature space into testing procedure. Finally, analyzing the most favorable distribution can help interpreting the reasons behind the outcome of the test.

The Wasserstein projection hypothesis test is the culmination of a benevolent motivation and effort, and it aims to furnish the developers, the regulators and the general public a quantitative method to verify certain notions of fairness in the classification setting. At the same time, we acknowledge the risks and limitations of the results presented in this paper.

First, it is essential to keep in mind that this paper focuses on *probabilistic* notions of fairness, in particular, we provide the Wasserstein statistical test for probabilistic equality of opportunity

and probabilistic equalized odds. Probabilistic notions are only approximations of the original definitions, and the employment of probabilistic notions are solely for the technical purposes. Due to the sensitivity of the test result on the choice of fairness notions, a test that is designed for probabilistic notions may not be applicable to test for original notions of fairness due to the interplay with the threshold τ and the radical difference of both the test statistic and the limiting distribution. If a logistic classifier h_β is rejected using our framework for probabilistic equal opportunity, it does *not* necessarily imply that the classifier h_β fails to satisfy the equal opportunity criterion, and vice versa. The same argument holds when we test for probabilistic equalized odds.

Second, the outcome of the Wasserstein projection test is dependent on the choice of the underlying metric on the feature, the sensitive attribute and the label spaces. Indeed, the test outcome can change if we switch the metric of the feature space, for example, from the Euclidean norm to a 1-norm. In the scope of this paper, we do not study how sensitive the test outcome is with respect to the choice of the metric, nor can we make any recommendation on the optimal choice of the metric. Nevertheless, it is reasonable to recommend that the metric should be chosen judiciously, and the action of tuning the metric in order to obtain favorable test outcome should be prohibited.

Third, to simplify the computation, we have assumed absolute trust on the sensitive attributes and the label. The users of our test should be mindful if there is potential corruption to these values. Moreover, our test is constructed under the assumption that there is no missing values in the test data. This assumption, unfortunately, may not hold in real-world implementations. Constructing statistical test which is robust to adversarial attacks and missing data using the Wasserstein projection framework is an interesting research direction.

Fourth, the statistical test in this paper is for a simple null hypothesis. In practice, the regulators may be interested in a relaxed fairness test in which the difference of the conditional expectations is upper bounded by a fixed positive constant ϵ . The extension of the Wasserstein hypothesis testing framework for a composite null hypothesis is non-trivial, thus we leave this idea for future study.

Finally, any auditing process for algorithmic fairness can become a dangerous tool if it falls into the hand of unqualified or vicious inspectors. The results in this paper are developed to broaden our scientific understanding, and we recommend that the test and its outcomes should be used as an informative reference, but *not* as an absolute certification to promote any particular classifier or as a justification for any particular classification decision.

We thus sincerely recommend that the tools proposed in this paper be exercised with utmost consideration.

ACKNOWLEDGMENTS

Research supported by the Swiss National Science Foundation under NCCR Automation, grant agreement 51NF40_180545. Material in this paper is based upon work supported by the Air Force Office of Scientific Research under award number FA9550-20-1-0397. Additional support is gratefully acknowledged from NSF grants 1915967, 1820942, 1838676, and also from the China Merchant Bank. Finally, we would like to thank Nian Si and Michael Sklar for helpful comments and discussions.

REFERENCES

- [1] David Alvarez-Melis, Tommi S Jaakkola, and Stefanie Jegelka. 2017. Structured optimal transport. *arXiv preprint arXiv:1712.06199* (2017).
- [2] Solon Barocas and Andrew D Selbst. 2016. Big data's disparate impact. *California Law Review* 104 (2016), 671–732.
- [3] Rachel KE Bellamy, Kuntal Dey, Michael Hind, Samuel C Hoffman, Stephanie Houde, Kalapriya Kannan, Pranay Lohia, Jacquelyn Martino, Sameep Mehta, Aleksandra Mojsilovic, et al. 2018. AI Fairness 360: An extensible toolkit for detecting, understanding, and mitigating unwanted algorithmic bias. *arXiv preprint arXiv:1810.01943* (2018).
- [4] Jean-David Benamou, Guillaume Carlier, Marco Cuturi, Luca Nenna, and Gabriel Peyré. 2015. Iterative Bregman projections for regularized transportation problems. *SIAM Journal on Scientific Computing* 37, 2 (2015), A1111–A1138.
- [5] Richard Berk, Hoda Heidari, Shahin Jabbari, Michael Kearns, and Aaron Roth. 2018. Fairness in criminal justice risk assessments: The state of the art. *Sociological Methods & Research* (2018), 0049124118782533.
- [6] Emily Black, Samuel Yeom, and Matt Fredrikson. 2020. FlipTest: fairness testing via optimal transport. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*. 111–121.
- [7] Jose Blanchet, Yang Kang, and Karthyek Murthy. 2019. Robust Wasserstein profile inference and applications to machine learning. *Journal of Applied Probability* 56, 3 (2019), 830–857.
- [8] Mathieu Blondel, Vivien Seguy, and Antoine Rolet. 2018. Smooth and sparse optimal transport. In *International Conference on Artificial Intelligence and Statistics*. 880–889.
- [9] Joy Buolamwini and Timnit Gebru. 2018. Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Conference on Fairness, Accountability and Transparency*. 77–91.
- [10] Toon Calders and Sicco Verwer. 2010. Three naive Bayes approaches for discrimination-free classification. *Data Mining and Knowledge Discovery* 21, 2 (2010), 277–292.
- [11] Elsa Cazelles, Vivien Seguy, Jérémie Bigot, Marco Cuturi, and Nicolas Papadakis. 2018. Geodesic PCA versus log-PCA of histograms in the Wasserstein space. *SIAM Journal on Scientific Computing* 40, 2 (2018), B429–B456.
- [12] Alexandra Chouldechova. 2017. Fair prediction with disparate impact: A study of bias in recidivism prediction instruments. *Big Data* 5, 2 (2017), 153–163.
- [13] Alexandra Chouldechova and Aaron Roth. 2020. A snapshot of the frontiers of fairness in machine learning. *Commun. ACM* 63, 5 (2020), 82–89.
- [14] Sam Corbett-Davies, Emma Pierson, Avi Feller, Sharad Goel, and Aziz Huq. 2017. Algorithmic decision making and the cost of fairness. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 797–806.
- [15] Nicolas Courty, Rémi Flamary, Devis Tuia, and Alain Rakotomamonjy. 2016. Optimal transport for domain adaptation. *IEEE transactions on pattern analysis and machine intelligence* 39, 9 (2016), 1853–1865.
- [16] Marco Cuturi. 2013. Sinkhorn Distances: Lightspeed Computation of Optimal Transport. In *Advances in Neural Information Processing Systems*. 2292–2300.
- [17] Jeffrey Dastin. 2018. Amazon scraps secret AI recruiting tool that showed bias against women. *San Fransico, CA: Reuters*. Retrieved on October 9 (2018), 2018.
- [18] Amit Datta, Michael Carl Tschantz, and Anupam Datta. 2015. Automated experiments on ad privacy settings: A tale of opacity, choice, and discrimination. *Proceedings on Privacy Enhancing Technologies* 2015, 1 (2015), 92–112.
- [19] Cyrus DiCiccio, Sriram Vasudevan, Kinjal Basu, Krishnarum Kenthapadi, and Deepak Agarwal. 2020. Evaluating fairness using permutation tests. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 1467–1477.
- [20] Pavel Dvurechensky, Alexander Gasnikov, and Alexey Kroshnin. 2018. Computational optimal transport: Complexity by accelerated gradient descent is better than by Sinkhorn's algorithm. *arXiv preprint arXiv:1802.04367* (2018).
- [21] Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard Zemel. 2012. Fairness through awareness. In *Proceedings of the 3rd innovations in theoretical computer science conference*. 214–226.
- [22] Mahyar Fazlyab, Manfred Morari, and George J. Pappas. 2019. Safety Verification and Robustness Analysis of Neural Networks via Quadratic Constraints and Semidefinite Programming. *arXiv preprint arXiv:1903.01287* (2019).
- [23] Michael Feldman, Sorelle A. Friedler, John Moeller, Carlos Scheidegger, and Suresh Venkatasubramanian. 2015. Certifying and Removing Disparate Impact. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 259–268.
- [24] Sira Ferradans, Nicolas Papadakis, Gabriel Peyré, and Jean-François Aujol. 2014. Regularized discrete optimal transport. *SIAM Journal on Imaging Sciences* 7, 3 (2014), 1853–1882.
- [25] Rémi Flamary, Marco Cuturi, Nicolas Courty, and Alain Rakotomamonjy. 2018. Wasserstein discriminant analysis. *Machine Learning* 107, 12 (2018), 1923–1945.
- [26] Rui Gao, Xi Chen, and Anton J Kleywegt. 2017. Wasserstein distributional robustness and regularization in statistical learning. *arXiv preprint arXiv:1712.06050* (2017).
- [27] Sahaj Garg, Vincent Perot, Nicole Limtiaco, Ankur Taly, Ed H Chi, and Alex Beutel. 2019. Counterfactual fairness in text classification through robustness. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*. 219–226.
- [28] Aude Genevay, Marco Cuturi, Gabriel Peyré, and Francis Bach. 2016. Stochastic optimization for large-scale optimal transport. In *Advances in Neural Information Processing Systems*. 3440–3448.
- [29] Paula Gordaliza, Eustasio Del Barrio, Gamboa Fabrice, and Jean-Michel Loubes. 2019. Obtaining Fairness using Optimal Transport Theory. In *Proceedings of the 36th International Conference on Machine Learning*. 2357–2365.
- [30] Nina Grgic-Hlaca, Muhammad Bilal Zafar, Krishna P Gummadi, and Adrian Weller. 2016. The case for process fairness in learning: Feature selection for fair decision making. In *NIPS Symposium on Machine Learning and the Law*, Vol. 1. 2.
- [31] Moritz Hardt, Eric Price, Eric Price, and Nati Srebro. 2016. Equality of Opportunity in Supervised Learning. In *Advances in Neural Information Processing Systems* 29. 3315–3323.
- [32] Nhat Ho, Xuan Long Nguyen, Mikhail Yurochkin, Hung Hai Bui, Viet Huynh, and Dinh Phung. 2017. Multilevel clustering via Wasserstein means. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*. JMLR. org, 1501–1509.
- [33] David W Hosmer Jr, Stanley Lemeshow, and Rodney X Sturdivant. 2013. *Applied Logistic Regression*. John Wiley & Sons.
- [34] Philips George John, Deepak Vijaykeerthy, and Diptikalyan Saha. 2020. Verifying Individual Fairness in Machine Learning Models. In *Conference on Uncertainty in Artificial Intelligence*. PMLR, 749–758.
- [35] Nathan Kallus, Xiaojie Mao, and Angela Zhou. 2019. Assessing algorithmic fairness with unobserved protected class using data combination. *arXiv preprint arXiv:1906.00285* (2019).
- [36] Jon Kleinberg, Jens Ludwig, Sendhil Mullainathan, and Ashesh Rambachan. 2018. Algorithmic fairness. In *AEA Papers and Proceedings*, Vol. 108. 22–27.
- [37] Jon Kleinberg, Sendhil Mullainathan, and Manish Raghavan. 2016. Inherent trade-offs in the fair determination of risk scores. *arXiv preprint arXiv:1609.05807* (2016).
- [38] Soheil Kolouri, Se Rim Park, Matthew Thorpe, Dejan Slepcev, and Gustavo K Rohde. 2017. Optimal mass transport: Signal processing and machine-learning applications. *IEEE signal processing magazine* 34, 4 (2017), 43–59.
- [39] Soheil Kolouri and Gustavo K Rohde. 2015. Transport-based single frame super resolution of very low resolution face images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 4876–4884.
- [40] Daniel Kuhn, Peyman Mohajerini Esfahani, Viet Anh Nguyen, and Soroosh Shafieezadeh-Abadeh. 2019. Wasserstein distributionally robust optimization: Theory and applications in machine learning. In *Operations Research & Management Science in the Age of Analytics*. INFORMS, 130–166.
- [41] Zachary Lipton, Julian McAuley, and Alexandra Chouldechova. 2018. Does mitigating ML's impact disparity require treatment disparity?. In *Advances in Neural Information Processing Systems*. 8125–8135.
- [42] Michael Lohaus, Michaël Perrot, and Ulrike von Luxburg. 2020. Too Relaxed to Be Fair. In *International Conference on Machine Learning*.
- [43] Arjun K Manrai, Birgit H Funke, Heidi L Rehm, Morten S Olesen, Bradley A Maron, Peter Szolovits, David M Margulies, Joseph Loscalzo, and Isaac S Kohane. 2016. Genetic misdiagnoses and the potential for health disparities. *New England Journal of Medicine* 375, 7 (2016), 655–665.
- [44] Ninareh Mehrabi, Fred Morstatter, Nripsuta Saxena, Kristina Lerman, and Aram Galstyan. 2019. A survey on bias and fairness in machine learning. *arXiv preprint arXiv:1908.09635* (2019).
- [45] Gaspard Monge. 1781. Mémoire sur la théorie des déblais et des remblais. *Histoire de l'Académie Royale des Sciences de Paris* (1781).
- [46] MultiMedia LLC. 2016 (accessed June 4, 2020). *Machine Bias*. Available at <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>.
- [47] Zak Murez, Soheil Kolouri, David Kriegman, Ravi Ramamoorthi, and Kyungnam Kim. 2018. Image to image translation for domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 4500–4509.
- [48] Katta G Murty and Santosh N Kabadi. 1985. *Some NP-complete problems in quadratic and nonlinear programming*. Technical Report.
- [49] XuanLong Nguyen. 2013. Convergence of latent mixing measures in finite and infinite mixture models. *The Annals of Statistics* 41, 1 (2013), 370–400.
- [50] Nicolas Papadakis and Julien Rabin. 2017. Convex Histogram-Based Joint Image Segmentation with Regularized Optimal Transport Cost. *Journal of Mathematical Imaging and Vision* 59, 2 (2017), 161–186.
- [51] Ofir Pele and Michael Werman. 2008. A linear time histogram metric for improved sift matching. In *European conference on computer vision*. Springer, 495–508.
- [52] O. Pele and M. Werman. 2009. Fast and robust Earth Mover's Distances. In *2009 IEEE 12th International Conference on Computer Vision*. 460–467. <https://doi.org/10.1109/ICCV.2009.5459199>
- [53] Gabriel Peyré and Marco Cuturi. 2019. Computational optimal transport. *Foundations and Trends® in Machine Learning* 11, 5-6 (2019), 355–607.
- [54] Geoff Pleiss, Manish Raghavan, Felix Wu, Jon Kleinberg, and Kilian Q Weinberger. 2017. On fairness and calibration. In *Advances in Neural Information Processing*

- Systems. 5680–5689.
- [55] Antoine Rolet, Marco Cuturi, and Gabriel Peyré. 2016. Fast dictionary learning with a smoothed Wasserstein loss. In *Artificial Intelligence and Statistics*. 630–638.
- [56] Yossi Rubner, Carlo Tomasi, and Leonidas J Guibas. 2000. The earth mover’s distance as a metric for image retrieval. *International journal of computer vision* 40, 2 (2000), 99–121.
- [57] Pedro Saleiro, Benedict Kuester, Loren Hinkson, Jesse London, Abby Stevens, Ari Anisfeld, Kit T Rodolfa, and Rayid Ghani. 2018. Aequitas: A bias and fairness audit toolkit. *arXiv preprint arXiv:1811.05577* (2018).
- [58] Bernhard Schmitzer. 2016. A sparse multiscale algorithm for dense optimal transport. *Journal of Mathematical Imaging and Vision* 56, 2 (2016), 238–259.
- [59] Vivien Seguy and Marco Cuturi. 2015. Principal geodesic analysis for probability measures under the optimal transport metric. In *Advances in Neural Information Processing Systems*. 3312–3320.
- [60] James E Smith. 1995. Generalized Chebychev inequalities: Theory and applications in decision analysis. *Operations Research* 43, 5 (1995), 807–825.
- [61] Justin Solomon, Fernando De Goes, Gabriel Peyré, Marco Cuturi, Adrian Butscher, Andy Nguyen, Tao Du, and Leonidas Guibas. 2015. Convolutional Wasserstein distances: Efficient optimal transportation on geometric domains. *ACM Transactions on Graphics (TOG)* 34, 4 (2015), 66.
- [62] Justin Solomon, Raif Rustamov, Leonidas Guibas, and Adrian Butscher. 2014. Earth mover’s distances on discrete surfaces. *ACM Transactions on Graphics (TOG)* 33, 4 (2014), 67.
- [63] Guillaume Tartavel, Gabriel Peyré, and Yann Gousseau. 2016. Wasserstein loss for image synthesis and restoration. *SIAM Journal on Imaging Sciences* 9, 4 (2016), 1726–1755.
- [64] Bahar Taskesen, Viet Anh Nguyen, Daniel Kuhn, and Jose Blanchet. 2020. A Distributionally Robust Approach to Fair Classification. *arXiv preprint arXiv:2007.09530* (2020).
- [65] Matthew Thorpe, Serim Park, Soheil Kolouri, Gustavo K Rohde, and Dejan Slepčev. 2017. A Transportation L^p Distance for Signal Analysis. *Journal of mathematical imaging and vision* 59, 2 (2017), 187–210.
- [66] Florian Tramer, Vaggelis Atlidakis, Roxana Geambasu, Daniel Hsu, Jean-Pierre Hubaux, Mathias Humbert, Ari Juels, and Huang Lin. 2017. FairTest: Discovering unwarranted associations in data-driven applications. In *2017 IEEE European Symposium on Security and Privacy (EuroS&P)*. IEEE, 401–416.
- [67] James Wexler, Mahima Pushkarna, Tolga Bolukbasi, Martin Wattenberg, Fernanda Viégas, and Jimbo Wilson. 2019. The what-if tool: Interactive probing of machine learning models. *IEEE transactions on visualization and computer graphics* 26, 1 (2019), 56–65.
- [68] Songkai Xue, Mikhail Yurochkin, and Yuekai Sun. 2020. Auditing ML Models for Individual Bias and Unfairness. *arXiv preprint arXiv:2003.05048* (2020).
- [69] Mikhail Yurochkin, Amanda Bower, and Yuekai Sun. 2020. Training individually fair ML models with sensitive subspace robustness. In *International Conference on Learning Representations*.
- [70] 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*. 1171–1180.
- [71] Muhammad Bilal Zafar, Isabel Valera, Manuel Gomez Rodriguez, and Krishna P Gummadi. 2017. Fairness constraints: Mechanisms for fair classification. *AISTATS* (2017).
- [72] C. Zhao and Y. Guan. 2018. Data-driven risk-averse stochastic optimization with Wasserstein metric. *Operations Research Letters* 46, 2 (2018), 262 – 267.

A APPENDIX - PROOFS

A.1 Proofs of Section 2

PROOF OF LEMMA 3.2. Because the fairness constraints are similar in both sets \mathcal{F}_h and $\mathcal{F}_h(\hat{p}^N)$, it thus suffice to verify that \mathbb{Q} satisfies the marginal conditions $\mathbb{Q}(A = a, Y = y) = \hat{p}_{ay}^N$ for all $(a, y) \in \mathcal{A} \times \mathcal{Y}$. By the definition of the Wasserstein distance and the ground metric c , there exists a coupling π such that

$$\mathbb{W}(\hat{\mathbb{P}}^N, \mathbb{Q})^2 = \mathbb{E}_\pi[(\|X' - X\| + \infty|A' - A| + \infty|Y' - Y|)^2]$$

and the marginal distribution of π are $\hat{\mathbb{P}}^N$ and \mathbb{Q} , respectively. By the law of total probability and because $\hat{\mathbb{P}}^N$ is an empirical distribution, we can write $\pi = N^{-1} \sum_{i=1}^N \delta_{(\hat{x}_i, \hat{a}_i, \hat{y}_i)} \otimes \mathbb{Q}_i$, where \mathbb{Q}_i denotes the conditional distributions of (X, A, Y) given $(X', A', Y') = (\hat{x}_i, \hat{a}_i, \hat{y}_i)$ for all $i \in [N]$.

Suppose without any loss of generality that there exists a tuple $(a, y) \in \mathcal{A} \times \mathcal{Y}$ such that $\mathbb{Q}(A = a, Y = y) > \hat{p}_{ay}^N$. This means

$$\begin{aligned} \mathbb{Q}(A = a, Y = y) &= \frac{1}{N} \sum_{i=1}^N \mathbb{Q}_i(A = a, Y = y) \\ &> \frac{1}{N} \sum_{i=1}^N \mathbb{1}_{(a, y)}(\hat{a}_i, \hat{y}_i). \end{aligned}$$

This implies that there must exist an index $i^* \in [N]$ with $(\hat{a}_{i^*}, \hat{y}_{i^*}) \neq (a, y)$, and that

$$\mathbb{Q}_{i^*}(A = a, Y = y) > 0.$$

However, this further implies that

$$\begin{aligned} \mathbb{W}(\hat{\mathbb{P}}^N, \mathbb{Q})^2 &= \frac{1}{N} \sum_{i=1}^N \mathbb{E}_{\mathbb{Q}_i}[(\|\hat{x}_i - X\| + \infty|\hat{a}_i - A| + \infty|\hat{y}_i - Y|)^2] \\ &\geq \frac{1}{N} \mathbb{E}_{\mathbb{Q}_{i^*}}[(\|\hat{x}_{i^*} - X\| + \infty|\hat{a}_{i^*} - A| + \infty|\hat{y}_{i^*} - Y|)^2] \\ &\geq \frac{1}{N} \mathbb{Q}_{i^*}(A = a, Y = y) (\infty(\hat{a}_{i^*} - a) + \infty(\hat{y}_{i^*} - y))^2 \\ &= \infty, \end{aligned}$$

where the equality follows from the decomposition of π using the law of total probability and the first inequality follows because the transportation cost is nonnegative. This contradicts the fact that $\mathbb{W}(\hat{\mathbb{P}}^N, \mathbb{Q}) < \infty$. \square

A.2 Proofs of Section 4

Before proving Proposition 4.2, we first prove a preparatory lemma that verifies the Slater condition of the conic optimization problem. To shorten the notation, we write $\xi = (X, A, Y)$ and denote $\Xi = \mathcal{X} \times \mathcal{A} \times \mathcal{Y}$, $\hat{\Xi}_N = \{(\hat{x}_i, \hat{a}_i, \hat{y}_i)\}_{i=1}^N$. We assume that $N \geq 2$ and $\hat{\xi}_i = (\hat{x}_i, \hat{a}_i, \hat{y}_i)$ are distinct. We use $\mathcal{M}_+(\Xi \times \hat{\Xi}_N)$ to denote the set of all nonnegative measures on $\Xi \times \hat{\Xi}_N$.

Lemma A.1 (Slater condition - Probabilistic equal opportunity). Suppose that $\beta \neq 0$, $\hat{p}_{11}^N \in (0, 1)$ and $\hat{p}_{01}^N \in (0, 1)$. Define the function

$$f_\beta(X, A, Y) \triangleq \frac{1}{\hat{p}_{11}^N} h_\beta(X) \mathbb{1}_{(1,1)}(A, Y) - \frac{1}{\hat{p}_{01}^N} h_\beta(X) \mathbb{1}_{(0,1)}(A, Y),$$

and let f be a vector-valued function $f : \Xi \times \hat{\Xi}_N \rightarrow \mathbb{R}^{N+1}$

$$f(\xi, \xi') = \begin{pmatrix} \mathbb{1}_{\hat{\xi}_1}(\xi') \\ \vdots \\ \mathbb{1}_{\hat{\xi}_N}(\xi') \\ f_\beta(\xi) \end{pmatrix}.$$

Then we have

$$\begin{pmatrix} 1/N \\ \vdots \\ 1/N \\ 0 \end{pmatrix} \in \text{int} \{ \mathbb{E}_\pi[f(\xi, \xi')] : \pi \in \mathcal{M}_+(\Xi \times \hat{\Xi}_N) \}.$$

PROOF OF LEMMA A.1. It suffices to show that for any

$$q \in \left(\frac{1}{2N}, \frac{3}{2N} \right)^N \times \left(-\frac{1}{4}, \frac{1}{4} \right),$$

there exists a nonnegative measure $\pi \in \mathcal{M}_+(\Xi \times \hat{\Xi}_N)$ such that $q = \mathbb{E}_\pi[f(\xi, \xi')]$. We will verify this claim by constructing π explicitly. To this end, define the following locations

$$x_{ay} \in \mathcal{X} \quad \forall (a, y) \in \mathcal{A} \times \mathcal{Y},$$

and set $\pi \in \mathcal{M}_+(\Xi \times \hat{\Xi}_N)$ explicitly as

$$\pi(\xi = (x_{\hat{a}_i \hat{y}_i}, \hat{a}_i, \hat{y}_i), \xi' = (\hat{x}_i, \hat{a}_i, \hat{y}_i)) = q_i,$$

and π is 0 everywhere else. By construction, one can verify that $\mathbb{E}_\pi[\mathbb{1}_{\hat{\xi}_i}(\xi')] = q_i$ for all $i \in [N]$. If we define the following index sets $\mathcal{I}_{ay} = \{i \in [N] : \hat{a}_i = a, \hat{y}_i = y\}$, then

$$\mathbb{E}_\pi[f_\beta(\xi)] = (\hat{p}_{11}^N)^{-1} h_\beta(x_{11}) \sum_{i \in \mathcal{I}_{11}} q_i - (\hat{p}_{01}^N)^{-1} h_\beta(x_{01}) \sum_{i \in \mathcal{I}_{01}} q_i.$$

It now remains to find the locations of x_{11} and x_{01} to balance the above equation. We have the following two cases.

(1) Suppose that $q_{N+1} \geq 0$. In this case, choose $x_{01} \in \mathcal{X}$ such that $h_\beta(x_{01}) = \frac{1}{\delta}$. The condition $\mathbb{E}_\pi[f_\beta(\xi)] = q_{N+1}$ requires that

$$h_\beta(x_{11}) = \frac{q_{N+1} + \frac{1}{\delta} (\hat{p}_{01}^N)^{-1} \sum_{i \in \mathcal{I}_{01}} q_i}{(\hat{p}_{11}^N)^{-1} \sum_{i \in \mathcal{I}_{11}} q_i}.$$

Because $q_{N+1} \geq 0$ and q_i are strictly positive, the term on the right hand side is strictly positive. Moreover, we have

$$(\hat{p}_{01}^N)^{-1} \sum_{i \in \mathcal{I}_{01}} q_i < \frac{3}{2} \quad \text{and} \quad (\hat{p}_{11}^N)^{-1} \sum_{i \in \mathcal{I}_{11}} q_i > \frac{1}{2}$$

for any feasible value of q_i , which implies that

$$0 < \frac{q_{N+1} + \frac{1}{\delta} (\hat{p}_{01}^N)^{-1} \sum_{i \in \mathcal{I}_{01}} q_i}{(\hat{p}_{11}^N)^{-1} \sum_{i \in \mathcal{I}_{11}} q_i} < \frac{\frac{1}{4} + \frac{1}{4}}{\frac{1}{2}} = 1.$$

This implies the existence of $x_{11} \in \mathcal{X}$ so that $\mathbb{E}_\pi[f_\beta(\xi)] = q_{N+1}$.

(2) Suppose that $q_{N+1} < 0$. In this case, we can choose $x_{11} \in \mathcal{X}$ such that $h_\beta(x_{11}) = \frac{1}{\delta}$. A similar argument as in the previous case implies the existence of $x_{01} \in \mathcal{X}$ such that $\mathbb{E}_\pi[f_\beta(\xi)] = q_{N+1}$.

Combining the two cases leads to the postulated results. \square

We are now ready to prove Proposition 4.2.

PROOF OF PROPOSITION 4.2. For the purpose of this proof, we define the function $\lambda : \mathcal{A} \times \mathcal{Y} \rightarrow \mathbb{R}$ as

$$\lambda(a, y) = \frac{\mathbb{1}_{(1,1)}(a, y)}{\hat{p}_{11}^N} - \frac{\mathbb{1}_{(0,1)}(a, y)}{\hat{p}_{01}^N}. \quad (10)$$

By definition of the squared distance function \mathcal{R}^{OPP} , we have

$$\begin{aligned} & \mathcal{R}^{\text{OPP}}(\hat{\mathbb{P}}^N, \hat{p}^N) \\ &= \begin{cases} \inf_{\mathbb{Q} \in \mathcal{P}} & \mathbb{W}(\hat{\mathbb{P}}^N, \mathbb{Q})^2 \\ \text{s.t.} & (\hat{p}_{11}^N)^{-1} \mathbb{E}_{\mathbb{Q}}[h_\beta(X) \mathbb{1}_{(1,1)}(A, Y)] \\ &= (\hat{p}_{01}^N)^{-1} \mathbb{E}_{\mathbb{Q}}[h_\beta(X) \mathbb{1}_{(0,1)}(A, Y)] \\ & \mathbb{Q}(A = a, Y = y) = \hat{p}_{ay}^N \quad \forall a \in \mathcal{A}, y \in \mathcal{Y} \end{cases} \\ &= \begin{cases} \inf_{\pi} & \mathbb{E}_\pi[c((X', A', Y'), (X, A, Y))^2] \\ \text{s.t.} & \pi \in \mathcal{P}((\mathcal{X} \times \mathcal{A} \times \mathcal{Y}) \times (\mathcal{X} \times \mathcal{A} \times \mathcal{Y})) \\ & \mathbb{E}_\pi[f_\beta(X, A, Y)] = 0 \\ & \pi(A = a, Y = y) = \hat{p}_{ay}^N \quad \forall a \in \mathcal{A}, y \in \mathcal{Y} \\ & \mathbb{E}_\pi[\mathbb{1}_{(\hat{x}_i, \hat{a}_i, \hat{y}_i)}(X', A', Y')] = 1/N \quad \forall i \in [N], \end{cases} \end{aligned}$$

where the function f_β is defined as

$$\begin{aligned} f_\beta(x, a, y) &\triangleq (\hat{p}_{11}^N)^{-1} h_\beta(x) \mathbb{1}_{(1,1)}(a, y) - (\hat{p}_{01}^N)^{-1} h_\beta(x) \mathbb{1}_{(0,1)}(a, y) \\ &= h_\beta(x) \lambda(a, y), \end{aligned} \quad (11)$$

and $\mathcal{P}(\mathcal{S})$ denotes the set of all joint probability measures supported on \mathcal{S} . Because of the infinity individual cost on \mathcal{A} and \mathcal{Y} by the definition of cost in (2), any joint measure π with finite objective value should satisfies $\pi(A = a, Y = y) = \hat{\mathbb{P}}^N(A' = a, Y' = y) = \hat{p}_{ay}^N$ for any $a \in \mathcal{A}$ and $y \in \mathcal{Y}$. Thus, the set of constraints $\pi(A = a, Y = y) = \hat{p}_{ay}^N$ can be eliminated without alternating the optimization problem. We thus have

$$\begin{aligned} & \mathcal{R}^{\text{OPP}}(\hat{\mathbb{P}}^N, \hat{p}^N) \\ &= \begin{cases} \inf_{\pi} & \mathbb{E}_\pi[c((X', A', Y'), (X, A, Y))^2] \\ \text{s.t.} & \pi \in \mathcal{P}((\mathcal{X} \times \mathcal{A} \times \mathcal{Y}) \times (\mathcal{X} \times \mathcal{A} \times \mathcal{Y})) \\ & \mathbb{E}_\pi[f_\beta(X, A, Y)] = 0 \\ & \mathbb{E}_\pi[\mathbb{1}_{(\hat{x}_i, \hat{a}_i, \hat{y}_i)}(X', A', Y')] = 1/N \quad \forall i \in [N]. \end{cases} \end{aligned}$$

To shorten the notations, we use $\Xi = \mathcal{X} \times \mathcal{A} \times \mathcal{Y}$ and $\hat{\Xi}_N = \{(\hat{x}_i, \hat{a}_i, \hat{y}_i)\}$. Moreover, define the vector \bar{q} and the vector-valued Borel measurable function on $\Xi \times \hat{\Xi}_N$ as

$$\bar{q} = \begin{pmatrix} 0 \\ 1/N \\ \vdots \\ 1/N \end{pmatrix} \quad f(\xi, \xi') = \begin{pmatrix} f_\beta(\xi) \\ \mathbb{1}_{\hat{\xi}_i}(\xi') \\ \vdots \\ \mathbb{1}_{\hat{\xi}_N}(\xi') \end{pmatrix}.$$

By using the introduced notation, we can reformulate the above optimization problem as

$$\inf \{ \mathbb{E}_\pi[c(\xi, \xi')^2] : \pi \in \mathcal{M}_+(\Xi \times \hat{\Xi}_N), \mathbb{E}_\pi[f(\xi, \xi')] = \bar{q} \}$$

which is a problem of moments. By Lemma A.1, the above optimization problem satisfies the Slater condition, thus the strong duality result [60, Section 2.2] implies that

$$\begin{aligned} & \mathcal{R}^{\text{OPP}}(\hat{\mathbb{P}}^N, \hat{p}^N) \\ &= \begin{cases} \sup & \frac{1}{N} \sum_{i=1}^N b_i \\ \text{s.t.} & b \in \mathbb{R}^N, \gamma \in \mathbb{R} \\ & \sum_{i=1}^N b_i \mathbb{1}_{(\hat{x}_i, \hat{a}_i, \hat{y}_i)}(x', a', y') - \gamma f_\beta(x, a, y) \\ & \leq c((x', a', y'), (x, a, y))^2 \\ & \forall (x, a, y), (x', a', y') \in \mathcal{X} \times \mathcal{A} \times \mathcal{Y}. \end{cases} \end{aligned} \quad (12)$$

Note that the problem in (12) can be equivalently represented as

$$\begin{aligned} & \begin{cases} \sup & \frac{1}{N} \sum_{i=1}^N b_i \\ \text{s.t.} & b \in \mathbb{R}^N, \gamma \in \mathbb{R} \\ & b_i - \gamma f_\beta(x_i, a_i, y_i) \leq c((\hat{x}_i, \hat{a}_i, \hat{y}_i), (x_i, a_i, y_i))^2 \\ & \forall (x_i, a_i, y_i) \in \mathcal{X} \times \mathcal{A} \times \mathcal{Y}, \forall i \in [N] \end{cases} \\ &= \sup_{\gamma \in \mathbb{R}} \frac{1}{N} \sum_{i=1}^N \inf_{x_i \in \mathcal{X}} \{ \|x_i - \hat{x}_i\|^2 + \gamma f_\beta(x_i, \hat{a}_i, \hat{y}_i) \}. \end{aligned} \quad (13)$$

Because f_β has the form (11), we have the equivalent problem

$$\sup_{\gamma \in \mathbb{R}} \frac{1}{N} \sum_{i=1}^N \inf_{x_i \in \mathcal{X}} \{ \|x_i - \hat{x}_i\|^2 + \gamma \lambda(\hat{a}_i, \hat{y}_i) h_\beta(x_i) \}.$$

For any $i \in \mathcal{I}_0$, $\lambda(\hat{a}_i, \hat{y}_i) = 0$, and in this case we have the optimal solution of x_i satisfies $x_i^* = \hat{x}_i$. As a consequence, the summation collapses to a partial sum over \mathcal{I}_1 . This observation completes the proof. \square

PROOF OF THEOREM 4.4. Leveraging equation (13), we can express

$$\mathcal{R}^{\text{opp}}(\hat{\mathbb{P}}^N, \hat{p}^N) = \sup_{\gamma} \mathbb{E}_{\hat{\mathbb{P}}^N} \left[\inf_{\Delta} \gamma h_\beta(X + \Delta) \left(\frac{\mathbb{1}_{(1,1)}(A, Y)}{\hat{p}_{11}^N} - \frac{\mathbb{1}_{(0,1)}(A, Y)}{\hat{p}_{01}^N} \right) + \|\Delta\|^2 \right].$$

We define

$$H^N \triangleq \frac{1}{\sqrt{N}} \sum_{i=1}^N h_\beta(\hat{x}_i) \left(\frac{\mathbb{1}_{(1,1)}(\hat{a}_i, \hat{y}_i)}{\hat{p}_{11}^N} - \frac{\mathbb{1}_{(0,1)}(\hat{a}_i, \hat{y}_i)}{\hat{p}_{01}^N} \right),$$

and using this expression we can reformulate $\mathcal{R}^{\text{opp}}(\hat{\mathbb{P}}^N, \hat{p}^N)$ as

$$\sup_{\gamma} \left\{ \frac{1}{\sqrt{N}} \gamma H^N + \mathbb{E}_{\hat{\mathbb{P}}^N} \left[\inf_{\Delta} \gamma [h_\beta(X + \Delta) - h_\beta(X)] \times \left(\frac{\mathbb{1}_{(1,1)}(A, Y)}{\hat{p}_{11}^N} - \frac{\mathbb{1}_{(0,1)}(A, Y)}{\hat{p}_{01}^N} \right) + \|\Delta\|^2 \right] \right\}.$$

Because h_β is a sigmoid function, it is differentiable, and by the fundamental theorem of calculus, we have for any $x \in \mathcal{X}$,

$$h_\beta(x + \Delta) - h_\beta(x) = \int_0^1 \nabla h_\beta(x + t\Delta) \cdot \Delta dt,$$

where \cdot represents the inner product on \mathbb{R}^d . By applying variable transformations $\gamma \leftarrow \gamma\sqrt{N}$ and $\Delta \leftarrow \Delta\sqrt{N}$, we have

$$\begin{aligned} N \times \mathcal{R}^{\text{opp}}(\hat{\mathbb{P}}^N, \hat{p}^N) &= \sup_{\gamma} \left\{ \gamma H^N + \mathbb{E}_{\hat{\mathbb{P}}^N} \left[\inf_{\Delta} \gamma \int_0^1 \nabla h_\beta \left(X + t \frac{\Delta}{\sqrt{N}} \right) \cdot \Delta dt \right. \right. \\ &\quad \left. \left. \left(\frac{\mathbb{1}_{(1,1)}(A, Y)}{\hat{p}_{11}^N} - \frac{\mathbb{1}_{(0,1)}(A, Y)}{\hat{p}_{01}^N} \right) + \|\Delta\|^2 \right] \right\} \\ &= \sup_{\gamma} \left\{ \gamma H^N + \frac{1}{N} \sum_{i=1}^N \inf_{\Delta_i} \gamma \int_0^1 \nabla h_\beta \left(\hat{x}_i + t \frac{\Delta_i}{\sqrt{N}} \right) \cdot \Delta_i dt \times \right. \\ &\quad \left. \left(\frac{\mathbb{1}_{(1,1)}(\hat{a}_i, \hat{y}_i)}{\hat{p}_{11}^N} - \frac{\mathbb{1}_{(0,1)}(\hat{a}_i, \hat{y}_i)}{\hat{p}_{01}^N} \right) + \|\Delta_i\|^2 \right\}, \end{aligned}$$

where the second equality follows by the definition of the empirical distribution $\hat{\mathbb{P}}^N$. For any values of $\hat{p}_{01}^N > 0$ and $\hat{p}_{11}^N > 0$, we have for any $\gamma \neq 0$,

$$\begin{aligned} \mathbb{P} \left(\left\| \gamma \nabla h_\beta(X) \left(\frac{\mathbb{1}_{(1,1)}(A, Y)}{\hat{p}_{11}^N} - \frac{\mathbb{1}_{(0,1)}(A, Y)}{\hat{p}_{01}^N} \right) \right\|_* = 0 \right) &= \mathbb{P}((\hat{p}_{11}^N)^{-1} \mathbb{1}_{(1,1)}(A, Y) = (\hat{p}_{01}^N)^{-1} \mathbb{1}_{(0,1)}(A, Y)) \\ &= \mathbb{P}(Y = 0) < 1, \end{aligned}$$

which implies that

$$\mathbb{P} \left(\left\| \gamma \nabla h_\beta(X) \left(\frac{\mathbb{1}_{(1,1)}(A, Y)}{\hat{p}_{11}^N} - \frac{\mathbb{1}_{(0,1)}(A, Y)}{\hat{p}_{01}^N} \right) \right\|_* > 0 \right) > 0.$$

This coincides with Assumption A4 in [7]. Using the same argument as in the proof of [7, Theorem 3], we can show that the optimal solution for γ and Δ_i belong to a compact set with high probability. Moreover, we have

$$\begin{aligned} &\frac{\mathbb{1}_{(1,1)}(\hat{a}_i, \hat{y}_i)}{\hat{p}_{11}^N} - \frac{\mathbb{1}_{(0,1)}(\hat{a}_i, \hat{y}_i)}{\hat{p}_{01}^N} \\ &= \frac{\mathbb{1}_{(1,1)}(\hat{a}_i, \hat{y}_i)}{p_{11}} (1 - o_{\mathbb{P}}(1)) - \frac{\mathbb{1}_{(0,1)}(\hat{a}_i, \hat{y}_i)}{p_{01}} (1 - o_{\mathbb{P}}(1)), \end{aligned}$$

and thus

$$\begin{aligned} N \times \mathcal{R}^{\text{opp}}(\hat{\mathbb{P}}^N, \hat{p}^N) &= \sup_{\gamma} \left\{ \gamma H^N + \frac{1}{N} \sum_{i=1}^N \inf_{\Delta_i} \gamma \int_0^1 \nabla h_\beta \left(\hat{x}_i + t \frac{\Delta_i}{\sqrt{N}} \right) \cdot \Delta_i dt \times \right. \\ &\quad \left. \left(\frac{\mathbb{1}_{(1,1)}(\hat{a}_i, \hat{y}_i)}{p_{11}} - \frac{\mathbb{1}_{(0,1)}(\hat{a}_i, \hat{y}_i)}{p_{01}} \right) + \|\Delta_i\|^2 + o_{\mathbb{P}}(1) \right\}. \end{aligned}$$

In the next step, fix any tuple $(a, y) \in \mathcal{A} \times \mathcal{Y}$, and denote the following constant

$$M_1 = |p_{11}^{-1} \mathbb{1}_{(1,1)}(a, y) - p_{01}^{-1} \mathbb{1}_{(0,1)}(a, y)|.$$

We find

$$\begin{aligned} &\| [\nabla h_\beta(x + \Delta) - \nabla h_\beta(x)] (p_{11}^{-1} \mathbb{1}_{(1,1)}(a, y) - p_{01}^{-1} \mathbb{1}_{(0,1)}(a, y)) \|_* \\ &= |h_\beta(x + \Delta) - h_\beta(x) - h_\beta(x + \Delta)^2 + h_\beta(x)^2| \|\beta\|_* M_1 \\ &\leq (|h_\beta(x + \Delta) - h_\beta(x)| + |h_\beta(x + \Delta)^2 - h_\beta(x)^2|) \|\beta\|_* M_1. \end{aligned}$$

Because the sigmoid function is slope-restricted in the interval $[0, 1]$ [22, Proposition 2], we have

$$0 \leq \frac{h_\beta(x + \Delta) - h_\beta(x)}{\beta^\top \Delta} \leq 1,$$

which implies that

$$|h_\beta(x + \Delta) - h_\beta(x)| \leq |\beta^\top \Delta| \leq \|\beta\|_* \|\Delta\|,$$

where the second inequality follows from Hölder inequality. Using a similar argument, we have

$$\begin{aligned} |h_\beta(x + \Delta)^2 - h_\beta(x)^2| &\leq (h_\beta(x + \Delta) + h_\beta(x)) |h_\beta(x + \Delta) - h_\beta(x)| \\ &\leq 2 \|\beta\|_* \|\Delta\|. \end{aligned}$$

Combining these inequalities, we conclude that

$$\begin{aligned} &\| [\nabla h_\beta(x + \Delta) - \nabla h_\beta(x)] (p_{11}^{-1} \mathbb{1}_{(1,1)}(a, y) - p_{01}^{-1} \mathbb{1}_{(0,1)}(a, y)) \|_2 \\ &\leq 3 \|\beta\|_*^2 M_1 \|\Delta\|, \end{aligned}$$

and thus Assumption 6' in [7] is satisfied. If $H^N \xrightarrow{d} \tilde{Z}$ for some random variable \tilde{Z} , then [7, Lemma 4] asserts that

$$\begin{aligned} & N \times \mathcal{R}^{\text{OPP}}(\hat{\mathbb{P}}^N, \hat{p}^N) \\ & \xrightarrow{d} \sup_{\gamma \in \mathbb{R}} \left\{ \gamma \tilde{Z} - \frac{\gamma^2}{4} \mathbb{E}_{\mathbb{P}} \left[\left\| \nabla h_{\beta}(X) \left(\frac{\mathbb{1}_{(1,1)}(A, Y)}{p_{11}} - \frac{\mathbb{1}_{(0,1)}(A, Y)}{p_{01}} \right) \right\|_*^2 \right] \right\} \\ & = \left(\mathbb{E}_{\mathbb{P}} \left[\left\| \nabla h_{\beta}(X) \left(\frac{\mathbb{1}_{(1,1)}(A, Y)}{p_{11}} - \frac{\mathbb{1}_{(0,1)}(A, Y)}{p_{01}} \right) \right\|_*^2 \right] \right)^{-1} \tilde{Z}^2, \end{aligned}$$

where the equality sign follows from the fact that for any realization of \tilde{Z} , the optimal solution of γ is

$$\gamma^*(\tilde{Z}) = \frac{2\tilde{Z}}{\mathbb{E}_{\mathbb{P}} \left[\left\| \nabla h_{\beta}(X) \left(\frac{\mathbb{1}_{(1,1)}(A, Y)}{p_{11}} - \frac{\mathbb{1}_{(0,1)}(A, Y)}{p_{01}} \right) \right\|_*^2 \right]}.$$

We now study the limit distribution \tilde{Z} . In the next step, we study the limit of H^N .

$$\begin{aligned} H^N &= \frac{1}{\sqrt{N}} \sum_{i=1}^N h_{\beta}(\hat{x}_i) \left(\frac{\mathbb{1}_{(1,1)}(\hat{a}_i, \hat{y}_i)}{\hat{p}_{11}^N} - \frac{\mathbb{1}_{(0,1)}(\hat{a}_i, \hat{y}_i)}{\hat{p}_{01}^N} \right) \\ &= \frac{1}{\hat{p}_{11}^N \hat{p}_{01}^N} \times \frac{1}{\sqrt{N}} \sum_{i=1}^N h_{\beta}(\hat{x}_i) \left(\hat{p}_{01}^N \mathbb{1}_{(1,1)}(\hat{a}_i, \hat{y}_i) - \hat{p}_{11}^N \mathbb{1}_{(0,1)}(\hat{a}_i, \hat{y}_i) \right) \\ &= \frac{1}{\hat{p}_{11}^N \hat{p}_{01}^N} \times \left(\frac{1}{\sqrt{N}} \sum_{i=1}^N h_{\beta}(\hat{x}_i) (p_{01} \mathbb{1}_{(1,1)}(\hat{a}_i, \hat{y}_i) - p_{11} \mathbb{1}_{(0,1)}(\hat{a}_i, \hat{y}_i)) \right. \\ & \quad \left. + \sqrt{N}(\hat{p}_{01}^N - p_{01}) \frac{1}{N} \sum_{i=1}^N \mathbb{1}_{(1,1)}(\hat{a}_i, \hat{y}_i) h_{\beta}(\hat{x}_i) \right. \\ & \quad \left. - \sqrt{N}(\hat{p}_{11}^N - p_{11}) \frac{1}{N} \sum_{i=1}^N \mathbb{1}_{(0,1)}(\hat{a}_i, \hat{y}_i) h_{\beta}(\hat{x}_i) \right) \end{aligned}$$

By Slutsky's theorem, we have

$$\begin{aligned} & \sqrt{N}(\hat{p}_{01}^N - p_{01}) \times \\ & \quad \frac{1}{N} \sum_{i=1}^N \left(\mathbb{1}_{(1,1)}(\hat{a}_i, \hat{y}_i) h_{\beta}(\hat{x}_i) - \mathbb{E}_{\mathbb{P}}[\mathbb{1}_{(1,1)}(A, Y) h_{\beta}(X)] \right) = o_{\mathbb{P}}(1), \\ & \sqrt{N}(\hat{p}_{11}^N - p_{11}) \times \\ & \quad \frac{1}{N} \sum_{i=1}^N \left(\mathbb{1}_{(0,1)}(\hat{a}_i, \hat{y}_i) h_{\beta}(\hat{x}_i) - \mathbb{E}_{\mathbb{P}}[\mathbb{1}_{(0,1)}(A, Y) h_{\beta}(X)] \right) = o_{\mathbb{P}}(1). \end{aligned}$$

Under the null hypothesis $\mathcal{H}_0^{\text{OPP}}$, we have

$$\begin{aligned} H^N &= \frac{1}{\hat{p}_{11}^N \hat{p}_{01}^N} \times \left[\frac{1}{\sqrt{N}} \sum_{i=1}^N h_{\beta}(\hat{x}_i) \left(p_{01} \mathbb{1}_{(1,1)}(\hat{a}_i, \hat{y}_i) - p_{11} \mathbb{1}_{(0,1)}(\hat{a}_i, \hat{y}_i) \right) \right. \\ & \quad \left. + \sqrt{N} \left(\frac{1}{N} \sum_{i=1}^N \mathbb{1}_{(0,1)}(\hat{a}_i, \hat{y}_i) - p_{01} \right) \mathbb{E}_{\mathbb{P}}[\mathbb{1}_{(1,1)}(A, Y) h_{\beta}(X)] \right. \\ & \quad \left. - \sqrt{N} \left(\frac{1}{N} \sum_{i=1}^N \mathbb{1}_{(1,1)}(\hat{a}_i, \hat{y}_i) - p_{11} \right) \mathbb{E}_{\mathbb{P}}[\mathbb{1}_{(0,1)}(A, Y) h_{\beta}(X)] \right] + o_{\mathbb{P}}(1) \\ &= \frac{1}{\hat{p}_{11}^N \hat{p}_{01}^N} \times \left[\frac{1}{\sqrt{N}} \sum_{i=1}^N h_{\beta}(\hat{x}_i) \left(p_{01} \mathbb{1}_{(1,1)}(\hat{a}_i, \hat{y}_i) - p_{11} \mathbb{1}_{(0,1)}(\hat{a}_i, \hat{y}_i) \right) \right. \\ & \quad \left. + \frac{1}{\sqrt{N}} \sum_{i=1}^N \left(\mathbb{1}_{(0,1)}(\hat{a}_i, \hat{y}_i) - p_{01} \right) \mathbb{E}_{\mathbb{P}}[\mathbb{1}_{(1,1)}(A, Y) h_{\beta}(X)] \right. \\ & \quad \left. - \frac{1}{\sqrt{N}} \sum_{i=1}^N \left(\mathbb{1}_{(1,1)}(\hat{a}_i, \hat{y}_i) - p_{11} \right) \mathbb{E}_{\mathbb{P}}[\mathbb{1}_{(0,1)}(A, Y) h_{\beta}(X)] \right] + o_{\mathbb{P}}(1) \\ & \xrightarrow{d} \tilde{Z}, \end{aligned}$$

where $\tilde{Z} \sim \frac{1}{p_{11}p_{01}} \mathcal{N}(0, \sigma^2)$, $\sigma^2 = \text{Cov}(Z)$, where Z is defined as in the theorem statement. Defining θ completes the proof. \square

A.3 Proofs of Section 5

The proof of Proposition 5.2 necessitates the following preparatory lemma. We use the same notations with Lemma A.1.

Lemma A.2 (Slater condition - Probabilistic equalized odds). Suppose that $\beta \neq 0$ and $\hat{p}_{ay}^N \in (0, 1)$ for all $(a, y) \in \mathcal{A} \times \mathcal{Y}$. Define the functions

$$\begin{aligned} f_{\beta}(X, A, Y) &\triangleq \frac{1}{\hat{p}_{11}^N} h_{\beta}(X) \mathbb{1}_{(1,1)}(A, Y) - \frac{1}{\hat{p}_{01}^N} h_{\beta}(X) \mathbb{1}_{(0,1)}(A, Y), \\ g_{\beta}(X, A, Y) &\triangleq \frac{1}{\hat{p}_{10}^N} h_{\beta}(X) \mathbb{1}_{(1,0)}(A, Y) - \frac{1}{\hat{p}_{00}^N} h_{\beta}(X) \mathbb{1}_{(0,0)}(A, Y), \end{aligned}$$

and let f be a vector-valued function $f : \Xi \times \hat{\Xi}_N \rightarrow \mathbb{R}^{N+2}$

$$f(\xi, \xi') = \begin{pmatrix} \mathbb{1}_{\hat{\xi}_i}(\xi') \\ \vdots \\ \mathbb{1}_{\hat{\xi}_N}(\xi') \\ f_{\beta}(\xi) \\ g_{\beta}(\xi) \end{pmatrix}$$

Then we have

$$\begin{pmatrix} 1/N \\ \vdots \\ 1/N \\ 0 \\ 0 \end{pmatrix} \in \text{int} \{ \mathbb{E}_{\pi} [f(\xi, \xi')] : \pi \in \mathcal{M}_+(\Xi \times \hat{\Xi}_N) \}.$$

PROOF OF LEMMA A.2. It suffices to show that for any

$$q \in \left(\frac{1}{2N}, \frac{3}{2N} \right)^N \times \left(-\frac{1}{4}, \frac{1}{4} \right)^2,$$

there exists a nonnegative measure $\pi \in \mathcal{M}_+(\Xi \times \hat{\Xi}_N)$ such that $q = \mathbb{E}_\pi[f(\xi, \xi')]$. The proof follows a similar argument as that of Lemma A.1 by noticing that

$$\mathbb{E}_\pi[g_\beta(\xi)] = (\hat{p}_{10}^N)^{-1} h_\beta(x_{10}) \sum_{i \in \mathcal{I}_{10}} q_i - (\hat{p}_{00}^N)^{-1} h_\beta(x_{00}) \sum_{i \in \mathcal{I}_{00}} q_i,$$

and the specification of x_{10} and x_{00} can be achieved using similar steps. \square

PROOF OF PROPOSITION 5.2. To ease the exposition, we let the function $\Lambda : \mathcal{A} \times \mathcal{Y} \rightarrow \mathbb{R}^2$ be defined as

$$\Lambda(a, y) = \begin{pmatrix} (\hat{p}_{11}^N)^{-1} \mathbb{1}_{(1,1)}(a, y) - (\hat{p}_{01}^N)^{-1} \mathbb{1}_{(0,1)}(a, y) \\ (\hat{p}_{10}^N)^{-1} \mathbb{1}_{(1,0)}(a, y) - (\hat{p}_{00}^N)^{-1} \mathbb{1}_{(0,0)}(a, y) \end{pmatrix}.$$

Moreover, we define f_β as in (11), and additionally define g_β as

$$g_\beta(x, a, y) = (\hat{p}_{10}^N)^{-1} h_\beta(x) \mathbb{1}_{(1,0)}(a, y) - (\hat{p}_{00}^N)^{-1} h_\beta(x) \mathbb{1}_{(0,0)}(a, y).$$

From the definition of $\mathcal{R}^{\text{odd}}(\hat{\mathbb{P}}^N, \hat{p}^N)$, we have

$$\begin{aligned} & \mathcal{R}^{\text{odd}}(\hat{\mathbb{P}}^N, \hat{p}^N) \\ &= \begin{cases} \inf_{\mathbb{Q} \in \mathcal{P}} \mathbb{W}(\hat{\mathbb{P}}^N, \mathbb{Q})^2 \\ \text{s.t. } (\hat{p}_{11}^N)^{-1} \mathbb{E}_{\mathbb{Q}}[h_\beta(X) \mathbb{1}_{(1,1)}(A, Y)] \\ \quad = (\hat{p}_{01}^N)^{-1} \mathbb{E}_{\mathbb{Q}}[h_\beta(X) \mathbb{1}_{(0,1)}(A, Y)] \\ (\hat{p}_{10}^N)^{-1} \mathbb{E}_{\mathbb{Q}}[h_\beta(X) \mathbb{1}_{(1,0)}(A, Y)] \\ \quad = (\hat{p}_{00}^N)^{-1} \mathbb{E}_{\mathbb{Q}}[h_\beta(X) \mathbb{1}_{(0,0)}(A, Y)] \\ \mathbb{Q}(A = a, Y = y) = \hat{p}_{ay}^N \quad \forall a \in \mathcal{A}, y \in \mathcal{Y} \end{cases} \\ &= \begin{cases} \inf_{\pi} \mathbb{E}_\pi[c((X', A', Y'), (X, A, Y))^2] \\ \text{s.t. } \pi \in \mathcal{P}((\mathcal{X} \times \mathcal{A} \times \mathcal{Y}) \times (\mathcal{X} \times \mathcal{A} \times \mathcal{Y})) \\ \mathbb{E}_\pi[f_\beta(X, A, Y)] = 0 \\ \mathbb{E}_\pi[g_\beta(X, A, Y)] = 0 \\ \pi(A = a, Y = y) = \hat{p}_{ay}^N \quad \forall a \in \mathcal{A}, y \in \mathcal{Y} \\ \mathbb{E}_\pi[\mathbb{1}_{(\hat{x}_i, \hat{a}_i, \hat{y}_i)}(X', A', Y')] = 1/N \quad \forall i \in [N]. \end{cases} \end{aligned}$$

To shorten the notations, we use $\Xi = \mathcal{X} \times \mathcal{A} \times \mathcal{Y}$ and $\hat{\Xi}_N = \{(\hat{x}_i, \hat{a}_i, \hat{y}_i)\}$. Moreover, define the vector \bar{q} and the vector-valued Borel measurable function on $\Xi \times \hat{\Xi}_N$ as

$$\bar{q} = \begin{pmatrix} 0 \\ 0 \\ 1/N \\ \vdots \\ 1/N \end{pmatrix} \quad f(\xi, \xi') = \begin{pmatrix} f_\beta(\xi) \\ g_\beta(\xi) \\ \mathbb{1}_{\hat{\xi}_i}(\xi') \\ \vdots \\ \mathbb{1}_{\hat{\xi}_N}(\xi') \end{pmatrix}.$$

By using the introduced notation, we can reformulate the above optimization problem as

$$\inf \{ \mathbb{E}_\pi[c(\xi, \xi')^2] : \pi \in \mathcal{M}_+(\Xi \times \hat{\Xi}_N), \mathbb{E}_\pi[f(\xi, \xi')] = \bar{q} \}$$

which is a problem of moments. By Lemma A.2, the above optimization problem satisfies the Slater condition, thus the strong duality

result [60, Section 2.2] implies that

$$\begin{aligned} & \mathcal{R}^{\text{odd}}(\hat{\mathbb{P}}^N, \hat{p}^N) \\ &= \begin{cases} \sup \frac{1}{N} \sum_{i=1}^N b_i \\ \text{s.t. } b \in \mathbb{R}^N, \gamma \in \mathbb{R}, \zeta \in \mathbb{R} \\ \sum_{i=1}^N b_i \mathbb{1}_{(\hat{x}_i, \hat{a}_i, \hat{y}_i)}(x', a', y') - \gamma f_\beta(x, a, y) - \zeta g_\beta(x, a, y) \\ \quad \leq c((x', a', y'), (x, a, y))^2 \\ \quad \forall (x, a, y), (x', a', y') \in \mathcal{X} \times \mathcal{A} \times \mathcal{Y} \end{cases} \\ &= \begin{cases} \sup \frac{1}{N} \sum_{i=1}^N b_i \\ \text{s.t. } b \in \mathbb{R}^N, \gamma \in \mathbb{R}, \zeta \in \mathbb{R} \\ b_i - \gamma f_\beta(x_i, a_i, y_i) - \zeta g_\beta(x_i, a_i, y_i) \\ \quad \leq c((\hat{x}_i, \hat{a}_i, \hat{y}_i), (x_i, a_i, y_i))^2 \\ \quad \forall (x_i, a_i, y_i) \in \mathcal{X} \times \mathcal{A} \times \mathcal{Y}, \forall i \in [N] \end{cases} \\ &= \sup_{\gamma, \zeta} \frac{1}{N} \sum_{i=1}^N \inf_{x_i \in \mathcal{X}} \{ \|x_i - \hat{x}_i\|^2 + \gamma f_\beta(x_i, \hat{a}_i, \hat{y}_i) + \zeta g_\beta(x_i, \hat{a}_i, \hat{y}_i) \}, \end{aligned}$$

By definition of f_β, g_β and the parameters λ_i , we have

$$\gamma f_\beta(x_i, \hat{a}_i, \hat{y}_i) + \zeta g_\beta(x_i, \hat{a}_i, \hat{y}_i) = (\gamma \lambda_i \mathbb{1}_1(\hat{y}_i) + \zeta \lambda_i \mathbb{1}_0(\hat{y}_i)) h_\beta(x_i).$$

The proof is complete. \square

PROOF OF LEMMA 5.3. Because $[N] = \mathcal{I}_0 \cup \mathcal{I}_1$, we can write

$$\begin{aligned} & \mathcal{R}^{\text{odd}}(\hat{\mathbb{P}}^N, \hat{p}^N) \\ &= \sup_{\gamma \in \mathbb{R}} \frac{1}{N} \sum_{i \in \mathcal{I}_1} \inf_{x_i \in \mathcal{X}} \{ \|x_i - \hat{x}_i\|^2 + \gamma \lambda_i h_\beta(x_i) \} \\ & \quad + \sup_{\zeta \in \mathbb{R}} \frac{1}{N} \sum_{i \in \mathcal{I}_0} \inf_{x_i \in \mathcal{X}} \{ \|x_i - \hat{x}_i\|^2 + \zeta \lambda_i h_\beta(x_i) \}. \end{aligned}$$

Note that the first supremum coincides with $\mathcal{R}^{\text{opp}}(\hat{\mathbb{P}}^N, \hat{p}^N)$, and the second supremum is U_N . Under the Euclidean norm assumption, we can use Lemma B.1 to reformulate the inner infimum problems for U_N , which leads to (9). \square

PROOF OF THEOREM 5.4. By applying a similar duality argument as in the proof of Theorem 4.4, we can reformulate $\mathcal{R}^{\text{odd}}(\hat{\mathbb{P}}^N, \hat{p}^N)$ as

$$\begin{aligned} & \mathcal{R}^{\text{odd}}(\hat{\mathbb{P}}^N, \hat{p}^N) \\ &= \sup_{\gamma, \zeta} \mathbb{E}_{\hat{\mathbb{P}}^N} \left[\inf_{\Delta} \left\{ \gamma h_\beta(X + \Delta) \left(\frac{\mathbb{1}_{(1,1)}(A, Y)}{\hat{p}_{11}^N} - \frac{\mathbb{1}_{(0,1)}(A, Y)}{\hat{p}_{01}^N} \right) \right. \right. \\ & \quad \left. \left. + \zeta h_\beta(X + \Delta) \left(\frac{\mathbb{1}_{(1,0)}(A, Y)}{\hat{p}_{10}^N} - \frac{\mathbb{1}_{(0,0)}(A, Y)}{\hat{p}_{00}^N} \right) + \|\Delta\|^2 \right\} \right] \\ &= \sup_{\gamma, \zeta} \left\{ \frac{1}{\sqrt{N}} (\zeta H_0^N + \gamma H_1^N) + \right. \\ & \quad \left. \mathbb{E}_{\hat{\mathbb{P}}^N} \left[\inf_{\Delta} \left(\gamma [h_\beta(X + \Delta) - h_\beta(X)] \left(\frac{\mathbb{1}_{(1,1)}(A, Y)}{\hat{p}_{11}^N} - \frac{\mathbb{1}_{(0,1)}(A, Y)}{\hat{p}_{01}^N} \right) \right. \right. \right. \\ & \quad \left. \left. + \zeta [h_\beta(X + \Delta) - h_\beta(X)] \left(\frac{\mathbb{1}_{(1,0)}(A, Y)}{\hat{p}_{10}^N} - \frac{\mathbb{1}_{(0,0)}(A, Y)}{\hat{p}_{00}^N} \right) \right. \right. \\ & \quad \left. \left. + \|\Delta\|^2 \right) \right] \right\} \end{aligned}$$

with the random variables H_0^N and H_1^N being defined as

$$H_0^N \triangleq \frac{1}{\sqrt{N}} \sum_{i=1}^N h_\beta(\hat{x}_i) \left(\frac{\mathbb{1}_{(1,0)}(\hat{a}_i, \hat{y}_i)}{\hat{p}_{10}^N} - \frac{\mathbb{1}_{(0,0)}(\hat{a}_i, \hat{y}_i)}{\hat{p}_{00}^N} \right),$$

$$H_1^N \triangleq \frac{1}{\sqrt{N}} \sum_{i=1}^N h_\beta(\hat{x}_i) \left(\frac{\mathbb{1}_{(1,1)}(\hat{a}_i, \hat{y}_i)}{\hat{p}_{11}^N} - \frac{\mathbb{1}_{(0,1)}(\hat{a}_i, \hat{y}_i)}{\hat{p}_{01}^N} \right).$$

Notice that the condition

$$\mathbb{P} \left(\left\| \begin{pmatrix} \gamma_1 \\ \gamma_0 \end{pmatrix}^\top \Lambda(A, Y) \nabla h_\beta(X) \right\|_* > 0 \right) > 0$$

is satisfied for any $(\gamma_0, \gamma_1) \neq 0$. Using the same argument as in the proof of [7, Theorem 3], we can show that the optimal solution for γ, ζ and Δ_i belong to a compact set with high probability. As $\hat{p}_{ay} - p_{ay} = o_{\mathbb{P}}(1)$ for any $(a, y) \in \mathcal{A} \times \mathcal{Y}$, we have

$$\begin{aligned} & N \times \mathcal{R}^{\text{odd}}(\hat{\mathbb{P}}^N, \hat{p}^N) \\ &= \sup_{\gamma, \zeta} \left\{ \gamma H_1^N + \zeta H_0^N + \frac{1}{N} \sum_{i=1}^N \inf_{\Delta_i} \gamma \int_0^1 \nabla h_\beta \left(\hat{x}_i + t \frac{\Delta_i}{\sqrt{N}} \right) \cdot \Delta_i dt \times \right. \\ & \quad \left. \begin{pmatrix} \gamma \\ \zeta \end{pmatrix}^\top \begin{pmatrix} p_{11}^{-1} \mathbb{1}_{(1,1)}(\hat{a}_i, \hat{y}_i) - p_{01}^{-1} \mathbb{1}_{(0,1)}(\hat{a}_i, \hat{y}_i) \\ p_{10}^{-1} \mathbb{1}_{(1,0)}(\hat{a}_i, \hat{y}_i) - p_{00}^{-1} \mathbb{1}_{(0,0)}(\hat{a}_i, \hat{y}_i) \end{pmatrix} + \|\Delta_i\|^2 + o_{\mathbb{P}}(1) \right\}. \end{aligned}$$

Using a similar argument, we can bound

$$\begin{aligned} & \|\nabla h_\beta(x + \Delta) - \nabla h_\beta(x)\| (p_{10}^{-1} \mathbb{1}_{(1,0)}(a, y) - p_{00}^{-1} \mathbb{1}_{(0,0)}(a, y)) \|_2 \\ & \leq 3 \|\beta\|_*^2 M_0 \|\Delta\| \end{aligned}$$

for some constant M_0 , and thus Assumption 6' in [7] is satisfied. If $H_0^N \xrightarrow{d} H_0$ and $H_1^N \xrightarrow{d} H_1$ for some random variables H_0 and H_1 , then [7, Lemma 4] asserts that

$$\begin{aligned} & N \times \mathcal{R}^{\text{odd}}(\hat{\mathbb{P}}^N, \hat{p}^N) \xrightarrow{d} \\ & \sup_{\gamma, \zeta} \{ \gamma H_1 + \zeta H_0 + \\ & \mathbb{E}_{\mathbb{P}} \left\{ \left\| \begin{pmatrix} \gamma \\ \zeta \end{pmatrix}^\top \begin{pmatrix} p_{11}^{-1} \mathbb{1}_{(1,1)}(A, Y) - p_{01}^{-1} \mathbb{1}_{(0,1)}(A, Y) \\ p_{10}^{-1} \mathbb{1}_{(1,0)}(A, Y) - p_{00}^{-1} \mathbb{1}_{(0,0)}(A, Y) \end{pmatrix} \nabla h_\beta(X) \right\|_*^2 \right\} \}. \end{aligned}$$

Using the same limiting argument as in the proof of Theorem 4.4, we have the characterization of H_1 and H_0 as in the statement of the theorem. \square

B APPENDIX - AUXILIARY RESULT

The following lemma is used repeatedly to prove Lemmas 4.3 and 5.3.

Lemma B.1. For any $\omega \in \mathbb{R}$, $\hat{x} \in \mathbb{R}^P$ and $\beta \in \mathbb{R}^P$, we have

$$\begin{aligned} & \inf_{x \in \mathbb{R}^P} \|x - \hat{x}\|_2^2 + \frac{\omega}{1 + \exp(-\beta^\top x)} \\ &= \min_{k \in [0, \frac{1}{8}]} \omega^2 \|\beta\|_2^2 k^2 + \frac{\omega}{1 + \exp(-\beta^\top \hat{x} + k\omega \|\beta\|_2^2)}. \end{aligned} \quad (14)$$

PROOF OF LEMMA B.1. Any $x \in \mathbb{R}^P$ can be written using the orthogonal decomposition as $x = \hat{x} - k\omega\beta - k'\beta^\perp$ for some $k \in \mathbb{R}$,

$k' \in \mathbb{R}$ and β^\perp perpendicular to β , that is, $\beta^\top(\beta^\perp) = 0$. Optimizing over x is equivalent to jointly optimizing over k, k' and β^\perp as

$$\begin{aligned} & \inf \|k\omega\beta + k'\beta^\perp\|_2^2 + \frac{\omega}{1 + \exp(-\beta^\top \hat{x} + k\omega \|\beta\|_2^2)} \\ & \text{s.t. } k \in \mathbb{R}, k' \in \mathbb{R}, \beta^\perp \in \mathbb{R}^P, \beta^\top(\beta^\perp) = 0. \end{aligned}$$

After extending the norm, and by noticing that the optimal solution in k' and β^\perp should satisfy $k'\beta^\perp = 0$, the above optimization problem is equivalent to

$$\begin{aligned} & \inf k^2 \omega^2 \|\beta\|_2^2 + \frac{\omega}{1 + \exp(-\beta^\top \hat{x} + k\omega \|\beta\|_2^2)} \\ & \text{s.t. } k \in \mathbb{R}. \end{aligned}$$

Let $L(k)$ be the objective function of the above optimization problem, we have

$$\begin{aligned} \nabla_k L(k) &= 2\omega^2 \|\beta\|_2^2 k - \frac{\omega^2 \|\beta\|_2^2 \exp(-\beta^\top \hat{x} + k\omega \|\beta\|_2^2)}{(1 + \exp(-\beta^\top \hat{x} + k\omega \|\beta\|_2^2))^2} \\ &= \omega^2 \|\beta\|_2^2 (2k - \sigma(k)(1 - \sigma(k))), \end{aligned}$$

where for the purpose of this proof, we define $\sigma(k)$ as

$$\sigma(k) \triangleq \frac{1}{1 + \exp(-\beta^\top \hat{x} + k\omega \|\beta\|_2^2)} \in (0, 1).$$

Notice that $\sigma(k)(1 - \sigma(k)) \in (0, \frac{1}{4})$ for any value of $k \in \mathbb{R}$. Because $\nabla_k L(k)$ is continuous in k , $\nabla_k L(k) \leq 0$ for any $k \leq 0$, and $\nabla_k L(k) \geq 0$ for any $k \geq \frac{1}{8}$, one can conclude that there exists an optimal solution k^* that lies in the compact range $[0, \frac{1}{8}]$. This completes the proof. \square

Let $L(k)$ be the objective function of the optimization problem (14). Figure 5 visualizes several instances of $L(k)$ for different values of inputs β, \hat{x} and ω . Note that $L(k)$ is non-convex in k , and the optimizer of $L(k)$ is not necessarily unique as indicated in Figure 5d.

C APPENDIX - NUMERICAL RESULTS

We use the synthetic experiment from [71] to generate unfairness landscapes provided in Figure 1. We set the true distributions of the class labels $\mathbb{P}(Y = 0) = \mathbb{P}(Y = 1) = 1/2$, and conditioning on Y , the feature X has


$$X|Y = 1 \sim \mathcal{N}([2; 2], [5, 1; 1, 5]),$$

$$X|Y = 0 \sim \mathcal{N}([-2; -2], [10, 1; 1, 3]).$$


Then, we draw sensitive attribute of each sample x from a Bernoulli distribution, that is

$$\mathbb{P}(A = 1|X = x') = pdf(x'|Y = 1) / (pdf(x'|Y = 1) + pdf(x'|Y = 0)),$$


where $x' = [\cos(\pi/4), \sin(\pi/4); \sin(\pi/4), \cos(\pi/4)]x$ is a rotated version of the feature vector x and $pdf(\cdot|Y = y)$ is the Gaussian probability density function of X given $Y = y$.




(a) $\beta = (0, 1)^T$, $\hat{x} = (-2, 10)^T$, $\omega = 17.6$



(b) $\beta = (-5, 5)^T$, $\hat{x} = (3, 5)^T$, $\omega = 4$



(c) $\beta = (-6, 5)^T$, $\hat{x} = (3, 5)^T$, $\omega = 4$



(d) $\beta = (-4.7, 5)^T$, $\hat{x} = (3, 5)^T$, $\omega = 4$

Figure 5: Plots of $L(k)$ with respect to k for different values of β , \hat{x} and ω .