

On the (In)Compatibility between Group Fairness and Individual Fairness

Shizhou Xu

*Department of Mathematics
University of California Davis
Davis, CA 95616-5270, USA*

SHZXU@UCDAVIS.EDU

Thomas Strohmer

*Department of Mathematics
Center of Data Science and Artificial Intelligence Research
University of California Davis
Davis, CA 95616-5270, USA*

STROHMER@MATH.UCDAVIS.EDU

Abstract

We study the compatibility between the optimal statistical parity solutions and individual fairness. While individual fairness seeks to treat similar individuals similarly, optimal statistical parity aims to provide similar treatment to individuals who share relative similarity within their respective sensitive groups. The two fairness perspectives, while both desirable from a fairness perspective, often come into conflict in applications. Our goal in this work is to analyze the existence of this conflict and its potential solution. In particular, we establish sufficient (sharp) conditions for the compatibility between the optimal (post-processing) statistical parity L^2 learning and the (K -Lipschitz or (ϵ, δ)) individual fairness requirements. Furthermore, when there exists a conflict between the two, we first relax the former to the Pareto frontier (or equivalently the optimal trade-off) between L^2 error and statistical disparity, and then analyze the compatibility between the frontier and the individual fairness requirements. Our analysis identifies regions along the Pareto frontier that satisfy individual fairness requirements. (Lastly, we provide individual fairness guarantees for the composition of a trained model and the optimal post-processing step so that one can determine the compatibility of the post-processed model.) This provides practitioners with a valuable approach to attain Pareto optimality for statistical parity while adhering to the constraints of individual fairness.

1. Introduction

Fairness in machine learning has gained increasing attention due to the increasing trend of decision-making and information sharing based on machine learning or artificial intelligence (AI) assistance in our daily lives, especially due to the current development of large language models (LLMs) and generative AI. The ethics of such decision-making or information-sharing process becomes a key to not only the development of machine learning but also the growth of our society in a healthy form.

There are two fairness concepts which look at fairness from different perspectives:

- I. **Group fairness** aims to enforce the (conditional) learning outcome to be equal in distribution or statistics among sensitive groups.

II. **Individual fairness** aims to guarantee that individuals who share similar qualification data would receive similar learning outcome.

Unfortunately, although both concepts are desirable in terms of fairness, the two can potentially conflict with each other. To see this, consider a learning outcome that does not satisfy the statistical parity (Definition 1.1), then it becomes necessary to move the individuals from different sensitive groups in different directions on the learning outcome space, such as $\{0, 1\}$ in classification and \mathbb{R} in 1-dimensional regression. However, such an enforcement of statistical parity can easily lead to a significant violation of individual fairness. Similarly, individual fairness tends to keep individuals sharing similar qualifications close on the learning outcome space. Therefore, given sensitive information being excluded from qualification, such closeness preservation is likely to extend the statistical disparity among different sensitive groups on the independent variable (excluding sensitive information) space to the learning outcome and hence violates group fairness definitions such as statistical parity.

This, naturally, gives rise to the following question, which remains open on the current frontiers of machine learning fairness [6, Section 3.1]: *When can one enjoy the best of both group fairness and individual fairness?* We aim to provide a theoretically provable answer to this question when adopting statistical parity as the group fairness definition. In particular, we provide sufficient conditions for the compatibility between individual fairness and the Pareto optimal (with respect to utility) statistical parity L^2 -objective learning.

Definition 1.1 (Statistical parity) *Given a prediction random variable \hat{Y} and its corresponding sensitive random variable Z , the tuple (\hat{Y}, Z) satisfies statistical parity if*

$$\hat{Y} \perp Z.$$

That is, given any set $A \in \sigma(\hat{Y}), B \in \sigma(Z)$, we have

$$\mathbb{P}(\{\hat{Y} \in A\} \cap \{Z \in B\}) = \mathbb{P}(\{\hat{Y} \in A\})\mathbb{P}(\{Z \in B\}).$$

Here, $(\Omega, \mathcal{F}, \mathbb{P})$ is a probability space. $\hat{Y} : \Omega \rightarrow \mathcal{Y}$ and $Z : \Omega \rightarrow \mathcal{Z}$ are the random variables (or equivalently measurable functions) that map the elements from the underlying probability space Ω to the state spaces $(\mathcal{Y}, \sigma(\hat{Y}), \mathbb{P} \circ \hat{Y}^{-1})$ and $(\mathcal{Z}, \sigma(Z), \mathbb{P} \circ Z^{-1})$. $\sigma(S)$ denotes the sigma-algebra generated by the random variable S for $S \in \{\hat{Y}, Z\}$.

1.1 Related Work and Contribution

Beginning with the celebrated work ‘‘Fairness through Awareness’’ [9], there is now a sizable body of research studying group fair machine learning solutions, where group fairness is defined by statistical parity. The resulting approaches can be categorized into the following:

- I **Pre-processing:** The data are deformed before the training step. The goal is to preserve as much information as possible and also keep the deformed data representation independent of the sensitive variable [5, 18, 30, 28].
- II **In-processing:** The fairness definition is quantified and then integrated into the training process by penalizing unfair outcomes [2, 29]. It seems to be the natural

way in machine learning to add a penalty term related to unfairness quantification. Unfortunately, there is currently no theoretical guarantee of such machine-learned outcome in terms of both performance and fairness.

III Post-processing: The definition of fairness is enforced directly on the learning outcome [16, 7, 13, 26, 28]. This is the most straightforward approach because the methods directly deform the provided learning outcome to satisfy the statistical parity constraint.

In recent years, the connection between probability metric spaces and group fairness has received significant attention due to the following remarkable results: the optimal fair distribution of supervised learning, such as classification [16, 31] and regression [7, 13], can be characterized as the Fréchet mean of the learning outcome marginals on the Wasserstein space (or, more generally, metric space of probability measures), which is also known as the Wasserstein barycenter in the optimal transport literature. Thereafter, [28] generalizes the post-processing barycenter characterization to all supervised learning models that tend to estimate conditional expectation, including all classification and regression, via a pre-processing (or synthetic data) approach. Furthermore, [28] provides a provable Pareto frontier, which extends the barycenter characterization to the provable optimal trade-off between utility loss and any statistical disparity level in both post-processing and pre-processing approaches.

In this work, we study the compatibility between the optimal learning outcome for statistical parity and individual fairness in a post-processing setting where utility is quantified by the L^2 loss and fairness is defined by statistical parity. In particular, we study the condition under which the optimal post-processing group fair L^2 learning (see Problem 1 below) satisfies K-Lipschitz individual fairness or (ϵ, δ) individual fairness, respectively. In other words, we aim to fulfill both group fairness *and* individual fairness at the lowest utility cost.

The exploration of compatibility between group fairness, defined by statistical parity, and individual fairness can be traced back to at least [9], which highlights the potential conflict between these two fairness concepts and proposes methods to achieve both. Subsequent research has delved into this conflict heuristically [3, 6, 8, 17] and experimentally [32]. Despite these efforts, the question of when we can enjoy the best of both remains open [6, Section 3.1]. To our knowledge, only [9, 21] have endeavored to achieve both statistical parity and individual fairness, yet their approaches remain experimental. Our contribution lies in being the first to provide a theoretical analysis of the (in)compatibility and, if compatible, to propose provable methods for achieving the optimal trade-off between group fairness and individual fairness.

There is another line of work addresses the conflict between group fairness and individual fairness from the fair audit or multi-calibration perspective, which aims to provide similar treatment to an infinite class of groups defined by some class of functions of bounded complexity. For more details, interested readers can refer to [19, 15]. In this work, our focus is on group fairness defined by statistical parity.

1.2 Generalized Individual Fairness Definitions

For individual fairness, we consider K-Lipschitz Individual Fairness [9] and (ϵ, δ) Individual Fairness [11]. Both definitions share the same heuristics on treating similar individual

similarly. Both definitions can be considered as constraints on learned functions between the independent variable metric space $(\mathcal{X}, d_{\mathcal{X}})$ and the dependent variable metric space $(\mathcal{Y}, d_{\mathcal{Y}})$:

- A function $f : \mathcal{X} \rightarrow \mathcal{Y}$ satisfies *K-Lipschitz Individual Fairness (K-Lipschitz-IF)* if, for all $x_1, x_2 \in \mathcal{X}$, there exists $K \in \mathbb{R}^+$ such that

$$d_{\mathcal{Y}}(f(x_1), f(x_2)) < K d_{\mathcal{X}}(x_1, x_2).$$

- A function $f : \mathcal{X} \rightarrow \mathcal{Y}$ satisfies (ϵ, δ) *Individual Fairness ((ϵ, δ)-IF)* if, for all $x_1, x_2 \in \mathcal{X}$, there exists $(\epsilon, \delta) \in (\mathbb{R}^+)^2$ such that

$$d_{\mathcal{X}}(x_1, x_2) < \epsilon \implies d_{\mathcal{Y}}(f(x_1), f(x_2)) < \delta.$$

But the above individual fairness definitions are not general enough to perform compatibility analysis with respect to the optimal trade-off between utility and statistical parity for L^2 -objective learning. As shown in [7, 13, 28], the optimal statistical parity L^2 learning outcome and the Pareto frontier requires functions depending on the sensitive information or, in other words, functions taking (x, z) as argument:

$$f : \mathcal{X} \times \mathcal{Z} \rightarrow \mathcal{Y}. \tag{1}$$

For more general machine learning than the L^2 -objective ones, the logic underlying fairness through awareness [9] also suggests one to apply models depending on the sensitive information: in order to remove the undesirable influence of a sensitive variable on other variables, the model has to first acknowledge such a sensitive variable.

Remark 1.1 (Fairness through awareness) *“Fairness through Awareness” [9] is mostly referred as a work on individual fairness. But, at least to our understanding, the fundamental idea behind the work is to leverage the knowledge of the sensitive information and unfairness to diminish the unfairness in learning outcome. To that end, the authors in [9] proposed a method applying the sensitive information and randomness to achieve both statistical parity and individual fairness, which is also the goal of the present work. In fact, “Fairness through Awareness” is one of the main inspirations of our compatibility study.*

Therefore, we generalize the above fairness definitions so that they could be suitable for functions that depend on the sensitive variable.

Definition 1.2 (Uniform K-Lipschitz-IF) *A function $f : \mathcal{X} \times \mathcal{Z} \rightarrow \mathcal{Y}$ satisfies uniform K-Lipschitz Individual Fairness if, for all $x_1, x_2 \in \mathcal{X}$, there exists $K \in \mathbb{R}^+$ such that*

$$\sup_{z_1, z_2} d_{\mathcal{Y}}(f(x_1, z_1), f(x_2, z_2)) < K d_{\mathcal{X}}(x_1, x_2).$$

Definition 1.3 (Uniform (ϵ, δ) -IF) *A function $f : \mathcal{X} \times \mathcal{Z} \rightarrow \mathcal{Y}$ satisfies uniform (ϵ, δ) Individual Fairness if, for all $x_1, x_2 \in \mathcal{X}$, there exists $(\epsilon, \delta) \in (\mathbb{R}^+)^2$ such that*

$$d_{\mathcal{X}}(x_1, x_2) < \epsilon \implies \sup_{z_1, z_2} d_{\mathcal{Y}}(f(x_1, z_1), f(x_2, z_2)) < \delta.$$

To see the above definitions are generalizations of the corresponding original, one can consider $f : \mathcal{X} \rightarrow \mathcal{Y}$ as the subset of $f : \mathcal{X} \times \mathcal{Z} \rightarrow \mathcal{Y}$ which remains constant when z changes. Thereby, the above generalized definitions reduce to the original ones. In the rest of the present work, we stick with the generalized individual definitions.

Remark 1.2 (Main difference between (ϵ, δ) -IF and K-Lipschitz-IF) *Notice that uniform K-Lipschitz-IF implies uniform $(\epsilon, K\epsilon)$ -IF. Hence, the latter is usually considered as a relaxed version of the former. The key difference is that the (ϵ, δ) -IF definition allows different learning outcomes assigned to individuals with the same independent variable value x but different sensitive information z . That is, (ϵ, δ) -IF allows different learning outcome to be assigned to individuals (for example $\{1, 2\}$) who share the same qualification ($x_1 = x_2$) but have different sensitive information ($z_1 \neq z_2$):*

$$z_1 \neq z_2 \implies f(x_1, z_1) \neq f(x_2, z_2)$$

even if $x_1 = x_2$. This relaxation is compatible with the fundamental idea underlying the optimal statistical parity L^2 learning (or more generally fairness through awareness) to let sensitive information dependent functions be applied to the same x to achieve statistical parity. As we will show later, such a relaxation results in different compatibility with the optimal statistical parity L^2 learning between the two individual fairness definitions.

Finally, since we assume $(\mathcal{Y}, \|\cdot\|)$ is a Euclidean space, the distance metric $d_{\mathcal{Y}}$ is induced by the Euclidean norm in our post-processing setting. It follows that the individual fairness constraints become:

- A function $f : \mathcal{Y} \times \mathcal{Z} \rightarrow \mathcal{Y}$ satisfies uniform K-Lipschitz Individual Fairness if, for all $y_1, y_2 \in \mathcal{Y}$, there exists $K \in \mathbb{R}^+$ such that

$$\sup_{z_1, z_2} \|f(y_1, z_1) - f(y_2, z_2)\| < K\|y_1 - y_2\|.$$

- A function $f : \mathcal{Y} \times \mathcal{Z} \rightarrow \mathcal{Y}$ satisfies uniform (ϵ, δ) Individual Fairness if, for all $y_1, y_2 \in \mathcal{Y}$, there exists $(\epsilon, \delta) \in (\mathbb{R}^+)^2$ such that

$$\|y_1 - y_2\| < \epsilon \implies \sup_{z_1, z_2} \|f(y_1, z_1) - f(y_2, z_2)\| < \delta.$$

Lastly, we apply Definition 1.3 to define the admissible set of functions or maps in the post-processing step to satisfy (ϵ, δ) -IF:

Definition 1.4 ((ϵ, δ) -IF constrained admissible set)

$$\mathcal{D}_{(\epsilon, \delta)\text{-IF}} := \{f \in L^2(\mathcal{Y} \times \mathcal{Z}, \mathcal{Y}) : \|y_1 - y_2\| \leq \epsilon \implies \sup_{z_1, z_2} \|f(y_1, z_1) - f(y_2, z_2)\| \leq \delta\}$$

One can also define the admissible set to satisfy the K-Lipschitz-IF constraint. But, as shown later in Theorem 3.1, there exists an intrinsic incompatibility between the optimal statistical parity L^2 learning solution (or any non-trivial Pareto optimal solution) and the K-Lipschitz-IF constraint. Such intrinsic incompatibility prevents further analysis. Therefore, we skip the definition of a K-Lipschitz-IF constrained admissible set.

1.3 Problem Setting

To provide our problem setting, we first develop the L^2 objective function and fairness constraints that are suitable for our post-processing approach. First, by the L^2 -objective learning assumption, the post-processing objective is to minimize the utility sacrifice quantified by the L^2 distance between the provided \hat{Y} and the post-processed final output $f(\hat{Y}, Z)$. Here, $f : \mathcal{Y} \times \mathcal{Z} \rightarrow \mathcal{Y}$ is the post-processing step designed to diminish statistical disparity in the originally provided learning outcome \hat{Y} . Also, it follows from the triangle inequality that

$$\underbrace{\|Y - f(\hat{Y}, Z)\|_2}_{\text{total loss}} \leq \underbrace{\|Y - \hat{Y}\|_2}_{\text{training loss}} + \underbrace{\|\hat{Y} - f(\hat{Y}, Z)\|_2}_{\text{post-processing loss}}.$$

Here, $\|Y_1 - Y_2\|_2 := \int_{\Omega} \|Y_1(\omega) - Y_2(\omega)\| d\mathbb{P}(\omega)$, and $\|\cdot\|$ denotes the Euclidean norm on \mathcal{Y} . Therefore, provided only access to (\hat{Y}, Z) , the best a post-processing step can do is to minimize the second term on the right-hand side. (The first term is assumed to be optimized during the training step that results in the provided \hat{Y} .) Therefore, the post-processing step objective is

Problem 1 (Optimal post-processing statistical parity L^2 -objective learning)

$$\inf_{f \in L^2(\mathcal{Y} \times \mathcal{Z}, \mathcal{Y})} \{ \|\hat{Y} - f(\hat{Y}, Z)\|_2^2 : f(\hat{Y}, Z) \perp Z \}. \quad (2)$$

Here, $L^2(\mathcal{Y} \times \mathcal{Z}, \mathcal{Y})$ is the admissible function set consisting of all the square integrable measurable functions from $\mathcal{Y} \times \mathcal{Z}$ to \mathcal{Y} (See Remark 1.3 below for a generalization to all measurable functions) and constraint $f(\hat{Y}, Z) \perp Z$ guarantees the post-processed output satisfies statistical parity definition. \hat{Y} is the provided learning outcome, and $f(\hat{Y}, Z)$ is the post-processed learning outcome. Z or z -dependent functions $f(\cdot, z)$ are allowed due to the same reason underlying fairness through awareness [9] (See Remark 1.1 below for an explanation). The loss function $\|\hat{Y} - f(\hat{Y}, Z)\|_2^2$ aims to maximize utility by minimizing the L^2 -norm between the provided learning outcome \hat{Y} and post-processed outcome $f(\hat{Y}, Z)$.

Remark 1.3 (Choice of the admissible set) *We adopt all square-integrable measurable functions ($L^2(\mathcal{Y} \times \mathcal{Z}, \mathcal{Y})$) due to the recent development of neural networks, which are able to estimate arbitrary measurable functions [23]. We note here that Problem 1 does not change if we change L^2 to all $\mathcal{Y} \times \mathcal{Z}/\mathcal{Y}$ -measurable functions. That is, the optimal measurable function happens to be square-integrable under our assumptions. But uniqueness becomes almost sure uniqueness if one replaces $L^2(\mathcal{Y} \times \mathcal{Z}, \mathcal{Y})$ with the set of all $\mathcal{Y} \times \mathcal{Z}/\mathcal{Y}$ -measurable functions.*

Furthermore, when the optimal solution to Problem 1 is incompatible with individual fairness, we first relax the hard statistical parity constraint of individual fairness by allowing different statistical disparity tolerance levels to characterize a Pareto frontier between utility loss and statistical disparity, then study which portion of the frontier (equivalently, which of the Pareto solutions) is compatible with the individual fairness requirement.

In particular, we first quantify the statistical disparity of \hat{Y} with respect to Z using the notion of *Wasserstein disparity*, which measures the pairwise Wasserstein distance among

sensitive group (see Definition 2.1 for details). We denote the Wasserstein disparity by $D(\hat{Y}, Z)$ and for now only note that $D(\hat{Y}, Z) = 0 \iff \hat{Y} \perp Z$. See Section 2.1 for a more detailed justification of this quantification of statistical parity.

Therefore, by relaxing the statistical parity constraint to different tolerance levels of the Wasserstein disparity, $D(\hat{Y}, Z) < d$ for $d \in [0, \infty)$, we obtain the following

Problem 2 (Post-processing L^2 -objective learning Pareto frontier)

$$\inf_{f \in L^2(\mathcal{Y} \times \mathcal{Z}, \mathcal{Y})} \{ \|\hat{Y} - f(\hat{Y}, Z)\|_2^2 : D((f(\hat{Y}, Z), Z)) < d \}. \quad (3)$$

We note that Problem 2 reduces to Problem 1 at $d = 0$ due to the fact that $D(\hat{Y}, Z) = 0 \iff \hat{Y} \perp Z$. Indeed, as we will argue in Section 2.1, Problem 2 is a natural relaxation of Problem 1.

Furthermore, we claim that Problem 2 characterizes the Pareto frontier between utility loss (quantified by the L^2 norm) and statistical disparity (quantified by the Wasserstein disparity). Indeed, if we want to achieve lower utility loss than the infimum of Problem 2 for some fixed d , then it is necessary to increase the statistical disparity tolerance level above d . On the other hand, if we want to achieve a statistical disparity level lower than d , it is necessary to have the utility loss more than the infimum of Problem 2 at d . Therefore, Problem 2 characterizes the Pareto frontier or the optimal trade-off by definition.

Now, to analyze the compatibility between the optimal L^2 learning and individual fairness, we impose an additional individual fairness constraint on the admissible post-processing functions, namely we assume

$$\mathcal{D}_{(\epsilon, \delta)\text{-IF}} \subset L^2(\mathcal{Y} \times \mathcal{Z}, \mathcal{Y}),$$

where $\mathcal{D}_{(\epsilon, \delta)\text{-IF}}$ has been introduced in Definition 1.4. Hence, the compatibility between the optimal group fair L^2 learning and individual fairness can be studied by comparing Problem 1 to

Problem 3 ((ϵ, δ) -IF optimal post-processing statistical parity L^2 -objective learning)

$$\inf_{f \in \mathcal{D}_{(\epsilon, \delta)\text{-IF}}} \{ \|\hat{Y} - f(\hat{Y}, Z)\|_2^2 : f(\hat{Y}, Z) \perp Z \}. \quad (4)$$

Furthermore, the compatibility between the Pareto optimal group fair L^2 learning (equivalently, portion of the Pareto frontier) and individual fairness can be studied by comparing Problem 2 to

Problem 4 ((ϵ, δ) -IF post-processing L^2 -objective learning Pareto frontier)

$$\inf_{f \in \mathcal{D}_{(\epsilon, \delta)\text{-IF}}} \{ \|\hat{Y} - f(\hat{Y}, Z)\|_2^2 : D(f(\hat{Y}, Z), Z) < d \}. \quad (5)$$

More specifically, since Problem 1 and 2 can be considered as the respective relaxation of Problem 3 and 4, the optimal solution of the former necessarily results in lower or equal value than the optimal solution to the respective later. Hence, we study when the optimal solution of Problem 1 (respectively 2) equals the one of Problem 3 (respectively 4) to determine the compatibility, due to the uniqueness results of the optimal solutions for the former problems below. That is, we define compatibility as the followings:

I The optimal statistical parity L^2 learning is compatible with individual fairness if

$$\text{Problem 1} \equiv \text{Problem 3.} \tag{6}$$

II The Pareto optimal statistical parity L^2 learning at Wasserstein disparity tolerance level d is compatible with individual fairness if

$$\text{Problem 2} \equiv \text{Problem 4 at a fixed } d. \tag{7}$$

Remark 1.4 (Compatibility analysis for pre-processing) *While we focus on the post-processing setting, we note that the compatibility analysis derived in the present work can be easily modified for the following two cases: (1) the optimal (not conditioned on post-processing, in-processing, or pre-processing) statistical parity L^2 learning when the conditional expectation is available, and (2) an optimal pre-processing statistical parity L^2 learning or the optimal statistical parity data representation for L^2 learning. In case (1), we can replace \hat{Y} by the conditional expectation $\mathbb{E}(Y|X, Z)$ and apply the post-processing composition result in Theorem 5.1 or 5.2. In case (2), we can replace \hat{Y} by the qualification or independent variable X and apply the pre-processing composition result in Theorem 5.1 or 5.2. We refer interested readers to [28] for more details on the two cases. The present work focuses on the post-processing case due to its straight-forwardness. We defer the modifications of the compatibility analysis for the other two cases to further work.*

In the rest of the current section, we answer the following questions that motivates the problem setting in the present work:

- I Why are we considering the (Pareto) optimal statistical parity L^2 learning instead of the original statistical parity definition itself in the compatibility analysis?
- II Which individual fairness definitions are adopted in the post-processing compatibility analysis of the present work?
- III Why is individual fairness considered as an additional constraint while keeping statistical parity as the main constraint, but not the other way around?

Section 1.4 first shows that utility maximization enhances the original statistical parity definition by overcoming the major insufficiency of it as mentioned in [9, 14]. Thereby, the optimal statistical parity L^2 learning and Pareto frontier provide a better group fairness concept when comparing to the original statistical parity definition: treating relatively (within one’s own sensitively group) similar individuals similarly. See Remark 1.5 below for the comparison between relative similarity and absolute similarity. Section 1.2 introduces the two commonly used individual fairness definitions: K -Lipschitz and (ϵ, δ) individual fairness, then generalizes them due to technical reasons. Finally, Section 1.5 explains why we choose individual fairness as an additional constraint. Finally, we provide a road map of the main results in the present work.

1.4 Statistical Parity Enhanced by Utility Optimization

In this subsection, we show how utility optimization helps to overcome the major insufficiency of the original statistical parity concept as a group fairness definition. There are three major criticisms on statistical parity as a fairness definition [9, 14]: (1) reduced utility, (2) self-fulfilling prophecy, and (3) subset targeting. But, as shown in [28], the first two of the three insufficiencies can be (partially) fixed by maximizing the utility and adopting the Pareto frontier.

- I **Reduced utility:** Since statistical parity is an additional constraint on the learning outcome, it necessarily compromises utility. While the utility sacrifice is unavoidable, a Pareto frontier allows practitioners to choose a statistical disparity level that can be tolerated in their own use case and apply the Pareto optimal solution at the chosen tolerance level. In that case, practitioners can trade statistical disparity tolerance for utility and vice versa.
- II **Self-fulfilling prophecy:** The expression ‘self-fulfilling prophecy’ refers to random, careless, or malicious selection in some of the sensitive groups, even in the case statistical parity is satisfied. But the utility maximization enforces careful selection in all sensitive groups and therefore prevents random, careless, or malicious selection in any of the sensitive groups from happening, regardless of the size of the sensitive group.

We refer interested readers to [28] for a more detailed explanation on how the connection between utility optimization for statistical parity and the optimal multi-marginal matching (equivalently, n -coupling problem or multidimensional Monge-Kantorovich problem) solve the above two insufficiencies and provides explainability. Also, we refer readers to [12] for the technical details on the equivalence between the Wasserstein barycenter problem and the multidimensional Monge-Kantorovich problem.

Remark 1.5 (Relative vs absolute similarity) *Interestingly, by solving the above listed insufficiencies, the utility-maximized statistical parity also provides similar treatment to similar individuals, but similarity is now defined based on the position in the individual’s own sensitive group. It is different from the similarity in the individual fairness concept, which is defined based on the individual’s position in the whole group. Hence, the compatibility between individual fairness and the (Pareto) optimal (with respect to utility) statistical parity L^2 -objective learning we study here is, in its nutshell, the compatibility between the following two view point of fairness:*

- **Relative similarity:** *Individuals who share the similar positions in his or her own sensitive group should receive similar learning outcome.*
- **Absolute similarity:** *Individuals sharing the similar positions in the entire big group (without considering sensitive groups) should receive similar learning outcome.*

The following example provides an intuition of the relative similarity and how utility optimization solves the two mentioned insufficiencies.

Example 1.1 (College admission) *Let X be the qualification consisting of only the standard test score when admitted, Y the GPA after four years, the sensitive information Z be gender, and \hat{Y} be the predicted GPA in four years. In this case, group fairness requires the predicted GPA to have the same distribution for male students and female students, while individual fairness requires that the learned model gives students sharing similar standard test scores similar predicted GPA, regardless of their gender.*

For simplicity, we assume that: (1) the ratio between the admitted male and female students is 1 : 1, and (2) students in the k^{th} percentile (ranked by the predicted GPA) of the male (respectively female) group all share the same prediction $\hat{y}(m, k)$ (respectively $\hat{y}(f, k)$). The optimal post-processing statistical parity L^2 learning applies the following steps to achieve statistical parity with minimum L^2 loss:

- *Optimality: For each k , assign both the k^{th} percentile of male and female students the prediction $\bar{y}(k) := 0.5\hat{y}(m, k) + 0.5\hat{y}(f, k)$.*
- *Pareto optimality: For each k , assign the k^{th} percentile of male and female students the prediction $\hat{y}(m, k, t) := (1 - t)\hat{y}(m, k) + t\bar{y}(k)$ and $\hat{y}(f, k, t) := (1 - t)\hat{y}(f, k) + t\bar{y}(k)$ respectively, for $t \in [0, 1]$.*

Here, the factor 0.5 in the optimality is due to the first assumption of the 1 : 1 ratio between male and female students, and $t \in [0, 1]$ in the Pareto optimality is an interpolation parameter that determines the exact position on the Pareto frontier.

How does the above method solve the two insufficiencies? (1) The Pareto frontier provides the optimal partial remedy when it is necessary to sacrifice a significant amount of utility to achieve statistical parity. (2) The matching ensures that the students who are at the same percentile in his or her own group would share the same learning outcome, so that higher ranked students still receive higher predicted GPA within his/her own sensitive group.

1.5 Individual Fairness as an Additional Constraint

Now, we discuss the reason of choosing individual fairness as an additional constraint to an optimization problem aiming for group fairness, but not the other way around.

In our post-processing setting, we assume no knowledge of X , Y , or \mathcal{F} in the training process $\inf_{f \in \mathcal{F}} \|Y - f(X)\|_2^2$ or $\inf_{f \in \mathcal{F}} \|Y - f(X, Z)\|_2^2$. \hat{Y} is the only provided information together with the sensitive information, and hence considered as the qualification variable to quantify similarity in the post-processing setting. Therefore, there is no reason to deform the provided \hat{Y} for individual fairness purpose only. For example, the identity map automatically satisfies the 1-Lipschitz-IF or (ϵ, ϵ) -IF.

On the other hand, if we consider the post-processing step for statistical parity, the rescue step can violate the individual fairness constraint when deform the provided learning outcome. Therefore, it is natural to consider individual fairness as an additional constraint to the optimal statistical parity learning problem in our post-processing setting.

Lastly, we refer interested readers to [10] for more general and detailed discussion on why it is better to consider individual fairness as a necessary condition and hence an additional constraint.

To end the current section, we summarize the questions the present work targets and provide the corresponding informal answer.

1. When is the optimal post-processing statistical parity L^2 learning compatible with K-Lipschitz individual fairness?

Theorem 3.1 shows that, unless the learning outcome automatically satisfies statistical parity, neither the optimal post-processing statistical parity L^2 learning nor the Pareto frontier (any non-trivial optimal trade-off between utility loss and statistical disparity) is compatible with the K-Lipschitz individual fairness definition under mild assumptions. That is, the optimal statistical parity L^2 learning and the Pareto frontier has an inherent conflict with the K-Lipschitz individual fairness definition.

2. When is the optimal post-processing statistical parity L^2 learning compatible with (ϵ, δ) individual fairness?

Lemma 3.1 shows that, provided certain assumption of relationship between ϵ and δ , one can have both optimal statistical parity L^2 learning and (ϵ, δ) individual fairness.

3. If the optimal post-processing statistical parity L^2 learning and (ϵ, δ) individual fairness are not compatible, which portion of the Pareto frontier (or the Pareto optimal solutions) becomes compatible with the (ϵ, δ) individual fairness?

Theorem 4.1 shows that, in the case of incompatibility between optimal statistical parity L^2 learning and (ϵ, δ) individual fairness, it is guaranteed to have a non-trivial portion of Pareto frontier (between L^2 loss and statistical disparity) compatible with (ϵ, δ) individual fairness, when $\epsilon < \delta$. That is, (ϵ, δ) individual fairness does not affect the Pareto optimality up to reduction of a certain level of statistical disparity in the learning outcome.

4. Now, assume $\hat{Y} = g(X, Z)$ or $g(X)$ for some trained model $g : \mathcal{X} \times \mathcal{Z} \rightarrow \mathcal{Y}$, what is compatibility guarantee for the post-processed learning result which composes the trained model with the (Pareto) optimal post-processing learning steps? (Or, as mentioned in Remark 1.4, what is compatibility guarantee for the pre-processed learning result which composes the (Pareto) optimal pre-processing learning steps with the trained model?)

Theorem 5.1 and Theorem 5.2 provide compatibility analysis for the learning outcome which composes a trained model with a post-processing step (or composes a pre-processing step with a trained model) under the (ϵ, δ) individual fairness and K-Lipschitz individual fairness assumption, respectively. So that researchers and practitioners can obtain compatibility result when composing the proposed post-processing (or pre-processing after modification) with some in-processing methods with individual fairness guarantee.

The remainder of this paper is structured as follows: Section 2 introduces the problem setting, presents optimal transport preliminaries, and reviews (Pareto) optimal L^2 learning for statistical parity—a necessary foundation for our main results. Section 3 explores the compatibility between individual fairness (K-Lipschitz-IF and (ϵ, δ) -IF) and optimal statistical parity L^2 learning. In Section 4, we address cases where optimal L^2 fair learning conflicts with a fixed (ϵ, δ) -individual fairness requirement, examining the compatible portion of the Pareto frontier. Section 6 presents empirical studies validating the theoretical results.

2. Preliminaries on the (Pareto) Optimal Fair L^2 Learning

In this section, we first define a quantification of statistical disparity using tools from optimal transport theory to relax the hard independence constraint in the statistical parity definition. We then review the theoretical results on the existence and uniqueness of the optimal solutions to Problem 1 and 2, on which our compatibility result is developed.

2.1 Quantification of Statistical Disparity

Before defining the statistical disparity quantification or relaxation, we first provide a brief review of Wasserstein space and barycenter and fix the notations used in the present work. Given $\mu, \nu \in \mathcal{P}(\mathbb{R}^d)$ where $\mathcal{P}(\mathbb{R}^d)$ denotes the set of all the probability measures on \mathbb{R}^d ,

$$\mathcal{W}_2(\mu, \nu) := \left(\inf_{\lambda \in \Pi(\mu, \nu)} \left\{ \int_{\mathbb{R}^d \times \mathbb{R}^d} \|x_1 - x_2\|^2 d\lambda(x_1, x_2) \right\} \right)^{\frac{1}{2}}.$$

Here, $\Pi(\mu, \nu) := \{\pi \in \mathcal{P}((\mathbb{R}^d)^2) : \int_{\mathbb{R}^d} d\pi(\cdot, \nu) = \mu, \int_{\mathbb{R}^d} d\pi(u, \cdot) = \nu\}$. $(\mathcal{P}_2(\mathbb{R}^d), \mathcal{W}_2)$ is called the Wasserstein space, where

$$\mathcal{P}_2(\mathbb{R}^d) := \left\{ \mu \in \mathcal{P}(\mathbb{R}^d) : \int_{\mathbb{R}^d} \|x\|^2 d\mu < \infty \right\}.$$

Given $\{\mu_z\}_{z \in \mathcal{Z}} \subset (\mathcal{P}_2(\mathbb{R}^d), \mathcal{W}_2)$ for some index set \mathcal{Z} , their barycenter with weights $\lambda \in \mathcal{P}(\mathcal{Z})$ is

$$\bar{\mu} := \operatorname{argmin}_{\mu \in \mathcal{P}_2(\mathbb{R}^d)} \left\{ \int_{\mathcal{Z}} \mathcal{W}_2^2(\mu_z, \mu) d\lambda(z) \right\}. \quad (8)$$

In [1, 20], the authors prove the existence and uniqueness of $\bar{\mu}$ under the assumption that $\lambda(\{\mu_z \in \mathcal{P}_{2,ac}(\mathbb{R}^d)\}) > 0$, where $\mathcal{P}_{2,ac}(\mathbb{R}^d) := \{\mu \in \mathcal{P}_2 : \mu \ll \mathcal{L}\}$ and \mathcal{L} is the Lebesgue measure on \mathbb{R}^d . Therefore, in the rest of the present work, we often assume $\{\mu_z\}_z := \{\mathcal{L}(\hat{Y}_z)\}_z \subset \mathcal{P}_{2,ac}(\mathbb{R}^d)$, where \hat{Y}_z denotes the conditional distribution of \hat{Y} on the event $\{Z = z\}$ defined by the disintegration theorem [25]. For an empirical data set $\{(\hat{y}_i, z_i)\}_{i=1}^N$, \hat{Y}_z is the random variable with uniform distribution on $\{\hat{y}_i : z_i = z\}$.

In order to relax the hard statistical parity or independence constraint, one needs to define a quantification of statistical disparity. To that end, we apply the average pairwise Wasserstein (W_2) distance among (the provided predictions of) sensitive groups to quantify statistical disparity and show the desirable properties of the quantification.

Definition 2.1 (Wasserstein disparity)

$$D(\hat{Y}, Z) := \left(\int_{\mathcal{Z}^2} \mathcal{W}_2^2(\mu_{z_1}, \mu_{z_2}) d\lambda^{\otimes 2}((z_1, z_2)) \right)^{\frac{1}{2}}, \quad (9)$$

where $\mu_{z_s} := \mathbb{P} \circ \hat{Y}_{z_s}^{-1}$ for $s \in \{1, 2\}$ and $\lambda := \mathbb{P} \circ Z^{-1}$ denote the law or distribution of \hat{Y}_{z_s} and Z respectively.

(To clarify, we note that $\hat{Y}_{z_s}^{-1} : \sigma(\hat{Y}) \rightarrow \mathcal{F}$ finds the pre-image of an event on the state space $(\mathcal{Y}, \sigma(\hat{Y}))$ in the underlying probability space $(\Omega, \mathcal{F}, \mathbb{P})$.) The Wasserstein disparity defined above has the following properties:

1 Lemma 2.1 (Characterization of statistical parity)

$$D(\hat{Y}, Z) = 0 \text{ if and only if } \hat{Y} \perp Z.$$

Proof See Section A. ■

- 2 (Physics interpretation)** In physics, $D(\hat{Y}, Z)$ can be interpreted as the expected minimum amount of work required to remove the distributional discrepancy between two randomly chosen (according to the distribution of Z or λ) sensitive groups on the provided learning outcome \hat{Y} .
- 3 (Characterization of the Pareto frontier)** As shown in [28], by adopting the Wasserstein disparity to relax the hard independence constraint, the optimal trade-off between the L^2 loss and statistical disparity (quantified by Wasserstein disparity) is characterized by the geodesic path from the conditional (on the event $\{Z = z\}$) distributions of \hat{Y} to their barycenter on the Wasserstein space.

Due to the listed properties, the average pairwise Wasserstein definition of statistical disparity is a natural quantification when studying the trade-off between statistical disparity and L^2 loss.

2.2 Optimal Statistical Disparity L^2 Learning and the Pareto Frontier

In this subsection, we summarize the existence and uniqueness results of the optimal solutions to Problem 1 and 2 that are needed later in the compatibility study.

[28] shows that the optimal statistical parity L^2 learning problem has an unique solution that coincides with the Wasserstein barycenter. In particular, let $\mu_z := \mathcal{L}(\hat{Y}_z)$ and $T_z : \mathcal{Y} \rightarrow \mathcal{Y}$ be the optimal transport map [4] such that

$$(T_z)_\# \mu = \mu_z,$$

where $T_\# \mu := \mu \circ T^{-1}$ denotes the push-forward measure of μ under the map T , the following result characterizes the unique solution to Problem 1:

Lemma 2.2 (Optimal fair L^2 learning characterization [28]) *Assume that \hat{Y} has sensitive conditional distributions satisfying $\{\mathcal{L}(\hat{Y}_z)\}_{z \in \mathcal{Z}} =: \{\mu_z\}_{z \in \mathcal{Z}} \subset \mathcal{P}_{2,ac}(\mathcal{Y})$, then there exists a unique $f^* \in L^2(\mathcal{Y} \times \mathcal{Z}, \mathcal{Y})$ defined by*

$$f^*(\cdot, z) = T_z(\cdot) \tag{10}$$

for λ -a.e. $z \in \mathcal{Z}$ such that

$$\|\hat{Y} - f^*(\hat{Y}, Z)\|_2^2 = \inf_{f \in L^2(\mathcal{Y} \times \mathcal{Z}, \mathcal{Y})} \{\|\hat{Y} - f(\hat{Y}, Z)\|_2^2 : f(\hat{Y}, Z) \perp Z\} = \underbrace{\int_{\mathcal{Z}} \mathcal{W}_2^2(\mu_z, \bar{\mu}) d\lambda}_{\text{independence projection loss}} .$$

Remark 2.1 (Generalization to all measurable functions) *The result in Theorem 2.2 does not change if we replace the admissible set $L^2(\mathcal{Y} \times \mathcal{Z}, \mathcal{Y})$ with all measurable functions $f : \mathcal{Y} \times \mathcal{Z} \rightarrow \mathcal{Y}$, except the uniqueness result now becomes almost sure uniqueness. That is, f^* is the almost sure unique solution to*

$$\inf_{f: \mathcal{Y} \times \mathcal{Z} \rightarrow \mathcal{Y}} \{ \|\hat{Y} - f(\hat{Y}, Z)\|_2^2 : f(\hat{Y}, Z) \perp Z \}.$$

Therefore, in practice we can apply e.g. a neural network to approximate $\{T_z\}_z$ and therefore f^* without worrying about the square integrability.

From now on, we denote the minimum value of Problem 1 or independence projection loss by $V(\hat{Y}, Z)$. That is,

$$V(\hat{Y}, Z) := \min_{f \in L^2(\mathcal{Y} \times \mathcal{Z}, \mathcal{Y})} \{ \|\hat{Y} - f(\hat{Y}, Z)\|_2 : f(\hat{Y}, Z) \perp Z \}. \quad (11)$$

V is an important quantity when studying the Pareto frontier because $t := 1 - \frac{d}{\sqrt{2}V}$ serves as the time parameter for the constant-speed geodesics which characterizes the Pareto frontier. Furthermore, V has the following interpretation in physics: Given (\hat{Y}, Z) , V is the minimum amount of work or energy required to deform \hat{Y} to satisfy statistical parity.

If one does not require strict statistical disparity, the Wasserstein disparity tolerance level d can be non-zero. For $d \in [0, \infty)$, the Problem 2 characterizes the Pareto frontier between $\|\hat{Y} - f(\hat{Y}, Z)\|_2^2$, the L^2 utility loss resulting from the post-processing step f , and $D(f(\hat{Y}, Z), Z)$, which is the Wasserstein disparity remaining in the post-processed outcome $f(\hat{Y}, Z)$. Let $f^*(\cdot, z)$ denote solution to Problem 1 defined in Lemma 2.2 and let

$$f^*(t)(\cdot, z) := (1 - t) \text{Id} + t f^*(\cdot, z), t \in [0, 1]$$

denote the McCann interpolation [22] between the identity map (Id) and the optimal transport map ($f^*(\cdot, z) = T_z$) for λ -a.e. $z \in \mathcal{Z}$. The following result shows the closed-form unique solution to the Problem 2 for each disparity tolerance level $d \in [0, \infty)$.

Lemma 2.3 (Pareto optimal fair L^2 -objective learning [28]) *Given (\hat{Y}, Z) satisfying $\mu_z \in \mathcal{P}_{2,ac}$, λ -a.e. and V the independence projection loss defined in (11), then*

$$f_d(\hat{Y}, Z) := \begin{cases} f^*(1 - \frac{d}{\sqrt{2}V})(\hat{Y}, Z), & \text{if } d \in [0, \sqrt{2}V] \\ \hat{Y}, & \text{if } d \in (\sqrt{2}V, \infty) \end{cases} \quad (12)$$

are the unique solutions to Problem 2 for $d \in [0, \infty)$.

Provided the unique optimal solutions to Problem 1 and 2, we are ready to study the compatibility between the (Pareto) optimal statistical parity L^2 learning and the (K-Lipschitz-IF and (ϵ, δ) -IF) individual fairness definitions.

3. Compatibility between the Optimal Statistical Parity L^2 Learning and Individual Fairness

In this section, we study the compatibility between the optimal statistical parity L^2 -objective learning (the solution to Problem 1) and the two individual fairness definitions: K-Lipschitz-IF and (ϵ, δ) -IF.

3.1 Optimal Statistical Parity L^2 Learning and Lipschitz-IF

To start, we show the inherent incompatibility between the (Pareto) optimal statistical parity L^2 learning and the K-Lipschitz-IF definition for any $K > 0$. The incompatibility is due to the strict prevention from applying different function or maps to the same $y \in \mathcal{Y}$, which unfortunately is necessary to achieve (Pareto) optimality for statistical parity in L^2 learning.

Theorem 3.1 (Incompatibility of optimal statistical parity L^2 learning and Lipschitz-IF)

Under the following two assumptions

- 1 $\hat{Y} \not\perp Z$ (statistical parity is not automatically satisfied by the provided prediction \hat{Y}),
- 2 $\text{Domain}(f^*(\cdot, z)) = \mathcal{Y}, \lambda - a.e. z \in \mathcal{Z}$ (the optimal fair L^2 learning is capable of making predictions for all individuals, even unobserved),

neither f^ nor $\{fa\}_{d \in [0, \sqrt{2}V]}$ is compatible with the K-Lipschitz-IF definition for any $K > 0$.*

The above result shows that if the two assumptions are satisfied, then the K-Lipschitz-IF and optimal statistical parity L^2 learning are incompatible. To see the inherent conflict, it remains to show that the two assumptions are satisfied in most of the practical or interesting machine learning problems:

Remark 3.1 *Here, we discuss the practicality of the two assumptions in the Theorem 3.1*

- 1 *For assumption 1, we are excluding the trivial case that statistical parity is automatically satisfied by the provided learning outcome. In the trivial case, there is neither necessity to trade utility for statistical parity nor need for a post-processing step. Also, we argue that is very unlikely for one to have $\hat{Y} \perp Z$ in practice.*
- 2 *For assumption 2, although not impossible, it is highly unlikely that there are no counterparts (similar data points w.r.t. qualification or the provided learning outcome) among different sensitive groups. Moreover, even if different sensitive groups are mutually exclusive w.r.t. the qualification or the provided learning outcome in the sample data, one should not make that assumption for prediction purpose. The key of supervised learning is to make predictions. Hence, a desirable post-processing step (for a fair supervised learning model) should provide us a fair prediction for every possible learning outcome in any sensitive group.*

3.2 Optimal Fair L^2 Learning and (ϵ, δ) -IF

Due to the inherent incompatibility result above, we adopt the (ϵ, δ) -IF definition, which shares the same heuristic concept of individual fairness with the Lipschitz definition but adds more flexibility by allowing the individuals who share the same qualification data to be mapped to different learning outcome due to their different sensitive information.

The following result provides us a straight-forward sufficient condition for the compatibility between the optimal statistical parity L^2 learning and the (ϵ, δ) -IF.

Lemma 3.1 (Compatibility between optimal statistical parity L^2 learning and (ϵ, δ) -IF)
 Assume $\{\mu_z\}_z \subset \mathcal{P}_{2,ac}(\mathcal{Y})$, let f^* be the (unique) solution to Problem 1, and $L(f^*) := \sup_{(y,z)} \|f(y, z) - y\|$, then for all $(\epsilon, \delta) \in (\mathbb{R}^+)^2$ that satisfy $L(f^*) \leq \frac{\delta - \epsilon}{2}$, f^* is the unique solution to

$$\inf_{f \in \mathcal{D}_{(\epsilon, \delta)-IF}} \{\|\hat{Y} - f(\hat{Y}, Z)\|_2 : f(\hat{Y}, Z) \perp Z\}. \quad (13)$$

The above result shows that, if the solution of Problem 2.2 satisfies the assumption $\sup_{(y,z)} \|f^*(y, z) - y\| \leq \frac{\delta - \epsilon}{2}$, then we have

$$\text{Problem 1} \equiv \text{Problem 3} \quad (14)$$

which is equivalent to the compatibility between the optimal statistical parity L^2 learning and the (ϵ, δ) -IF requirement.

Proof

$$\begin{aligned} \|f^*(y_1, z_1) - f^*(y_2, z_2)\| &\leq \|f^*(y_1, z_1) - y_1\| + \|y_1 - y_2\| + \|y_2 - f^*(y_2, z_2)\| \\ &\leq 2(\sup \|f^*(y, z) - y\|) + \epsilon \\ &\leq 2\frac{\delta - \epsilon}{2} + \epsilon = \delta \end{aligned}$$

where the second inequality follows from the assumption. ■

Remark 3.2 (Sharpness of the assumption upper bound) *The upper bound condition is sharp without any further assumptions on the sensitive marginal distributions of the learning outcome. To construct a counter-example, one can consider a learning outcome with two sensitive Gaussian marginals, $\hat{Y}_i \sim \mathcal{N}(m_i, \sigma)$ for $i \in \{1, 2\}$, with the same standard deviation but different means with $\|m_1 - m_2\| = \delta$. Now, assume*

$$\sup_{(y,z)} \|f^*(y, z) - y\| = \frac{\delta - \epsilon}{2} - h,$$

we show that V cannot be achieved with such a restriction for any $h > 0$. Indeed, since \hat{Y}_i 's are Gaussian with the same standard deviation, their barycenter is also Gaussian with the same standard deviation. Therefore, the optimal transport maps $f^*(y, z = 1) := y + \frac{\mu_2 - \mu_1}{2}$ and $f^*(y, z = 2) := y + \frac{\mu_1 - \mu_2}{2}$ are rigid translations. Hence,

$$\{f^*(\cdot, z)\}_z \not\subset \mathcal{D}_{(\epsilon, \delta)-IF}$$

because f^* satisfies $\|f^*(y, z) - y\| = \frac{\|\mu_1 - \mu_2\|}{2} = \frac{\delta}{2}, \forall y \in \mathcal{Y}$ when $\epsilon = 0$. Hence, for any $h > 0$, we have

$$V < \inf_{f \in \mathcal{D}_{(\epsilon, \delta)-IF}} \{\|\hat{Y} - f(\hat{Y}, Z)\|_2 : f(\hat{Y}, Z) \perp Z\}.$$

That is, the optimal solution is not compatible with the (ϵ, δ) -IF definition if $\sup_{(y,z)} \|f(y, z) - y\| = \frac{\delta - \epsilon}{2} - h$ for any $h > 0$. Hence, the bound is sharp.

4. Compatibility between Pareto Frontier and (ϵ, δ) -IF

In this section, we study the case where the optimal fair L^2 learning is not guaranteed to be compatible with the (ϵ, δ) -IF requirement. When the optimal fair L^2 learning cannot be obtained due to the (ϵ, δ) -IF constraint on the admissible maps, the natural partial solution is to study which portion of the Pareto frontier is compatible with the (ϵ, δ) -IF requirement. That is, we study in which cases one has to give up the Pareto optimality between L^2 loss and statistical disparity to satisfy the (ϵ, δ) -IF requirement.

The following result provides a sufficient condition to find the portion of Pareto frontier that is guaranteed to be compatible with the (ϵ, δ) -IF requirement. Let f^* be the solution to Problem 1 as defined in Lemma 2.2, $L(f^*) := \sup_{(y,z)} \|f^*(y, z) - y\|$, V be the independence projection loss defined in equation (11), and f_d be the Pareto optimal solutions for $d \in [0, \infty)$ as defined in Lemma 2.3. Then we have the following result:

Theorem 4.1 (Compatible portion of Pareto frontier with (ϵ, δ) -IF) *Suppose that*

$$\text{Domain}(f^*(\cdot, z)) = \mathcal{Y}, \lambda - a.e. z \in \mathcal{Z},$$

then

$$f_d \text{ is } (\epsilon, \delta)\text{-IF for } \begin{cases} d \in [\sqrt{2}V(1 - \frac{\delta-\epsilon}{2L(f^*)}), \infty), & \text{if } \delta - \epsilon \in [0, 2L(f^*)) \\ d \in [0, \infty), & \text{if } \delta - \epsilon \in [2L(f^*), \infty) \end{cases} \quad (15)$$

for all $\epsilon > 0$ as $L(f_d) = (1 - \frac{d}{\sqrt{2}V})L(f^*) \leq \frac{\delta-\epsilon}{2}$.

In other words, Problem 2 \equiv Problem 4 if

$$\frac{d}{\sqrt{2}V} \in \begin{cases} [1 - \frac{\delta-\epsilon}{2L(f^*)}, \infty), & \text{if } \delta - \epsilon \in [0, 2L(f^*)) \\ [0, \infty), & \text{if } \delta - \epsilon \in [2L(f^*), \infty) \end{cases} \quad (16)$$

Here, the assumption is exactly the same as assumption 2 in Theorem 3.1. The practicality of the assumption is discussed in Remark 3.1. The above result shows that, if $\delta - \epsilon \in [0, 2L(f^*))$ (respectively $\delta - \epsilon \in [2L(f^*), \infty)$), then the Pareto optimal solutions (hence Pareto optimality due to uniqueness of the solutions) are compatible with the (ϵ, δ) -IF constraint when the statistical disparity tolerance level d is in the interval $[\sqrt{2}V(1 - \frac{\delta-\epsilon}{2L(f^*)}), \infty)$ (respectively $[0, \infty)$).

Remark 4.1 (Different Cases of (ϵ, δ) in Theorem 4.1) *Here, we discuss different cases of (ϵ, δ) in the above result:*

I $[\delta - \epsilon < 0]$ The above result provides sufficient conditions for compatibility without any further assumptions on the distribution of \hat{Y}_z . When $\delta - \epsilon < 0$, let $\epsilon = 0$, (ϵ, δ) -IF constraint requires

$$\sup_{z_1, z_2} \|f(y, z_1) - f(y, z_2)\| \leq \delta < 0, \quad (17)$$

which leads to a contradiction. Hence, there is no sufficient condition to guarantee any portion of the Pareto frontier to be compatible with (ϵ, δ) -IF constraint in this case.

- II [$\delta - \epsilon = 0$] In the case of $\delta = \epsilon$, the (ϵ, δ) -IF constraint does not allow any different maps to the same $y \in \mathcal{Y}$ when set $\epsilon = 0$. Therefore, the inherent incompatibility of K -Lipschitz-IF also applies to (ϵ, δ) -IF in this case. One has to set the d bigger than or equal to $\sqrt{2}V$ to tolerate the original prediction \hat{Y} . Any non-trivial trade-off is prohibited.
- III [$\delta - \epsilon \in (0, 2K)$] When $\delta - \epsilon \in (0, 2K)$, the (ϵ, δ) -IF constraint now allows sensitive information dependent maps, but the flexibility is not large enough to tolerate the case where one wants to trade a large amount of utility for statistical parity. Therefore, only a limited portion of the Pareto frontier starting from the original learning outcome is compatible with the (ϵ, δ) -IF constraint in this case.
- IV [$\delta - \epsilon \in [2K, \infty)$] When $\delta - \epsilon \in [2K, \infty)$, the optimal transport maps that result in the optimal statistical parity L^2 learning are compatible with the (ϵ, δ) -IF constraint. Therefore, all the tolerance levels are compatible in this case.

5. Composition Results

Since a post-processed model is a composition of a trained model and a post-processing step, it is important in practice to study the compatibility between the composition and the individual fairness constraints. Therefore, in this section, we provide a compatibility guarantee for the composition or, equivalently, the post-processed model based on the analysis results developed for the post-processing step in the previous sections.

In particular, we assume the trained model, $g : \mathcal{X} \times \mathcal{Z} \rightarrow \mathcal{Y}$ (or $g : \mathcal{X} \rightarrow \mathcal{Y}$ when considering $g(x, z)$ remains constant when z varies), is provided. Here, $(\mathcal{X}, d_{\mathcal{X}})$ is the qualification or dependent variable (metric) space and $(\mathcal{Y}, \|\cdot\|)$ is the independent variable (Euclidean) space. That is, the provided prediction \hat{Y} has the form $\hat{Y} = g(X, Z)$, where X is the qualification or dependent variable and Z is the sensitive variable. Now, by assuming g satisfies uniform (ϵ, δ) -IF:

$$d_{\mathcal{X}}(x_1, x_2) < \epsilon \implies \sup_{(z_1, z_2) \in \mathcal{Z}^2} \|g(x_1, z_1) - g(x_2, z_2)\| < \delta,$$

or uniform K -Lipschitz-IF:

$$\sup_{(z_1, z_2) \in \mathcal{Z}^2} \|g(x_1, z_1) - g(x_2, z_2)\| \leq K d_{\mathcal{X}}(x_1, x_2),$$

we provide compatibility guarantee to the post-processed model $f^* \circ g$ or $f_d \circ g$. Here, f^* and f_d are defined as in Lemma 2.2 and Lemma 2.3. The post-processing composition $f \circ g : \mathcal{X} \times \mathcal{Z} \rightarrow \mathcal{Y}$ is defined by $(x, z) \rightarrow f(g(x, z), z)$ for $f \in \{f^*, f_d\}$, where g is the trained model that use (x, z) or x for prediction and f is the post-processing step that applied z -dependent maps $f(\cdot, z)$ to deform $g(x, z)$ to satisfy statistical parity.

To prove the main compatibility result, we need the following lemma, which shows compatibility guarantee for the composition $(f \circ g(x, z) := f(g(x, z), z))$ of an arbitrary measurable function $g : \mathcal{X} \times \mathcal{Z} \rightarrow \mathcal{Y}$ and an arbitrary measurable function $f : \mathcal{Y} \times \mathcal{Z} \rightarrow \mathcal{Y}$.

Lemma 5.1 (Composition guarantee for (ϵ, δ) -IF functions) *Assume $g : \mathcal{X} \times \mathcal{Z} \rightarrow \mathcal{Y}$ satisfies uniform (ϵ, δ) -IF and define $L(f) := \sup_{(y,z)} \|f(y,z) - y\|$, define the post-processing composition $f \circ g$ by $(x, z) \rightarrow f \circ g(x, z) := f(g(x, z), z)$, then*

$$f \circ g \text{ satisfies uniform } (\epsilon, \delta + 2L(f))\text{-IF, for all } \epsilon > 0. \quad (18)$$

Furthermore, assume $g : \mathcal{X} \rightarrow \mathcal{Y}$ satisfies (ϵ, δ) -IF and define the pre-processing composition $g \circ f$ by $(x, z) \rightarrow g \circ f(x, z) := g(f(x, z))$, then

$$g \circ f \text{ satisfies uniform } (\epsilon - 2L(f), \delta)\text{-IF, for all } \epsilon > 2L(f) \quad (19)$$

Remark 5.1 (Composition analysis for pre-processing) *As mentioned in Remark 1.4, one can apply f^* or f_d as pre-processing steps which now are functions from $\mathcal{X} \times \mathcal{Z}$ to \mathcal{X} . See detailed explanation in [28]. Therefore, despite our focus on the post-processing case in the present work, Lemma 5.1, 5.2 and Theorem 5.1, 5.2 all include results for both post-processing ($f \circ g$) and pre-processing ($g \circ f$). Therefore, the composition results also include the pre-processing cases so that practitioners or researchers who are interested in pre-processing or synthetic data approach to both group fairness and individual fairness can modify the compatibility analysis and apply the respective composition result.*

The above result shows that: (1) For post-processing, if the trained model g satisfies (ϵ, δ) -IF, then it is guaranteed that the post-processed learning outcome $f \circ g$ satisfies $(\epsilon, \delta + 2L(f))$ -IF. (2) For pre-processing, if the trained model g satisfies (ϵ, δ) -IF for some $\epsilon > 2L(f)$, then it is guaranteed that the pre-processed learning outcome $f \circ g$ satisfies $(\epsilon - 2L(f), \delta)$ -IF.

But, in practice, we often fix (ϵ, δ) -IF requirement first and then look for the Pareto optimal solutions to have both individual fairness and diminished statistical disparity at the lowest utility loss. The following result shows which portions of the Pareto frontier (or, equivalently, which Pareto optimal solutions), when composed with a trained model g , result in a post-processed model compatible with the (ϵ, δ) -IF requirement.

Theorem 5.1 (Composition guarantee for (ϵ, δ) -IF trained model) *For any $(\epsilon, \delta) \in (\mathbb{R}^+)^2$, if $g : \mathcal{X} \times \mathcal{Z} \rightarrow \mathcal{Y}$ satisfies uniform (ϵ, δ_g) -IF for some $\delta_g < \delta$, then*

$$\{f_d \circ g\}_{d \in [\sqrt{2}V[1 - (\frac{\delta - \delta_g}{2L(f)} \wedge 1)], \infty)} \text{ satisfy uniform } (\epsilon, \delta)\text{-IF.}$$

Furthermore, if $g : \mathcal{X} \rightarrow \mathcal{Y}$ satisfies (ϵ_g, δ) -IF for some $\epsilon_g > \epsilon$, then

$$\{g \circ f_d\}_{d \in [\sqrt{2}V[1 - (\frac{\epsilon_g - \epsilon}{2L(f)} \wedge 1)], \infty)} \text{ satisfy uniform } (\epsilon, \delta)\text{-IF.}$$

That is, if we require the post-processed learning outcome to satisfy uniform (ϵ, δ) -IF, the above result together with Theorem 5.1 shows that one sufficient approach is to first require the trained model g to satisfy (ϵ, δ_g) -IF for some $\delta_g < \delta$ and then pick

$$d \in [\sqrt{2}V[1 - (\frac{\delta - \delta_g}{2L(f)} \wedge 1)], \infty)$$

so that the composed Pareto optimal solution $f_d \circ g$ satisfies is guaranteed to satisfy (ϵ, δ) -IF. On the other hand, if we require the pre-processed learning outcome to satisfy (ϵ, δ) -IF, the above result together with Theorem shows that one sufficient approach is to first require the trained model g to satisfy (ϵ_g, δ) -IF for some $\epsilon_g > \epsilon$ and then pick

$$d \in [\sqrt{2}V[1 - (\frac{\epsilon_g - \epsilon}{2L(f)} \wedge 1)], \infty)$$

so that the composed Pareto optimal solution $g \circ f_d$ satisfies is guaranteed to satisfy (ϵ, δ) -IF.

Similarly, we have compatibility guarantee when assuming g satisfies K -Lipschitz-IF:

Lemma 5.2 (Composition guarantee for K – Lipschitz-IF functions) *Assume $g : \mathcal{X} \times \mathcal{Z} \rightarrow \mathcal{Y}$ satisfies uniform K -Lipschitz-IF and define $L(f) := \sup_{(y,z)} \|f(y, z) - y\|$, then*

$$f \circ g \text{ satisfies uniform } (\epsilon, K\epsilon + 2L(f))\text{-IF}, \forall \epsilon > 0. \quad (20)$$

Furthermore, assume $g : \mathcal{X} \rightarrow \mathcal{Y}$ satisfies K -Lipschitz-IF, then

$$g \circ f \text{ satisfies uniform } (\epsilon - 2L(f), K\epsilon)\text{-IF}, \forall \epsilon > 2L(f) \quad (21)$$

Therefore, we are able to derive the following theorem to provide compatibility guarantee for the post-processed model as in the (ϵ, δ) -IF assumption case above.

Theorem 5.2 (Composition guarantee for K -Lipschitz-IF trained model) *For any $(\epsilon, \delta) \in (\mathbb{R}^+)^2$, if $g : \mathcal{X} \times \mathcal{Z} \rightarrow \mathcal{Y}$ satisfies uniform K -Lipschitz-IF for some $K \in (0, \frac{\delta}{\epsilon}]$, then*

$$\{f_d \circ g\}_{d \in [\sqrt{2}V[1 - (\frac{\delta - K\epsilon}{2L(f)} \wedge 1)], \infty)} \text{ satisfy uniform } (\epsilon, \delta)\text{-IF}.$$

Furthermore, if $g : \mathcal{X} \rightarrow \mathcal{Y}$ satisfies K -Lipschitz-IF for some $K \in \mathbb{R}^+$, then

$$\{g \circ f_d\}_{d \in [\sqrt{2}V[1 - (\frac{\delta}{2L(f)} - \epsilon \wedge 1)], \infty)} \text{ satisfy uniform } (\epsilon, \delta)\text{-IF}.$$

6. Empirical Study: Fair Supervised Learning

In this section, we provide an experimental study on the compatibility analysis developed in the previous sections on different data sets¹.

6.1 Benchmark Data and Comparison Methods

1 For uni-variate regression test, we implement the compatibility result in Theorem 4.1 via the optimal affine estimation of the Wasserstein barycenter that is proposed in [28]. One advantage of applying affine estimation of the optimal transport maps is

1. The code for the results of our experiments is available online at: https://github.com/xushizhou/compatibility_group_individual_fairness.

an accurate derivation of $L(f_{\text{affine}})$ via the fact that affine transport maps can be decomposed into a rigid translation and a linear map:

$$\begin{aligned} L(f_{\text{affine}}) &:= \sup_{y,z} \|f_{\text{affine}}(y, z) - y\| \\ &\leq \sup_z \|m_z - \bar{m}\| + \left(\sup_{(y,z)} \|f_{\text{linear}}(y, z) - y\|\right) \\ &\leq \sup_z \|m_z - \bar{m}\| + \left(\sup_z \|f_{\text{linear}}(\cdot, z) - \text{Id}\|_{op}\right) \left(\sup_{y \in \mathcal{Y}} \|y\|\right) \end{aligned}$$

where $m_z := \int_{\mathcal{Y}} y d\mu_z$, $\bar{\mu} := \int_{\mathcal{Y}} y d\bar{\mu}$, $\|\cdot\|_{op}$ denotes the operator norm, $\sup_z \|m_z - \bar{m}\|$ is the worst-case rigid translation, and $\left(\sup_z \|f_{\text{linear}}(\cdot, z) - \text{Id}\|_{op}\right) \left(\sup_{y \in \text{Domain}(\bar{Y})} \|y\|\right)$ is the worst-case linear deformation. In the experiments, we use empirical mean to estimate $\sup_z \|m_z - \bar{m}\|$, the largest absolute value of eigenvalue for operator norm, and the largest $\|y\|$ in the sample to estimate $\sup_{y \in \mathcal{Y}} \|y\|$. In practice, one should derive $\sup_{y \in \mathcal{Y}} \|y\|$ based on the specific application context.

In order to show the estimation accuracy of the optimal affine maps (“supervised learning name + post-proc. Pareto frontier Est. or Pseudo-barycenter”), we also include the exact cumulative distribution function matching method in [7] (“supervised learning name + chzhen”).

- 2 For the multi-variate supervised learning test we also implement the optimal affine method to estimate the post-processing Wasserstein barycenter, estimate the K via the above upper bound, and finally provide the compatibility portion of the Pareto frontier under different (ϵ, δ) -IF requirements.

Data set	Tests	Data size	dim(X)	dim(Y)
LSAC	linear regression, ANN	20454	9	1
CRIME	linear regression, ANN	1994	97	1
CRIME	linear regression, ANN	1994	87	11

- The CRIME dataset contains social, economic, law enforcement, and judicial data for U.S. communities in 1994 (1994 examples) [24].

In univariate L^2 learning, the goal is to predict the number of crimes per 10^5 population using the remaining dataset information. The sensitive variable is race, specifically whether the percentage of African American population exceeds 30

In multivariate supervised learning on the CRIME dataset, we retain the same sensitive variable (race). However, the learning task is to predict a vector representing local housing and rental market data, including low quartile occupied home value, median home value, high quartile home value, low quartile rent, median rent, high quartile rent, median gross rent, number of immigrants, median number of bedrooms, number of vacant households, and the number of crimes.

- The LSAC dataset comprises social, economic, and personal data of 20,454 law school students [27].
- In univariate modeling, the objective is to predict students' GPA using the remaining dataset variables. In this context, race serves as the sensitive variable on whether the student is non-white.

6.2 Numerical Result

The univariate supervised learning test results for the LSAC and CRIME datasets are presented in Figure 1 and Figure 2, respectively.

In these plots, the vertical axis represents the L^2 loss, while the horizontal axis represents the Wasserstein (\mathcal{W}_2) disparity. Achieving a result closer to the lower-left corner indicates better performance. In the top row, one-third of the Pareto frontier is guaranteed to be compatible with the shown (ϵ, δ) -IF. In the middle and bottom rows, one-half and two-thirds of the frontier, respectively, are guaranteed to be compatible with the corresponding (ϵ, δ) -IF.

We adopt two supervised learning methods: linear regression and artificial neural networks (ANN) with four stacked layers. Each of the first three layers contains of 32 units with Rectified Linear Unit (ReLU) activation, and the final layer consists of one unit with linear activation.

Furthermore, since linear regression satisfies Lipschitz-IF constraint, we can apply the composition result (Theorem 5.2) to provide individual fairness guarantee for the composed learning outcome.

In Figure 1, the top, middle, and bottom row shows the compatible portion of the Pareto frontier corresponding to $(\epsilon, \epsilon + \frac{2}{3}L(f^*))$, $(\epsilon, \epsilon + \frac{1}{2}L(f^*))$, $(\epsilon, \epsilon + \frac{4}{3}L(f^*))$ individual fairness constraint, respectively. Here, $L(f^*) = 0.959$ for linear regression prediction and $L(f^*) = 1.250$ for ANN prediction. The resulting compatible portion is the first $\frac{1}{3}$, $\frac{1}{2}$, $\frac{2}{3}$ part of the Pareto frontier. In general, Each percentage increase in $\frac{\delta - \epsilon}{2L(f^*)}$ results in one percentage larger portion of the Pareto frontier to be compatible, until achieving the end of the frontier which is the (estimation of the) optimal post-processing statistical parity L^2 learning. Also, notice that we use $(\epsilon, \epsilon + (\delta - \epsilon))$ in the plots because the compatibility is guaranteed for all $\epsilon \in [0, \infty)$.

For the composition, it follows from the K-Lipschitz condition for $K = 0.1265$ of the linear regression model and the composition theorem, we are able to conclude that the post-processed linear regression model corresponding to the top, middle, and bottom post-processing step satisfies $(\epsilon, 0.1265\epsilon + \frac{2}{3}L(f^*))$, $(\epsilon, 0.1265\epsilon + \frac{1}{2}L(f^*))$, $(\epsilon, 0.1265\epsilon + \frac{4}{3}L(f^*))$ individual fairness constraint respectively.

Moreover, if the goal is to make the post-processed linear regression model satisfy a fixed (ϵ, δ) -IF, then it follows from Theorem 5.2 that the Pareto optimal solutions f_d are guaranteed to satisfy the required (ϵ, δ) -IF constraint for all $d \in [\sqrt{2}V[1 - (\frac{\delta}{0.1265} - \frac{\epsilon}{2L(f^*)} \wedge 1)], \infty)$ where $V = 0.272$ and 0.287 for linear regression and ANN respectively.

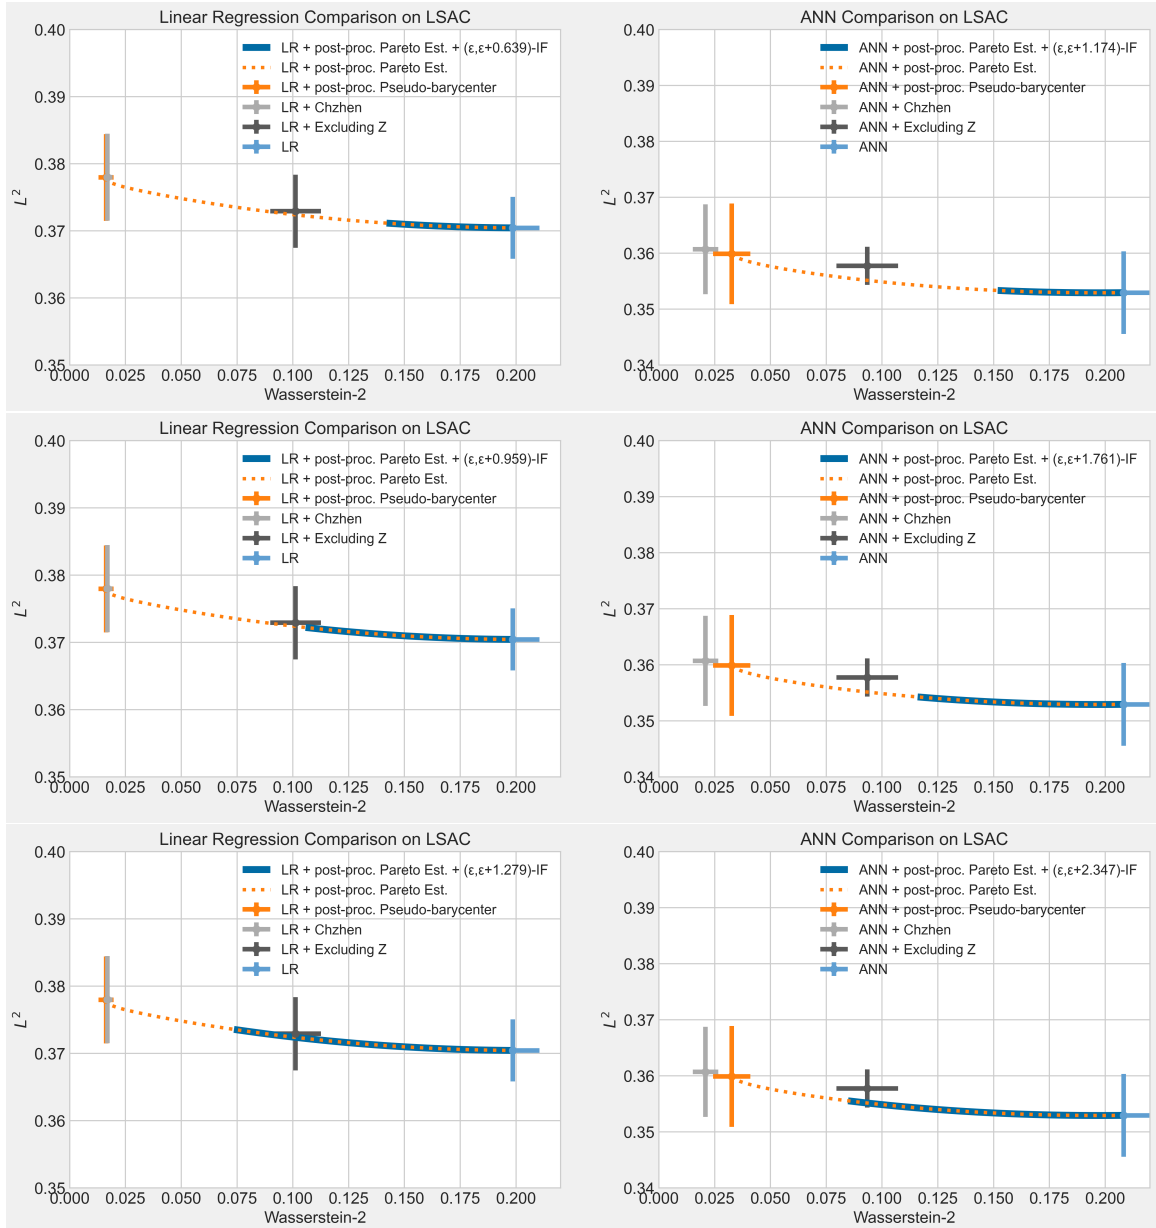


Figure 1: As shown in the univariate regression test on LSAC above, all three rows consist of the L^2 loss and Wasserstein disparity of the original prediction (LR or ANN), the prediction using data excluding Z (LR or ANN + Excluding Z), the exact post-processing \mathcal{W}_2 barycenter via cumulative distribution functions (cdf) matching approach (LR or ANN + chzhen2020fair), the optimal affine estimation of the post-processing \mathcal{W}_2 barycenter (LR or ANN + post-proc. Pseudo-barycenter), the Pareto frontier estimated by the optimal affine maps (LR or ANN + post-proc. Pareto Est.), and finally the portion of the estimated Pareto frontier that is compatible with the corresponding (ϵ, δ) -IF constraints. Here, $L(f^*) = 0.959$ for linear regression prediction and $L(f^*) = 1.250$ for ANN prediction. For each (ϵ, δ) -IF constraint, the compatible portion is the first $\frac{\delta - \epsilon}{2L(f^*)}$ part of the Pareto frontier. More generally, each percentage increase in $\frac{\delta - \epsilon}{2L(f^*)}$ results in one percentage larger portion of the Pareto frontier to be compatible. Also, the portion is guaranteed to satisfy $(\epsilon, \epsilon + (\delta - \epsilon))$ -IF for all $\epsilon \in [0, \infty)$

In Figure 2, the compatibility result for the CRIME experiment is shown analogously as in Figure 1, except now $L(f^*) = 1.045$ and $V = 0.382$ for linear regression, $L(f^*) = 1.385$

and $V = 0.389$ for ANN prediction, and the linear model satisfies $K = 1.030$ -Lipschitz condition.

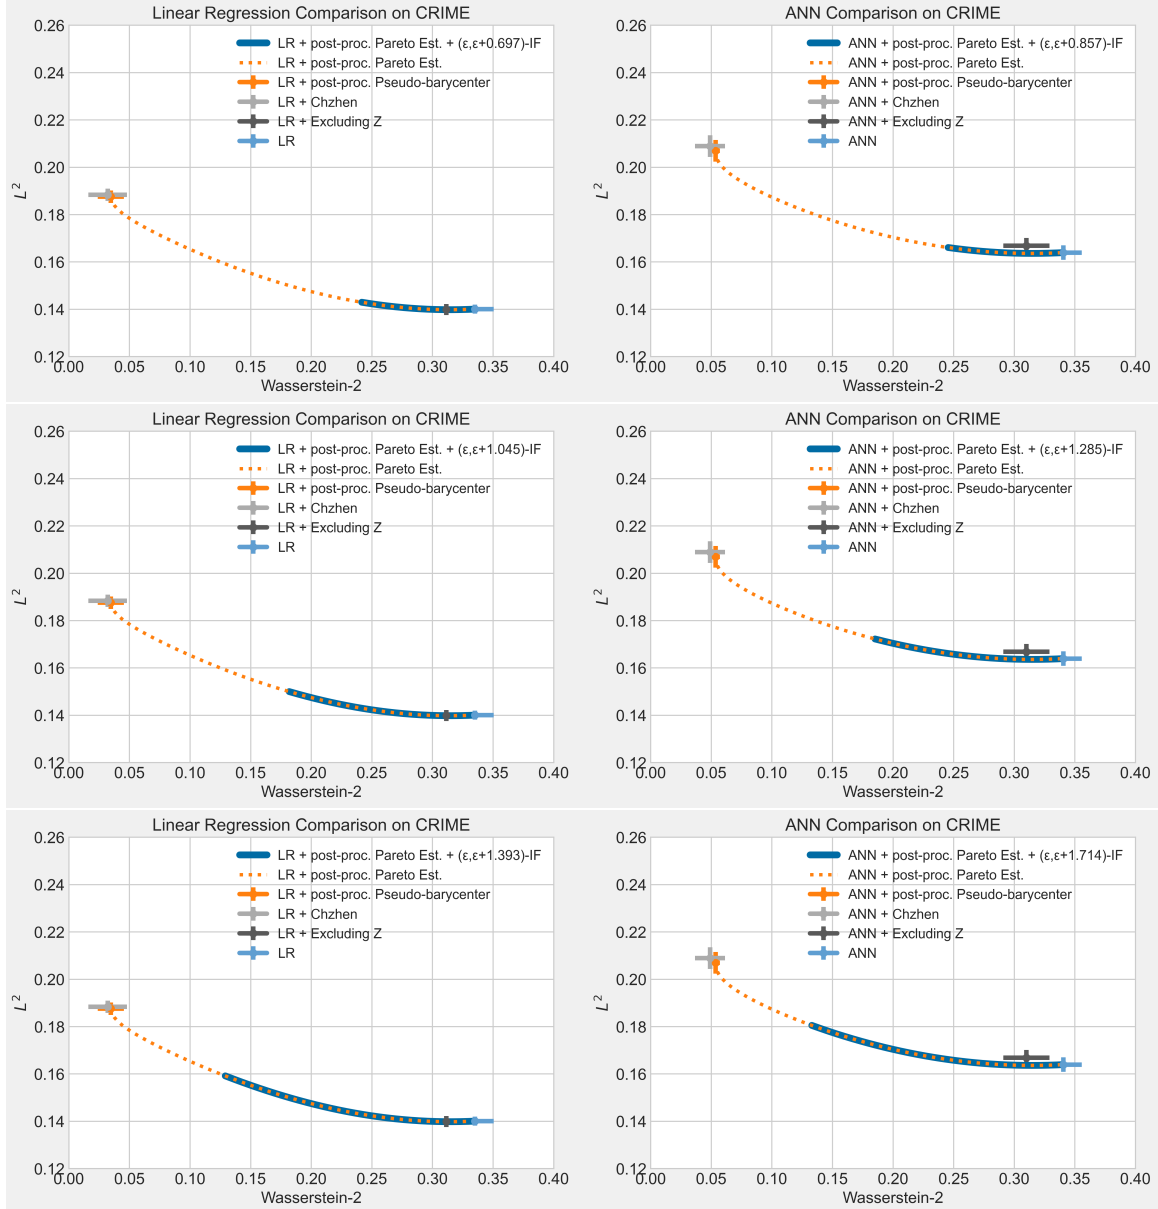


Figure 2: As shown in the univariate regression test on CRIME above, all three rows consist of the L^2 loss and Wasserstein disparity of the original prediction (LR or ANN), the prediction using data excluding Z (LR or ANN + Excluding Z), the exact post-processing \mathcal{W}_2 barycenter via cdfs matching approach (LR or ANN + chzhen2020fair), the optimal affine estimation of the post-processing \mathcal{W}_2 barycenter (LR or ANN + post-proc. Pseudo-barycenter), the Pareto frontier estimated by the optimal affine maps (LR or ANN + post-proc. Pareto Est.), and finally the portion of the estimated Pareto frontier that is compatible with the corresponding (ϵ, δ) -IF constraints. Here, $L(f^*) = 1.045$ for linear regression and $L(f^*) = 1.385$ for ANN prediction. Each percentage increase in $\frac{\delta - \epsilon}{2L(f^*)}$ results in one percentage larger portion of the Pareto frontier to be compatible.

The multivariate supervised learning test results are shown in Figure 3. The compatibility result can be concluded analogous to the one for the LSAC test. Except now

$L(f^*) = 3.396$ and $V = 0.527$ for linear regression, which satisfies 1.457-Lipschitz condition, and $L(f^*) = 4.434$ and $V = 0.545$ for ANN.

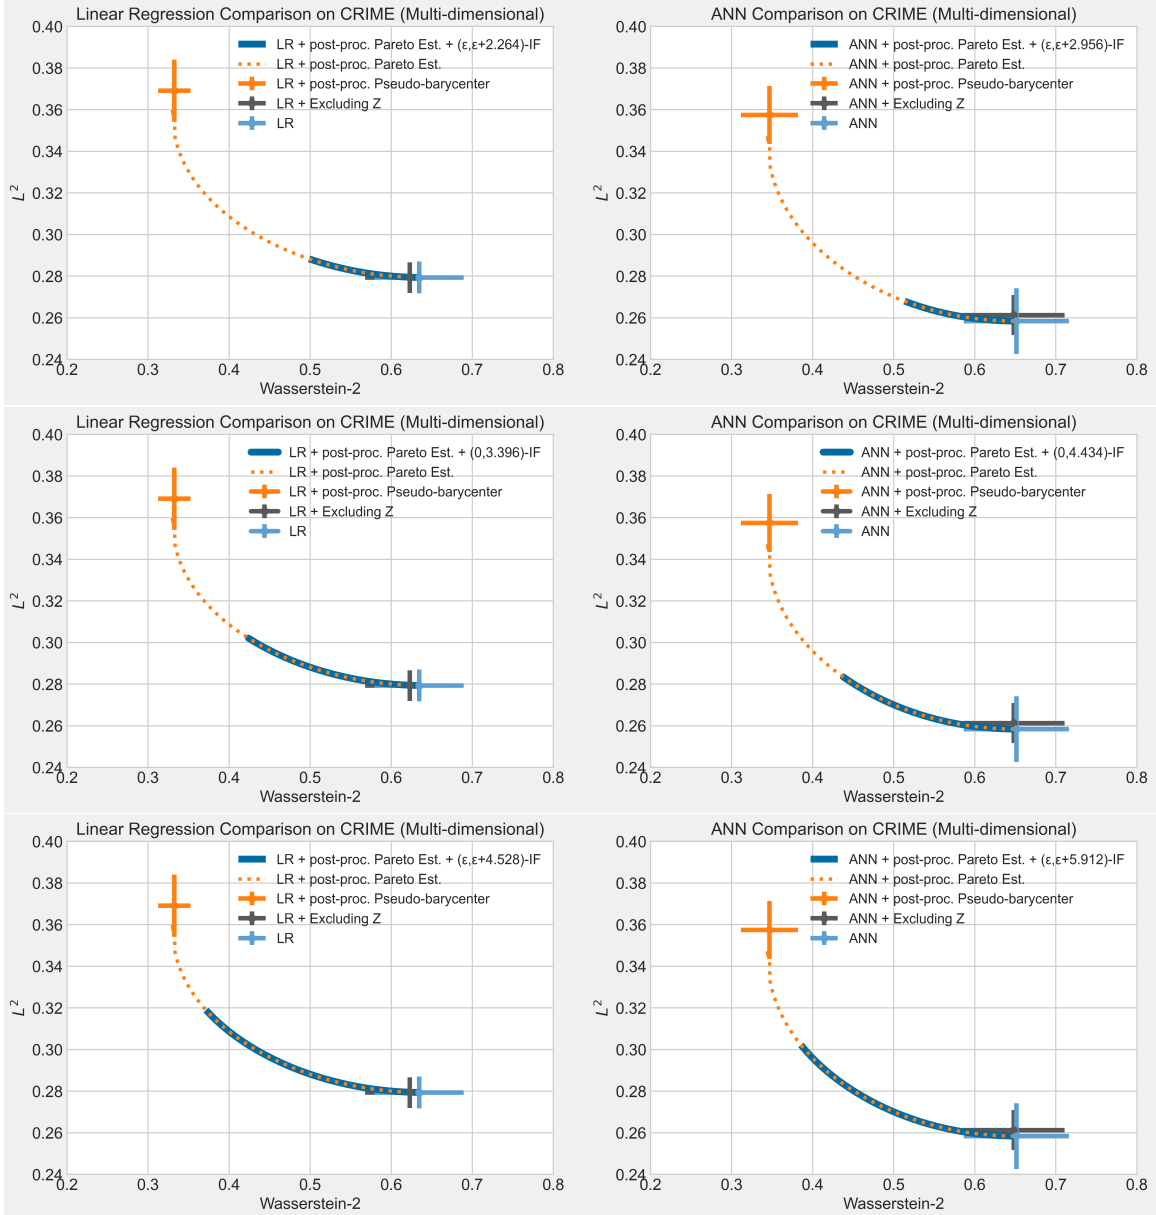


Figure 3: In the multivariate regression test on CRIME above, all three rows consist of the L^2 loss and Wasserstein disparity of the original prediction (LR or ANN), the prediction using data excluding Z (LR or ANN + Excluding Z), the optimal affine estimation of the post-processing W_2 barycenter (LR or ANN + post-proc. Pseudo-barycenter), the Pareto frontier estimated by the optimal affine maps (LR or ANN + post-proc. Pareto Est.), and finally the portion of the estimated Pareto frontier that is compatible with the corresponding (ϵ, δ) -IF constraints. Notice that excluding Z now removes only limited Wasserstein disparity due to the multidimensional dependent variable. Here, $L(f^*) = 3.396$ for linear regression and $L(f^*) = 4.434$ for ANN prediction. Each percentage increase in $\frac{\delta - \epsilon}{2L(f^*)}$ results in one percentage larger portion of the Pareto frontier to be compatible.

Acknowledgement

The authors acknowledge support from NSF DMS-2027248, NSF DMS-2208356, NSF CCF-1934568, and NIH R01HL16351.

A. Appendix: Proof of Results in Section 2

A.1 Proof of Lemma 2.1

Proof [$D(\hat{Y}, Z) = 0 \iff \mathcal{W}_2^2(\mathcal{L}(\hat{Y}_z), \bar{\mu}) = 0, \lambda\text{-a.e.} \iff \hat{Y} \perp Z$.] We first show that (1) $\mathcal{W}_2^2(\mathcal{L}(\hat{Y}_z), \bar{\mu}) = 0, \lambda\text{-a.e.} \iff \hat{Y} \perp Z$, then complete the proof by showing (2) $D(\hat{Y}, Z) = 0 \iff \mathcal{W}_2^2(\mathcal{L}(\hat{Y}_z), \bar{\mu}) = 0, \lambda\text{-a.e.}$. Here, $\bar{\mu}$ denotes the Wasserstein barycenter of the marginal distributions $\{\mathcal{L}(\hat{Y}_z)\}_z$.

(1) Assume $\hat{Y} \perp Z$, we have $\mathcal{L}(\hat{Y}_z) = \mathcal{L}(\hat{Y})$ $\lambda\text{-a.e.}$ which implies $\mathcal{W}_2^2(\mathcal{L}(\hat{Y}_z), \mathcal{L}(\hat{Y})) = 0, \lambda\text{-a.e.}$. Then, it follows from the uniqueness of $\bar{\mu}$ that $\bar{\mu} = \mathcal{L}(\hat{Y})$ and $\mathcal{W}_2^2(\mathcal{L}(\hat{Y}_z), \bar{\mu}) = 0, \lambda\text{-a.e.}$. For the other direction, assume that $\mathcal{W}_2^2(\mathcal{L}(\hat{Y}_z), \bar{\mu}) = 0$ $\lambda\text{-a.e.}$, then for all $A \in \sigma(\hat{Y})$ and $B \in \sigma(Z)$, we have

$$\begin{aligned} \mathbb{P}(\{\hat{Y} \in A\} \cap \{Z \in B\}) &= \int_B \mathbb{P}(\{\hat{Y}_z \in A\}) d\mathbb{P} \circ Z^{-1} \\ &= \int_B \bar{\mu}(A) d\mathbb{P} \circ Z^{-1} \\ &= \bar{\mu}(A) \int_B d\mathbb{P} \circ Z^{-1} \\ &= \mathbb{P}(\{\hat{Y} \in A\}) \mathbb{P}(\{Z \in B\}). \end{aligned}$$

This proves $\hat{Y} \perp Z$ and completes the proof for the first claim.

(2) Assume that \hat{Y} is not independent of Z , then $\mathcal{W}_2^2(\mathcal{L}(\hat{Y}_z), \bar{\mu}) \neq 0, \lambda\text{-a.e.}$. This implies

$$\begin{aligned} &\lambda(\{\mathcal{W}_2^2(\mathcal{L}(\hat{Y}_z), \bar{\mu}) > 0\}) > 0 \\ \iff &\lambda\left(\bigcup_{k=1}^{\infty} \{\mathcal{W}_2^2(\mathcal{L}(\hat{Y}_z), \bar{\mu}) \geq \frac{1}{k}\}\right) > 0 \\ \implies &\exists K < \infty \text{ such that } \lambda(\{\mathcal{W}_2^2(\mathcal{L}(\hat{Y}_z), \bar{\mu}) \geq \frac{1}{K}\}) > 0 \\ \implies &\int_{\mathcal{Z}} \mathcal{W}_2^2(\mathcal{L}(\hat{Y}_z), \bar{\mu}) d\lambda \geq \frac{1}{K} \lambda(\{\mathcal{W}_2^2(\mathcal{L}(\hat{Y}_z), \bar{\mu}) \geq \frac{1}{K}\}) > 0. \end{aligned}$$

For the other direction, assume that $\hat{Y} \perp Z$, then it follows from the first claim that $\mathcal{W}_2^2(\mathcal{L}(\hat{Y}_z), \bar{\mu}) = 0, \lambda\text{-a.e.}$ which further implies $\int_{\mathcal{Z}} \mathcal{W}_2^2(\mathcal{L}(\hat{Y}_z), \bar{\mu}) d\lambda = 0$. This proves the second claim. \blacksquare

B. Appendix: Proof of Results in Section 3

B.1 Proof of Theorem 3.1

Proof (Optimal L^2 learning) It follows from the assumption $\hat{Y} \not\perp Z$ that there exists an event $A \in \mathcal{B}_{\mathcal{Z}}$ such that $\lambda(A) > 0$ and $\mathcal{W}_2^2(Y_{z_1}, Y_{z_2}) \neq 0$ for some $z_1, z_2 \in A$. By Lemma

2.2, the optimal fair L^2 learning requires the set of optimal transport maps $\{f(\cdot, z)\}_z$ to map each sensitive marginal learning outcome to their Wasserstein barycenter:

$$f(Y_z, z) = \bar{Y}, \text{ for } \lambda - a.e. z \in \mathcal{Z}. \quad (22)$$

Since $\mathcal{W}_2^2(Y_{z_1}, Y_{z_2}) \neq 0$, we have

$$\max \{\mathcal{W}_2^2(Y_{z_1}, \bar{Y}), \mathcal{W}_2^2(Y_{z_2}, \bar{Y})\} > 0,$$

which further implies $f(\cdot, z_1) \neq f(\cdot, z_2)$. Now, it follows from the second assumption that there exists $y \in \mathcal{Y} = \text{Domain}(f(\cdot, z_1)) \cap \text{Domain}(f(\cdot, z_2))$ such that

$$\|f(y, z_1) - f(y, z_2)\|_2 > 0 = K\|y - y\|_2,$$

for any $K > 0$. That contradicts the Lipschitz individual fairness definition and hence completes the proof for the incompatibility between the optimal L^2 learning and the K-Lipschitz-IF.

(Pareto optimal solutions) By Lemma 2.3, the Pareto optimal solutions are achieved by the McCann interpolations: $f(1 - \frac{d}{\sqrt{2V}})(\cdot, z)$. But it follows from the assumptions that there exists $y \in \mathcal{Y} = \text{Domain}(f(\cdot, z_1)) \cap \text{Domain}(f(\cdot, z_2))$ such that $\|f(y, z_1) - f(y, z_2)\|_2 > 0$, which further implies that $\|f(1 - \frac{d}{\sqrt{2V}})(y, z_1) - f(1 - \frac{d}{\sqrt{2V}})(y, z_2)\|_2 > 0$ for any $d \in [0, \sqrt{2V})$. That is, $f(1 - \frac{d}{\sqrt{2V}})$ satisfies the K-Lipschitz-IF constraint only if $d \in [\sqrt{2V}, \infty)$. But the Pareto optimal solution set becomes $\{\hat{Y}\}$ for $d \in [\sqrt{2V}, \infty)$. That completes the proof. ■

C. Appendix: Proof of Results in Section 4

C.1 Proof of Theorem 4.1

Proof [$\delta - \epsilon \in [0, 2K]$] It suffices to show that the optimal transport map $\{f(t)(\cdot, z)\}_z := \{(1-t)\text{Id} + tf(\cdot, z)\}_z$ is a subset of $\mathcal{D}_{(\epsilon, \delta)\text{-IF}}$ given $t = 1 - \frac{d}{\sqrt{2V}}$. Indeed, when $\delta - \epsilon \in [0, 2K)$, we have

$$\begin{aligned} \frac{d}{\sqrt{2V}} \in [1 - \frac{\delta - \epsilon}{2K}, \infty) &\iff 1 - \frac{d}{\sqrt{2V}} \leq \frac{\delta - \epsilon}{2K} \\ &\implies (1 - \frac{d}{\sqrt{2V}}) \sup_{y,z} \|f(y, z) - y\| \leq \frac{\delta - \epsilon}{2} \\ &\implies \sup_{y,z} \|(1 - \frac{d}{\sqrt{2V}})f(y, z) + \frac{d}{\sqrt{2V}}y - y\| \leq \frac{\delta - \epsilon}{2} \\ &\iff \sup_{y,z} \|f(1 - \frac{d}{\sqrt{2V}})(y, z) - y\| \leq \frac{\delta - \epsilon}{2} \end{aligned}$$

Here, the last line is due to McCann interpolation: $(1 - \frac{d}{\sqrt{2V}})f(y, z) + \frac{d}{\sqrt{2V}}y = f(1 - \frac{d}{\sqrt{2V}})(y, z)$. Now, by Lemma 2.3, $\{f(1 - \frac{d}{\sqrt{2V}})(\cdot, z)\}_z$ are the Pareto optimal solutions at the tolerance level d , it follows from the proof of Lemma 3.1 that the corresponding Optimal

solutions satisfy the (ϵ, δ) -IF constraint for all d satisfies $\frac{d}{\sqrt{2V}} \in [1 - \frac{\delta - \epsilon}{2K}, \infty)$. That completes the proof for the case of $\delta - \epsilon \in [0, 2K)$.

$[\delta - \epsilon \in [2K, \infty)]$ When $\delta - \epsilon \in [2K, \infty)$, we have

$$K \leq \frac{\delta - \epsilon}{2}.$$

It again follows from the proof of Lemma 3.1 that $\{f(1 - \frac{d}{\sqrt{2V}})(\cdot, z)\}_z$ satisfies the (ϵ, δ) -IF constraint for all $d \in [0, \infty)$. We are done. \blacksquare

D. Appendix: Proof of Results in Section 5

D.1 Proof of Lemma 5.1

Proof ($f \circ g$) For the post-processing case, let $(\mathcal{X}, d_{\mathcal{X}})$ be a metric space, $(\mathcal{Y}, \|\cdot\|)$ be a Euclidean space. Now, choose $x_1, x_2 \in \mathcal{X}$ to satisfy $d_{\mathcal{X}}(x_1, x_2) < \epsilon$, it follows from the assumption of g that $\|g(x_1, z_1) - g(x_2, z_2)\| < \delta$. But we also have

$$\begin{aligned} & \|f(g(x_1, z_1), z_1) - f(g(x_2, z_2), z_2)\| \\ & \leq \|f(g(x_1, z_1), z_1) - g(x_2, z_2)\| + \|g(x_1, z_1) - g(x_2, z_2)\| + \|g(x_2, z_2) - f(g(x_2, z_2), z_2)\| \\ & < L + \delta + L = \delta + 2L(f) \end{aligned}$$

Since our choice of $x_1, x_2 \in \mathcal{X}$ satisfying $d_{\mathcal{X}}(x_1, x_2) < \epsilon$ is arbitrary, we have $f \circ g$ is $(\epsilon, \delta + 2L(f))$ -IF by definition. Finally, since our choice of $\epsilon > 0$ is arbitrary, it follows that $f \circ g$ is $(\epsilon, \delta + 2L(f))$ -IF, $\forall \epsilon > 0$.

$(g \circ f)$ For the pre-processing case, let $(\mathcal{X}, \|\cdot\|)$ be a Euclidean space, and let $(\mathcal{Y}, d_{\mathcal{Y}})$ be a metric space. Now, choose $x_1, x_2 \in \mathcal{X}$ to satisfy $\|x_1 - x_2\| < \epsilon - 2L(f)$ for some $\epsilon > 2L(f)$, then we have

$$\begin{aligned} & \|f(x_1, z_1) - f(x_2, z_2)\| \\ & \leq \|f(x_1, z_1) - x_1\| + \|x_1 - x_2\| + \|x_2 - f(x_2, z_2)\| \\ & < L(f) + \epsilon - 2L(f) + L(f) = \epsilon \end{aligned}$$

It follows from the assumption of g that

$$d_{\mathcal{Y}}(g(f(x_1, z_1)) - g(f(x_2, z_2))) < \delta.$$

Since our choice of $x_1, x_2 \in \mathcal{X}$ satisfying $\|x_1 - x_2\| < \epsilon - 2L(f)$ is arbitrary, that proves $f \circ g$ is $(\epsilon - 2L(f), \delta)$ -IF by definition. Finally, since our choice of $\epsilon > 2L(f)$ is arbitrary, it follows that $f \circ g$ is $(\epsilon - 2L(f), \delta)$ -IF, $\forall \epsilon > 2L(f)$. \blacksquare

D.2 Proof of Lemma 5.2

Proof The proof is similar to the proof of Lemma 5.1 above and left to the reader. \blacksquare

D.3 Proof of Theorem 5.1

Proof [Post-processing $f_d \circ g$] Let $(x_1, x_2) \in \mathcal{X} \times \mathcal{X}$ satisfy $d_{\mathcal{X}}(x_1, x_2) < \epsilon$. It follows from the assumption of (ϵ, δ_g) -IF of g and the triangle inequality that

$$\begin{aligned} & \|f_d \circ g(x_1, z_1) - f_d \circ g(x_2, z_2)\| \\ & \leq \|f_d \circ g(x_1, z_1) - g(x_1, z_1)\| + \|g(x_1, z_1) - g(x_2, z_2)\| + \|g(x_2, z_2) - f_d \circ g(x_2, z_2)\| \\ & < 2L(f_d) + \delta_g \end{aligned}$$

Now, if $\frac{\delta - \delta_g}{2L(f^*)} \geq 1$, then $2L(f_d) + \delta_g \leq 2L(f^*) + \delta_g \leq \delta, \forall d \in [0, \infty)$. This implies that $\|f_d \circ g(x_1, z_1) - f_d \circ g(x_2, z_2)\| < \delta$ for all $d \in [0, \infty)$. On the other hand, if $\frac{\delta - \delta_g}{2L(f^*)} < 1$, then for all $d \in [\sqrt{2V}(1 - \frac{\delta - \delta_g}{2L(f^*)}), \infty)$ we have

$$\begin{aligned} \frac{d}{\sqrt{2V}} \geq 1 - \frac{\delta - \delta_g}{2L(f^*)} & \implies 1 - \frac{d}{\sqrt{2V}} \leq \frac{\delta - \delta_g}{2L(f^*)} \\ & \implies 2L(f^*)(1 - \frac{d}{\sqrt{2V}}) \leq \delta - \delta_g \\ & \implies 2L(f_d) + \delta_g \leq \delta \\ & \implies \|f_d \circ g(x_1, z_1) - f_d \circ g(x_2, z_2)\| < \delta. \end{aligned}$$

Here, the second last line follows from $(1 - \frac{d}{\sqrt{2V}})L(f^*) = L(f_d)$.

[Pre-processing $g \circ f_d$] Let $(x_1, x_2) \in \mathcal{X} \times \mathcal{X}$ satisfy $d_{\mathcal{X}}(x_1, x_2) < \epsilon$. Now, if $\frac{\epsilon_g - \epsilon}{2L(f^*)} \geq 1$, then

$$2L(f_d) + \epsilon \leq 2L(f^*) + \epsilon \leq \epsilon_g, \forall d \in [0, \infty).$$

Hence, $\|f_d(x_1) - f_d(x_2)\| < \epsilon + 2L(f_d) \leq \epsilon_g$ and it follows from the assumption of (ϵ_g, δ) -IF of g that $d(g \circ f_d(x_1), g \circ f_d(x_2)) < \delta$. On the other hand, if $\frac{\epsilon_g - \epsilon}{2L(f^*)} < 1$, then for all $d \in [\sqrt{2V}(1 - \frac{\epsilon_g - \epsilon}{2L(f^*)}), \infty)$, we have

$$\begin{aligned} \frac{d}{\sqrt{2V}} \geq 1 - \frac{\epsilon_g - \epsilon}{2L(f^*)} & \implies 1 - \frac{d}{\sqrt{2V}} \leq \frac{\epsilon_g - \epsilon}{2L(f^*)} \\ & \implies 2L(f^*)(1 - \frac{d}{\sqrt{2V}}) \leq \epsilon_g - \epsilon \\ & \implies 2L(f_d) + \epsilon \leq \epsilon_g \end{aligned}$$

Hence, $\|f_d(x_1) - f_d(x_2)\| < \epsilon + 2L(f_d) \leq \epsilon_g$ and it follows from the assumption of (ϵ_g, δ) -IF of g that $d(g \circ f_d(x_1), g \circ f_d(x_2)) < \delta$. This completes the proof. ■

D.4 Proof of Theorem 5.2

Proof The proof is similar to the proof of Theorem 5.1 above and left to the reader. ■

References

- [1] M. Agueh and G. Carlier. Barycenters in the Wasserstein space. *SIAM Journal on Mathematical Analysis*, 43(2):904–924, 2011.
- [2] R. Berk, H. Heidari, S. Jabbari, M. Joseph, M. Kearns, J. Morgenstern, S. Neel, and A. Roth. A convex framework for fair regression. *arXiv preprint arXiv:1706.02409*, 2017.
- [3] R. Binns. On the apparent conflict between individual and group fairness. In *Proceedings of the 2020 conference on fairness, accountability, and transparency*, pages 514–524, 2020.
- [4] Y. Brenier. Polar factorization and monotone rearrangement of vector-valued functions. *Communications on pure and applied mathematics*, 44(4):375–417, 1991.
- [5] F. Calmon, D. Wei, B. Vinzamuri, K. Natesan Ramamurthy, and K. R. Varshney. Optimized pre-processing for discrimination prevention. *Advances in neural information processing systems*, 30, 2017.
- [6] A. Chouldechova and A. Roth. The frontiers of fairness in machine learning. *arXiv preprint arXiv:1810.08810*, 2018.
- [7] E. Chzhen, C. Denis, M. Hebiri, L. Oneto, and M. Pontil. Fair regression with Wasserstein barycenters. *Advances in Neural Information Processing Systems*, 33:7321–7331, 2020.
- [8] S. Corbett-Davies and S. Goel. The measure and mismeasure of fairness: A critical review of fair machine learning. *arXiv preprint arXiv:1808.00023*, 2018.
- [9] C. Dwork, M. Hardt, T. Pitassi, O. Reingold, and R. Zemel. Fairness through awareness. In *Proceedings of the 3rd innovations in theoretical computer science conference*, pages 214–226, 2012.
- [10] W. Fleisher. What’s fair about individual fairness? In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*, pages 480–490, 2021.
- [11] S. A. Friedler, C. Scheidegger, and S. Venkatasubramanian. The (im) possibility of fairness: Different value systems require different mechanisms for fair decision making. *Communications of the ACM*, 64(4):136–143, 2021.
- [12] W. Gangbo and A. Swiech. Optimal maps for the multidimensional monge-kantorovich problem. *Communications on Pure and Applied Mathematics: A Journal Issued by the Courant Institute of Mathematical Sciences*, 51(1):23–45, 1998.
- [13] T. L. Gouic, J.-M. Loubes, and P. Rigollet. Projection to fairness in statistical learning. *arXiv preprint arXiv:2005.11720*, 2020.
- [14] M. Hardt, E. Price, and N. Srebro. Equality of opportunity in supervised learning. *Advances in neural information processing systems*, 29, 2016.

- [15] U. Hébert-Johnson, M. Kim, O. Reingold, and G. Rothblum. Multicalibration: Calibration for the (computationally-identifiable) masses. In *International Conference on Machine Learning*, pages 1939–1948. PMLR, 2018.
- [16] R. Jiang, A. Pacchiano, T. Stepleton, H. Jiang, and S. Chiappa. Wasserstein fair classification. In *Uncertainty in artificial intelligence*, pages 862–872. PMLR, 2020.
- [17] M. Joseph, M. Kearns, J. H. Morgenstern, and A. Roth. Fairness in learning: Classic and contextual bandits. *Advances in neural information processing systems*, 29, 2016.
- [18] F. Kamiran and T. Calders. Data preprocessing techniques for classification without discrimination. *Knowledge and information systems*, 33(1):1–33, 2012.
- [19] M. Kearns, S. Neel, A. Roth, and Z. S. Wu. Preventing fairness gerrymandering: Auditing and learning for subgroup fairness. In *International conference on machine learning*, pages 2564–2572. PMLR, 2018.
- [20] T. Le Gouic and J.-M. Loubes. Existence and consistency of Wasserstein barycenters. *Probability Theory and Related Fields*, 168:901–917, 2017.
- [21] P. K. Lohia, K. N. Ramamurthy, M. Bhide, D. Saha, K. R. Varshney, and R. Puri. Bias mitigation post-processing for individual and group fairness. In *Icassp 2019-2019 ieee international conference on acoustics, speech and signal processing (icassp)*, pages 2847–2851. IEEE, 2019.
- [22] R. J. McCann. A convexity principle for interacting gases. *Advances in mathematics*, 128(1):153–179, 1997.
- [23] S. Park, C. Yun, J. Lee, and J. Shin. Minimum width for universal approximation. *arXiv preprint arXiv:2006.08859*, 2020.
- [24] M. Redmond and A. Baveja. A data-driven software tool for enabling cooperative information sharing among police departments. *European Journal of Operational Research*, 141(3):660–678, 2002.
- [25] F. Santambrogio. Optimal transport for applied mathematicians. *Birkäuser, NY*, 55 (58-63):94, 2015.
- [26] C. Silvia, J. Ray, S. Tom, P. Aldo, J. Heinrich, and A. John. A general approach to fairness with optimal transport. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34(04), pages 3633–3640, 2020.
- [27] L. F. Wightman. LSAC National Longitudinal Bar Passage Study. LSAC Research Report Series. 1998.
- [28] S. Xu and T. Strohmer. Fair data representation for machine learning at the Pareto frontier. *Journal of Machine Learning Research*, 24(331):1–63, 2023. URL <http://jmlr.org/papers/v24/22-0005.html>.

- [29] M. B. Zafar, I. Valera, M. Gomez Rodriguez, and K. P. Gummadi. Fairness beyond disparate treatment & disparate impact: Learning classification without disparate mistreatment. In *Proceedings of the 26th international conference on world wide web*, pages 1171–1180, 2017.
- [30] R. Zemel, Y. Wu, K. Swersky, T. Pitassi, and C. Dwork. Learning fair representations. In *International conference on machine learning*, pages 325–333. PMLR, 2013.
- [31] H. Zhao and G. J. Gordon. Inherent tradeoffs in learning fair representations. *The Journal of Machine Learning Research*, 23(1):2527–2552, 2022.
- [32] W. Zhou. *Group vs. individual algorithmic fairness*. PhD thesis, University of Southampton, 2022.