

UC Davis

UC Davis Electronic Theses and Dissertations

Title

Fairness in Machine Learning via Optimal Transport

Permalink

<https://escholarship.org/uc/item/8s79k0qs>

Author

Xu, Shizhou

Publication Date

2024

Peer reviewed|Thesis/dissertation

Fairness in Machine Learning via Optimal Transport

By

Shizhou Xu
DISSERTATION

Submitted in partial satisfaction of the requirements for the degree of

DOCTOR OF PHILOSOPHY

in

APPLIED MATHEMATICS

in the

OFFICE OF GRADUATE STUDIES

of the

UNIVERSITY OF CALIFORNIA

DAVIS

Approved:

Naoki Saito

Thomas Strohmer (Chair)

Qinglan Xia

Committee in Charge

2024

To my wife Yuan Ni,
my parents Gang Xu & Jianfeng Ni.

Contents

Abstract	v
Acknowledgments	vi
Chapter 1. Introduction	1
1.1. Fairness in Machine Learning	1
1.2. Optimization Problems with Sensitive Variable Independence Constraint	5
1.3. Challenges in Machine Learning Fairness	8
1.4. Corresponding Contributions in Machine Learning Fairness	11
1.5. Setting and Notation	15
1.6. Dissertation Organization	17
Chapter 2. Preliminaries: Optimal Transport	19
2.1. General Distribution Case	19
2.2. Location-Scale Case and Optimal Affine Transport	24
Chapter 3. Pareto Frontier for L^2 -objective Machine Learning	27
3.1. Wasserstein Barycenter Characterization of Optimal Fair Learning	27
3.2. Optimal Affine Estimation of Barycenter	30
3.3. Wasserstein Geodesics Characterization of Pareto Frontier	36
Chapter 4. Fair Data Representation for Conditional Expectation Estimation	47
4.1. Objective & Constraint for Fair Data Representation	47
4.2. Wasserstein Barycenter Pair Characterization	53
4.3. Gaussian Marginals: Exact Solution	57
4.4. General Distribution: Optimal Affine Estimation	61
4.5. Optimal Fair Data Representation at the Pareto Frontier	66
4.6. Algorithm Design	68

4.7. Empirical Study: Fair Supervised Learning	71
Chapter 5. (In)Compatibility between Group and Individual Fairness	82
5.1. Generalized Individual Fairness Definitions	83
5.2. Problem Setting	86
5.3. Preliminaries on the (Pareto) Optimal Fair L^2 Learning	93
5.4. Compatibility between the Optimal Statistical Parity L^2 Learning and Individual Fairness	96
5.5. Compatibility between Pareto Frontier and (ϵ, δ) -IF	99
5.6. Composition Results	102
5.7. Empirical Study: Group and Individual Fair Supervised Learning	107
Chapter 6. Equalized Odds	116
6.1. Problem Setting	116
6.2. Solution via Conditional Wasserstein Barycenter	119
Chapter 7. Future Plan	124
7.1. Future Plan on Machine Learning Fairness	124
7.2. Future Plan in Broader Directions	125
Bibliography	127

Abstract

As machine learning powered decision-making becomes increasingly important in our daily lives, it is imperative to strive for fairness of the underlying data processing. In this work, we apply optimal transport technique to develop provably trustworthy solutions to open challenges in fair machine learning:

- (*Statistical parity*) We first apply the optimal affine transport to approach the post-processing Wasserstein barycenter characterization of the optimal fair L^2 -objective supervised learning via a pre-processing data deformation. We call it Wasserstein pseudo-barycenter. Then, we prove that the Wasserstein geodesics from learning outcome marginals to their barycenter characterizes the Pareto frontier between L^2 -loss and total Wasserstein distance among the marginals. Thereby, an application of McCann interpolation generalizes the pseudo-barycenter to a family of data representations via which L^2 -objective supervised learning algorithms estimate the Pareto frontier. Numerical simulations underscore the advantages: composition flexibility, sensitive information protection, computational efficiency, and applicability to unsupervised learning.
- (*Compatibility between group and individual fairness*) 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 similarities within their respective sensitive groups. The two fairness perspectives, while both desirable, often come into conflict. We analyze the existence of this conflict and its potential solution: When there exists a conflict between the two, we first relax the former to the Pareto frontier (optimal trade-off) between L^2 error and statistical disparity, and then identify 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.
- (*Equalized odds*) We apply conditional Wasserstein barycenter to characterize the optimal solution to odds-equalized data representation, via which a broad family of the trained supervised learning models satisfies equalized odds, under mild assumptions.

Acknowledgments

Completing this dissertation has been an incredible journey, and I am deeply grateful to everyone who supported me along the way.

First and foremost, I'd like to express my deepest gratitude to my advisor, Prof. Thomas Strohmmer. Your support, both academically and personally, has been an anchor for me throughout this journey. The challenges brought by the pandemic starting in 2020 and the separation from my wife, Yuan Ni, were too tough for me to navigate while continuing my study in mathematics. It was you who helped me restart my PhD at UC Davis, allowing Yuan and me to reunite and continue our shared dreams in mathematics. In academia, your mentorship not only introduced me to intriguing problems but also granted me the freedom to explore them, fostering my independence as a researcher. Moreover, your commitment to academic integrity and guidance have left a lasting impact on my approach to research and work ethics. You taught me that intuition holds as much value as technical prowess, that research is about uncovering "signals" rather than adding "noise", and that the best research strikes a balance between theoretical intrigue and real-world applicability. Thank you for reshaping not only my outlook on research and academia but also on life itself.

I am also profoundly thankful to the UC Davis Mathematics Department, which has been like a family to me. Special thanks to my dissertation committee members, Prof. Naoki Saito and Prof. Qinglan Xia. Prof. Saito, thank you for teaching me harmonic analysis, providing constructive critiques of my dissertation, and giving me career advice. Prof. Xia, thank you for sharing your expertise in optimal transport and for your career guidance. I also appreciate the support from Prof. Martin Fraas and Prof. Eric Babson during my qualification exams and for their research insights. My gratitude also goes to Prof. Janko Gravner, Prof. Joseph Biello, and Prof. John Hunter for sharing their knowledge in probability and PDE. The supportive and stimulating academic environment provided by the department has been essential to my growth. As an international student far from home, I found a family-like atmosphere within this community, for which I am deeply grateful.

Additionally, I would like to thank the professors at NYU who supported me during my difficult times in academia and paved the path for my PhD studies. My gratitude extends to Prof. C. Sinan Güntürk, Prof. Michael L. Overton, and Prof. Quanyan Zhu for helping me keep my dream in mathematics alive in 2019 and 2020. I also thank Prof. Yuri Bakhtin, Prof. Fengbo Hang, Prof.

Esteban G. Tabak, Prof. S. R. Srinivasa Varadhan, and Prof. Lai-Sang Young for teaching me and shaping my standards in analysis, probability, optimal transport, and dynamical systems early in my studies. I wish I had possessed more academic maturity back then to fully appreciate the time spent with them.

Special thanks to my colleagues and friends, particularly Siddarth Chintamani, Gal Dimand, Xue Feng, Girish Kumar, Eli Moore, Tim & Sandy Moran, and Jingyang Shu for their unwavering support and companionship. Your friendship kept this journey both warm and enjoyable.

I would also like to thank my pets, Chubby and Otto, who are my first-ever pets, for their support during this long journey. Otto, thank you for keeping me physically and mentally healthy during the pandemic. Chubby, thank you for spending your life with me in Davis and waiting until the very end for me to complete my dissertation.

Lastly, I am eternally grateful to my family for their unconditional love and support. To my spouse, Yuan Ni, you have been my bedrock through all the ups and downs over the past years. You are the only one who truly understands how hard I have worked to reach this moment. I apologize for being childish during the journey and thank you for your understanding and patience in waiting for me to mature. To my parents and my sister, your unwavering belief in me and encouragement to pursue my dreams have been my guiding light.

To everyone who, in one way or another, has offered their support—whether or not your name appears here—thank you. You have all been instrumental in the completion of this dissertation.

Thank you all from the bottom of my heart.

(Reflecting on this acknowledgment, I've come to appreciate just how much Davis means to both Yuan and me. Saying farewell is never easy, but as they say, all good things must come to an end. I'll carry the cherished memories of our time in Davis with me wherever life takes me, revisiting them whenever I need a reminder of the joy we experienced here.)

CHAPTER 1

Introduction

1.1. Fairness in Machine Learning

Our society is increasingly influenced by artificial intelligence as (direct or indirect) decision-making and information-sharing processes become more reliant on statistical inference and machine learning, especially considering the recent development of large language models (LLMs) and generative AI. The potentially significant long-term impact from sequences of automated (facilitate of) decision-making and information-sharing has brought large concerns about bias and discrimination in machine learning [5, 48]. Machine learning based on unbiased algorithms can naturally inherit the historical biases that exist in data and hence reinforce the bias via automated decision-making process [13].

One straightforward partial remedy is to exclude the sensitive variables from the data set used in the learning and decision process. But such exclusion merely eliminates disparate treatment, which refers to direct discrimination, and leaves disparate impact, which refers to unintended or indirect discrimination, remaining in both data and learning outcome [25]. Examples of the legal doctrine of disparate impact include *Griggs v. Duke Powers Co.* [11] and *Ricci v. DeStefano* [1], where the decision is based on factors that are strongly correlated with race, such as intelligence qualification in the former and the racially disproportionate test result in the latter, are ruled illegal by the US supreme court. As a result, along with the trending development of automated decision-making and information-sharing, the need for more sophisticated but practical and explainable techniques has made fairness in machine learning an important research area [42].

There are two important but potentially conflicting fairness concepts that look at fairness from different perspectives:

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

- (*Individual fairness*) aims to guarantee that individuals who share similar qualification data would receive similar learning outcomes.

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 group fairness, 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 group fairness can easily lead to a significant violation of individual fairness. Similarly, individual fairness tends to keep individuals sharing similar qualifications close in the learning outcome space. Therefore, given sensitive information being excluded from qualification, such closeness preservation is likely to extend the distributional discrepancy 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 (Definition 1.2.1).

In this dissertation, we first target an important definition of group fairness: *statistical parity* [23], which is defined by statistical independence between the sensitive variable and the learning outcome. For example, an automated college admission outcome satisfies statistical parity for gender if the admission rate is the same across all gender groups in the application pool. In other words, the admission decision is statistically independent of the information of gender. The importance of statistical parity results from its close relation to disparate impact and hence long-term structural influence [57]. Here, we start with characterizing and deriving the optimal fair learning outcome when fairness is defined by statistical parity, then prove the Pareto frontier when trade-offs between statistical parity and utility are needed, and finally show how to achieve the optimal fair learning via a pre-processing data deformation. Thereafter, we focus on the compatibility between the Pareto frontier and individual fairness regulations. Last but not least, we study equalized odds [31], which is another key fairness definition inspired by the belief that a fair model should give the same prediction accuracy to different sensitive groups.

Therefore, in the remainder of my dissertation, fairness or group fairness and statistical parity are used interchangeably¹, except in Chapter 6 where fairness is defined by equalized odds.

¹There are many other notions of fairness than statistical parity, such as equalized odds or equal opportunity, which all have their benefits and shortcomings [17]. A discussion of the advantages or disadvantages of the different concepts of fairness is beyond the scope of this work.

Before further discussing statistical parity, we note that fairness in machine learning should not be defined by a single condition without considering the application context. The goal of the present work is to provide theoretically reliable and explainable tools to help practitioners obtain the optimal (w.r.t. utility) solutions at any chosen statistical disparity level, provided one chooses to adopt statistical parity (or limited statistical dependence between the learning outcome and the sensitive information) as a meaningful fairness definition in one’s particular application context.

Remark 1.1.1 below provides a more detailed discussion on statistical parity, namely how the utility optimization solves some major insufficiency of the original statistical parity definition and improves statistical parity to proportional equality, a fairness concept similar to equity in modern ethics which can be traced back to Aristotle and Plato [6, 20].

REMARK 1.1.1 (Statistical parity enhanced by utility optimization). *Statistical parity is one of the most important definitions of group fairness. It has advantages such as (1) legal support on mitigating adverse impact and (2) the long-term effect resulting from the enforced involvement of minority groups or diversity in learning outcomes via affirmative action [33]. On the other hand, there are three major criticisms about statistical parity that are often mentioned, e.g. see [23, 31]: (1) reduced utility, (2) self-fulfilling prophecy, (3) subset targeting. However, we notice that the first two are insufficiencies with respect to utility. Therefore, the proposed method mitigates these two insufficiencies.*

- *(Utility) The development of the Pareto frontier allows us to achieve a desirable statistical disparity level with theoretically provable minimum (hence necessary) utility sacrifice. Equivalently, practitioners can choose a tolerable utility sacrifice level so that the Pareto frontier will provide a learning outcome with the minimum statistical disparity while not violating the utility sacrifice tolerance.*
- *(Self-fulfilling prophecy) As mentioned in [23, 31], self-fulfilling prophecy results from random, careless, or malicious selection in minority groups. But the barycenter characterization method guarantees the optimal fair model to make good selections in all sensitive groups to maximize utility. Section 1.3 contribution point 4 and Section 2.1 provide, respectively, the intuitive and technical explanation of how the utility maximization enforces the model to give similar learning outcomes to data points sharing relatively (within their sensitive groups) similar qualifications.*

For example, if race is the sensitive information and an admission test score is the only qualification variable, a barycenter-characterized optimal fair admission model would give admission to the same percentage of top-score students in each of their racial groups.

Interestingly, the interpretation is consistent with the philosophical definition of fairness involving proportional equality: a model is fair (with respect to the sensitive information) if it distributes proportional chance or prediction to proportionally qualified independent variables within each of the sensitive groups.

Beginning with the celebrated work “Fairness through Awareness” [23], 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:

- (*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 [14, 36, 53, 56].
- (*In-processing*) The fairness definition is quantified and then integrated into the training process by penalizing unfair outcomes [8, 55]. 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 outcomes in terms of both performance and fairness.
- (*Post-processing*) The definition of fairness is enforced directly on the learning outcome [19, 29, 34, 47, 53]. 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 post-processing approach has received significant attention due to the following remarkable result: the optimal fair distribution of supervised learning, such as classification [34] and regression [19, 29], can be characterized as the Fréchet mean of the learning outcome marginals on the Wasserstein space, which is also known as the Wasserstein barycenter in the optimal transport literature. (See Remark 2.1.2 for more details on learning outcome marginals.) Thereafter, [53]² 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, [53] provides a provable Pareto frontier,

²The results in Chapter 3 and 4 are published as [53] in the Journal of Machine Learning Research (JMLR) in 2023.

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.

The following remark provides an intuition of the Wasserstein (\mathcal{W}_2) barycenter characterization, on which we develop our theoretical results and algorithms.

REMARK 1.1.2 (Intuition of Wasserstein barycenter characterization). *The Fréchet mean is the closest point to a set of points in a metric space and, therefore, a generalization of the mean on the Euclidean space to general metric spaces such as the Wasserstein space. Intuitively, one can consider the barycenter (Fréchet mean in Wasserstein space) characterization of optimal fair learning outcome as an analog of representing a set of points by their average, which thereby optimally (with respect to total moving distance) removes the disparity among those points, except that each point is now in Wasserstein space, and hence a distribution. See Section 1.4 point 4 below for more details.*

Despite the theoretical elegance of the post-processing barycenter characterization, challenges remain in theory and practice (see Section 1.3 for a detailed explanation of the challenges), especially compared to pre-processing or data representation methods.

Fair machine learning using a pre-processing approach has been considered in [14, 25, 30, 36, 47]. While the Wasserstein barycenter provides a mathematically rigorous characterization of the post-processing optimal learning outcome, optimal fair data representation for general supervised learning models still lacks a theoretical characterization. See, for example, [17, Section 3.4, 3.5] for more details on the current challenges in fair data representation design for general machine learning models beyond classification, not to mention data representations that provide the optimal trade-off between accuracy and fairness.

The goal of the present work is to develop an optimal fair data representation characterization so that a broad family of supervised learning models trained via fair data representation can result in the provably optimal fair learning outcome. The ultimate goal is to develop methods that enjoy both the mathematically rigorous characterization of optimal fair learning and the flexibility of pre-processing.

1.2. Optimization Problems with Sensitive Variable Independence Constraint

The statistical parity constraint for supervised learning or data representation in a nutshell is a constraint on the dependence between the learning outcome and a chosen sensitive variable:

DEFINITION 1.2.1 (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 $S \in \{\hat{Y}, Z\}$. Equivalently, the constraint limits the ability to access or reverse engineer the sensitive variable from the learning outcome or data representation. Therefore, although the theory and methods in the present work aim to solve current challenges in machine learning fairness, they can also be useful in other areas where sensitive or undesirable information needs to be eliminated within the existing learning outcome or data. One example of such an area other than fair machine learning is machine (feature) unlearning. It starts from [15] and now has a sizable body of research works.

Here, we summarize the constrained optimization problems solved in the present work. We prove existence (and uniqueness, if possible) results via a constructive characterization approach so that an explicit formula of the solutions becomes available. Practitioners and researchers interested in limiting the statistical dependence between the learning outcome or data representation and certain feature variables can directly refer to the corresponding section for results. We leave the underlying motivations resulting from machine learning fairness to the following two subsections.

In Section 3.1, we target the following problem:

PROBLEM 1 (Optimal fair L^2 -objective learning outcome).

$$(1.1) \quad \inf_{f \in L^2(\mathcal{X} \times \mathcal{Z}, \mathcal{Y})} \{\|Y - f(X, Z)\|_2^2 : f(X, Z) \perp Z\}$$

Here, Y is the dependent variable, and $f(X, Z)$ is an estimator that uses the independent variable X and sensitive variable Z to estimate Y . The loss function aims to maximize utility by minimizing

the L^2 -norm between Y and $f(X, Z)$:

$$\|Y - f(X, Z)\|_2^2 = \int_{\Omega} \|Y - f(X, Z)\|^2 d\mathbb{P}.$$

For $S \in \{X, Y, Z\}$, $S : \Omega \rightarrow \mathcal{S}$ is a random variable from Ω to the state space \mathcal{S} . $\|\cdot\|$ denotes the Euclidean norm. The constraint $f(X, Z) \perp Z$ guarantees that the final result is independent of the sensitive information Z and hence satisfies statistical parity. Finally, the admissible function space $L^2(\mathcal{X} \times \mathcal{Z}, \mathcal{Y})$ is the space of all square-integrable measurable functions from $\mathcal{X} \times \mathcal{Z}$ to \mathcal{Y} . (Our proof shows Problem 1 does not change if one allows all measurable functions $\mathcal{X} \times \mathcal{Z}$ to \mathcal{Y} . See Remark 1.2.1 below for more details.) The reason of allowing all measurable functions in our problem setting is due to the recent development of deep neural networks that are capable of estimating arbitrary measurable functions.

In Section 3.3, we relax the above strict independence constraint by applying a quantification of statistical disparity: the *Wasserstein disparity*, which is the average pairwise Wasserstein distance among conditional (on Z) distributions of $f(X, Z)$, denoted by $D(f(X, Z), Z)$. It has the following desirable properties: (1) $D(f(X, Z)) = 0$ if and only if $f(X, Z) \perp Z$. (2) The larger D is, the more disparities there are among the marginals (w.r.t. Z) of $f(X, Z)$. (3) D has a meaningful interpretation in physics as the minimum expected amount of work required to remove the distributional discrepancy between two randomly chosen sensitive groups on the learning outcome. Therefore, fixing a disparity tolerance level $d \in [0, \infty)$,

PROBLEM 2 (Optimal L^2 -objective learning Pareto frontier).

$$(1.2) \quad \inf_{f \in L^2(\mathcal{X} \times \mathcal{Z}, \mathcal{Y})} \{\|Y - f(X, Z)\|_2^2 : D(f(X, Z), Z) < d\}$$

gives us the corresponding Pareto optimal solution. That is, if one wants a lower L^2 -loss than provided by the infimum in Problem 2, then it is necessary to increase the tolerance level d . Equivalently, if one wants to lower the tolerance level d , then it is necessary to sacrifice more L^2 -loss than the infimum.

REMARK 1.2.1 (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 [43]. We note here that Problem 1 and 2 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.

In Section 4.2, we provide a theoretical characterization of the solution to

PROBLEM 3 (Optimal fair data representation for conditional expectation estimation).

$$(1.3) \quad \inf_{(\tilde{X}, \tilde{Y}) \in \mathcal{D}} \{ \|Y - \mathbb{E}(\tilde{Y}|\tilde{X})\|_2^2 : \tilde{X}, \mathbb{E}(\tilde{Y}|\tilde{X}, Z) \perp Z \},$$

where \mathcal{D} is the admissible data representation set we define later. Here, the objective function aims to maximize the potential utility remaining within the deformed data (\tilde{X}, \tilde{Y}) by minimizing the L^2 distance between the perfect estimator $\mathbb{E}(\tilde{Y}|\tilde{X})$ on (\tilde{X}, \tilde{Y}) and the original Y , so that better estimation of $\mathbb{E}(\tilde{Y}|\tilde{X})$ leads to better prediction of Y . The constraint $\tilde{X}, \mathbb{E}(\tilde{Y}|\tilde{X}, Z) \perp Z$ guarantees: (1) $f(\tilde{X}) \perp Z$ for $\forall f : \mathcal{X} \rightarrow \mathcal{Y}$, such that any estimator of $E(\tilde{Y}|\tilde{X})$ is independent of Z ; (2) The perfect adversarial estimator $\mathbb{E}(\tilde{Y}|\tilde{X}, Z)$ is independent of Z , so that a better estimation of $E(\tilde{Y}|\tilde{X}, Z)$ leads to more independence of Z (alignment between the training objective and independence constraint). In addition, one may choose the following alternative constraints according to the application context: (1) $\tilde{X} \perp Z$, which guarantees $f(\tilde{X}) \perp Z$ for all measurable f as mentioned above; (2) $(\tilde{X}, \tilde{Y}) \perp Z$, which guarantees any (adversarial) supervised or unsupervised learning on (\tilde{X}, \tilde{Y}) to be independent of Z . The first alternative is useful if only measurable functions of X are allowed, whereas the second should be applied when one does not know which features are dependent or independent. See Section 4.1 for a more detailed derivation and explanation of the data representation objective function and constraints.

1.3. Challenges in Machine Learning Fairness

Now, we go back to the motivation behind the optimization problems listed above: fair machine learning. We first summarize the limitations of the current post-processing characterization and the current methods based on it to estimate the optimal fair learning outcome.

To show the relevance and importance of the challenges, we note that both *Pareto frontier characterization* and *Optimal fair data representation characterization* below are mentioned as important

open problems at the NeurIPS FairML Workshop in 2022. Also, *Group fairness & individual fairness (in)compatibility* is listed as a major open problem in the recent machine learning fairness survey by Chouldechova and Roth [17, Section 3.1].

- (*Optimal fair learning characterization*) The post-processing barycenter characterization lacks theoretical and computational generalization to high-dimensional data spaces, such as text or image spaces. From a theoretical perspective, the current works [19, 29, 47] focus on classification and 1-dimensional regression. From a computational perspective, the current works apply the coupling of cumulative distribution functions (cdf) of the learning outcome sensitive conditionals to find the barycenter and the inverse of the cdf to compute the optimal transport map. Both the coupling and the inverse of the cdf are computationally expensive. Furthermore, since the inverse of the cdf cannot be generalized to high-dimensional spaces, the current methods lack the generalization to supervised learning with high-dimensional dependent variables.

Due to the recent development of generative AI models, it is now important to have fair machine learning methods for arbitrarily high-dimensional data. We hope the present work on the L^2 space can be a starting point for fair machine learning or data representation on more general spaces for high-dimensional data.

- (*Pareto frontier characterization*) The current post-processing barycenter characterization lacks both theoretical and computational generalization to (an estimation of) the optimal trade-off, also known as the Pareto frontier, between prediction accuracy and fairness. In theory, there is a lack of characterization of the Pareto frontier (optimal trade-off) between utility and fairness. Current works on the Pareto frontier, such as [47], apply tight inequalities based on the convexity of distance metrics to suggest the optimal trade-off coincides with the Wasserstein geodesic path. While such inequalities are tight for a broad type of metrics on the space of probability measures, they are *not tight* for the Wasserstein metric. Hence, the inequalities are not able to extend the mathematically rigorous Wasserstein barycenter characterization of the optimal fair learning outcome to a Pareto frontier. From a computational perspective, current methods, such as [47], apply interpolation between the inverses of the sensitive conditional cdf's (more specifically, interpolating the data points that share the same image under the sensitive conditional cdf's) to estimate the geodesics. In addition to the drawbacks mentioned above, the inverse of the

cdf also does not come with an explicit form, which makes the computation of an interpolation between two cdf inverses even more cumbersome.

- (*Optimal fair data representation characterization*) The post-processing nature of the characterization requires explicit or implicit sensitive information in the training and decision-making process. More specifically, in order to apply the barycenter characterization to find the optimal fair learning outcome or to make predictions to newly incoming data, one needs the following steps: (1) Estimate the conditional expectation and obtain its conditional distributions with respect to the sensitive information; (2) Find the Wasserstein barycenter of the sensitive conditionals of the conditional expectation estimation or the learning outcome; (3) Compute the optimal transport maps from each sensitive conditional to the barycenter; (4) Apply each transport map to the conditional with the matched sensitive information. Here, not only does the trained model still inherit unfairness, but it is also clear that sensitive information needs to be attached to both the dependent variable or incoming data and its learning outcome or prediction, until the very last post-processing step of finding the barycenter comes to the rescue. Hence, we say that the characterization has a post-processing nature. As a result, the user needs access to the sensitive information of each individual incoming data at every step during the learning process. Such a strong access to sensitive information makes the supervised learning process vulnerable to attack and sensitive information leakage.

The post-processing nature of the characterization also suffers from the lack of flexibility in model selection, modification, and composition. For model selection and modification, a practitioner would have to perform the post-processing step for every model and every modification in order to compare the corresponding optimal fair learning outcomes. See Table 4.5 for more details on the additive computational cost of the post-processing approach compared to the one-time cost of the proposed pre-processing approach. For model composition, we consider the simple example $task_2 \circ task_1$ where $task_i, i \in \{1, 2\}$ are trained supervised learning models. In practice, there is a good chance that $task_1$ and $task_2$ belong to different practitioners or organizations, denoted by practitioners 1 and 2, respectively. Therefore, to protect sensitive information from practitioner 2, practitioner 1 will perform the post-processing step to obtain a fair learning outcome and provide it as an input variable for the training task of practitioner 2. But unless $task_2$ needs no more input variables other than the dependent variables of $task_1$ (in

that case, $task_1$ would be fair data representation design), still practitioner 2 needs full access to the sensitive variable attached to its input data, which includes the desensitized $task_1$ output and other input variables. Such attachment makes the post-processing step performed by practitioner 1 meaningless. Considering the recent development of decentralized learning in practice, such drawback in model composition makes a model-independent fair data representation more applicable than a post-processing solution.

- (*Explainability*) Many of the current fair machine learning methods are proposed without utility guarantee or explainability. Such a lack of utility guarantee or explainability prevents the study of fair machine learning from practical use. For instance, Wells Fargo [57] concluded recently that current fair machine learning methods are black-box methods, and hence they hesitate to adopt fair machine learning techniques.
- (*Group fairness & individual fairness (in)compatibility*) As mentioned above, there is a potential conflict between group fairness and individual fairness as the enforcement of group fairness could result in opposite effects on individual learning outcome due to different sensitive information. This, naturally, gives rise to the following question, which remains open on the current frontiers of machine learning fairness [17, Section 3.1]: *When can one enjoy the best of both group fairness and individual fairness?*

1.4. Corresponding Contributions in Machine Learning Fairness

We provide a road map of the tools that we have developed in response to each of the listed challenges and how the present work combines all the tools to provide (exact solution and estimation of) the fair data representation at the Pareto frontier.

Here, we note that the results on *Optimal fair learning characterization*, *Pareto frontier characterization*, and *Optimal fair data representation characterization* are published as [53] in the Journal of Machine Learning Research (JMLR) in 2023, whereas the results on *Group fairness & individual fairness (in)compatibility* is submitted to the SIAM Journal on Mathematics of Data Science (SIMODS) [54].

- (*Optimal fair learning characterization*) In response to the theoretical part of the first challenge, Lemma 5.3.1 in Section 3.1 provides a characterization (with explicit construction) of the exact solution to Problem 1 (the optimal fair L^2 -objective learning). The result shows that the

infimum loss value of Problem 1 can be nicely decomposed into two parts: (1) L^2 orthogonal projection loss and (2) independence projection loss. Also, the result now allows the data spaces $\mathcal{X}, \mathcal{Y}, \mathcal{Z}$ to be $[k]^d, \mathbb{N}^d, [0, l]^d$, or \mathbb{R}^d for arbitrary dimension $d < \infty$.

To address the challenge of computing the Wasserstein barycenter in high-dimensional data spaces [3], we propose a method that applies affine transport maps to find the *optimal affine estimation* of the post-processing optimal fair L^2 -objective supervised learning outcome with an arbitrarily finite-dimensional dependent variable, which responds to the first challenge listed above. In particular, by restricting admissible transport maps to be affine and making a corresponding relaxation to the fairness constraint, we derive a relaxed version of Problem 1, stated as Problem 4. Applying the optimal affine transport maps [2], Definition 3.2.1 introduces the post-processing *pseudo-barycenter*, Lemma 3.2.1 shows the proposed pseudo-barycenter coincides with the true barycenter when the sensitive conditionals are Gaussian, and finally, Theorem 3.2.1 proves that the pseudo-barycenter is the optimal affine estimation of the true barycenter in the general conditional distribution case and provides the estimation error. Optimal affine transport and pseudo-barycenter have the advantage of computational efficiency, compared to the current methods, due to the explicit matrix form of the transport map and the nearly closed-form solution to the pseudo-barycenter.

The importance of optimal affine maps encompasses much more than a solution to the first challenge. The optimal affine maps together with McCann interpolation [41] help us in obtaining an explicit form of the geodesic path characterization of the Pareto frontier in Section 4. More importantly, Section 5 shows that optimal affine maps and the pseudo-barycenter are necessary tools to overcome the post-processing nature of the Wasserstein barycenter characterization by exploiting the linearity of conditional expectation and thereby generating optimal fair data representations.

- (*Pareto frontier characterization [53]*) In Section 4, we prove an exact characterization of the solution to Problem 2 (the optimal utility-parity trade-off or Pareto frontier) in response to the theoretical part of the second challenge. In particular, Theorem 3.3.1 shows that, when utility loss and disparity are quantified respectively by the L^2 distance (between the true outcome Y and the prediction $\hat{Y} = f(X, Z)$) and the average pairwise \mathcal{W}_2 distance among the sensitive conditionals of \hat{Y} , the optimal trade-off happens if and only if the conditionals of \hat{Y} travel

along the Wasserstein geodesic path from the conditionals of $\mathbb{E}(Y|X, Z)$ to their barycenter. Therefore, we say that the Pareto frontier is on the Wasserstein space. Corollary 3.3.1 then derives an explicit form of the Pareto optimal solution to Problem 2. The result is a natural extension to the post-processing Wasserstein barycenter characterization of the optimal fair learning outcome: the barycenter characterization coincides with the point at zero disparity on the Pareto frontier. Interestingly, our result shows that the Pareto frontier is linear.

To solve the computational challenge of the geodesic path, Remark 3.3.1 applies McCann interpolation together with the optimal affine maps and the pseudo-barycenter to derive a computationally efficient (nearly) closed-form formula to estimate the Pareto frontier, which results in Algorithm 1.

- (*Optimal fair data representation characterization [53]*) In response to the third challenge, the present work proposes in Section 4.1 Problem 3 (optimal fair data representation problem), which makes the objective function and the fairness (statistical parity) constraint *model-independent* and therefore suitable for fair data representation design. More specifically, by applying the Minkowski inequality, we use an objective function to maximize the potential utility remaining in the data. On the other hand, a fair data representation should provide a fairness guarantee to arbitrary L^2 -objective supervised learning models. Therefore, the present work proposes a pre-processing fairness constraint to guarantee fairness in the learning outcome of arbitrary L^2 -objective models trained via the fair data representation.

In Section 4.2, Lemma 4.2.4 first provides a characterization of the exact solution to Problem 3 under a mild assumption. Next, Definition 4.3.2 and Definition 4.3.1 define the dependent and independent pseudo-barycenter, respectively. Then, similar to solving a relaxation of the post-processing characterization to obtain the optimal affine estimation, Theorem 4.3.1 proves that the dependent and independent pseudo-barycenter pair coincides with the true solution to the optimal fair data representation when the conditional data distributions are Gaussian, and Theorem 4.4.1 proves that the pseudo-barycenter pair forms the optimal affine estimation of the optimal fair data representation.

To derive (an estimation of) fair data representation at the Pareto frontier, Corollary 4.5.1 in Section 4.5 first provides a characterization of the Pareto frontier for conditional expectation on a fixed sigma-algebra. Finally, combining optimal affine map, pseudo-barycenter, together

with a diagonal argument in Remark 4.5.1, we derive an estimation of the fair representation at the Pareto frontier, which results in Algorithm 1 and Algorithm 2.

Furthermore, in Section 4.7, experiments show that the proposed fair data representations preserve as large an amount of information (w.r.t. the L^2 objective) as the fairness constraint allows. Therefore, it provides a better and more flexible solution to fair learning compared to encoding-based data representations [14, 56], which encode the information of the original data into some binary feature variables designed to guarantee statistical parity for classification. Surprisingly, experiments also show that applying the pseudo-barycenter results in nearly zero utility loss compared to the post-processing barycenter characterization solution.

- (*Explainability*) In addition to the provable utility guarantee resulting from the Pareto frontier, the proposed method also has a meaningful interpretation from a *datapoint-wise perspective* in how it achieves the statistical parity requirement: A data point of the optimal fair learning outcome is the Euclidean average of the optimally matched data points from each of the sensitive groups. Here, matching means partitioning the original data set into subsets consisting of one point from each sensitive group. Each subset is called a match. The points within a match are called matched points. Optimality in matching is equivalent to minimization of the expected variance within a randomly chosen match. Such expected (hence total) variance minimization enforces points with similar relative positions in their sensitive marginal to form a match. For example, assume that there are two sensitive conditionals $A = \{1 \text{ (low in A)}, 4 \text{ (high in A)}\}$ and $B = \{2 \text{ (low in B)}, 3 \text{ (high in B)}\}$, then the optimal matching is

$$\{\{1 \text{ (low in A)}, 2 \text{ (low in B)}\}, \{3 \text{ (high in B)}, 4 \text{ (high in A)}\}\}$$

to minimize the expected or total variance within the matches. The optimal matching in high-dimensional L^2 spaces shares the same geometric intuition with the simple example. That is, from a point-wise perspective, the optimal fair learning achieves statistical parity by first matching the points with similar relative positions in their sensitive groups and then representing the matched ones with their Euclidean average.

- (*Group fairness & individual fairness (in)compatibility [54]*) In Chapter 5, we provide a theoretically provable answer to the question of (in)compatibility between group fairness and individual fairness, when adopting statistical parity as the group fairness definition. In particular, we prove

sufficient conditions for the compatibility between individual fairness and the Pareto optimal (with respect to utility) statistical parity L^2 -objective learning.

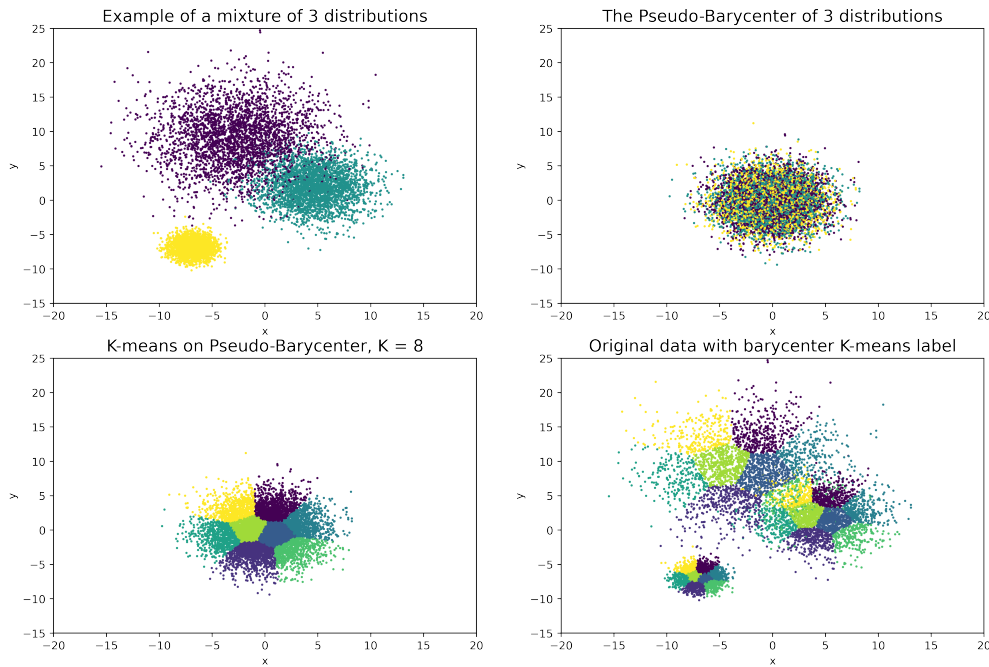


FIGURE 1.1. The upper left panel depicts three distributions, sampled from an isotropic Gaussian distribution with different first two moments. The upper right panel shows the barycenter of the three sample distributions. The lower left panel provides the K-means ($K=8$) clustering result on the barycenter. The lower right panel shows the original distribution with the barycenter K-means labeling.

1.5. Setting and Notation

In the rest of the work, $\mathcal{L}(X) = \mathbb{P} \circ X^{-1} : \mathcal{B}_{\mathcal{X}} \rightarrow [0, 1]$ denotes the distribution or law of X , which is a function that assigns each event in the Borel sigma-algebra, $\mathcal{B}_{\mathcal{X}}$, a probability. Let $\lambda := \mathcal{L}(Z)$ denote the law of the sensitive random variable to simplify notation. To remove sensitive information Z , the method we propose is to find a set of maps $T_x := \{T_x(\cdot, z)\}_z$ such that $T_x(\cdot, z) : \mathcal{X} \rightarrow \mathcal{X}$ pushes the conditional (on $\{Z = z\}$) distribution (see the definition of conditional distribution $\mathcal{L}(X_z)$ below) forward to a common probability measure $\mathcal{L}(\tilde{X})$ for λ -a.e. $z \in \mathcal{Z}$. Also, when restricting T to be a linear map or a matrix, we use $T \succ 0$ to denote T is positive definite, and $\|T\|_F$ to denote its Frobenius norm.

Given a measurable map $T : \mathcal{X} \rightarrow \mathcal{X}$ and a probability measure $\mu \in \mathcal{P}(\mathcal{X})$, $T_{\#}\mu$ denotes the push-forward probability measure that is defined as the following: for any event, A , in the Borel

sigma-algebra, $\mathcal{B}_{\mathcal{X}}$, $T_{\#}\mu(A) := \mu(T^{-1}(A))$. In the rest of the dissertation, we often say T pushes μ forward to $T_{\#}\mu$.

The conditional distributions $\{\mathcal{L}(X_z)\}_z$ are defined uniquely λ -a.e. by the disintegration theorem [46, Box 2.2]. Hence, $z \rightarrow \mathcal{L}(X_z)$ is Borel measurable and, for all Borel measurable sets $E \in \mathcal{B}_{\mathcal{X}}$, $\mathbb{P}(E) = \int_{\mathcal{Z}} \mathbb{P}(X_z^{-1}(E)) d\lambda(z)$. The application of the disintegration theorem aims to allow \mathcal{Z} to be uncountably infinite, such as the real line or the real vector space. In the practical case of a finite data set, when the data set (X, Z) is $\{(x_i, z_i)\}_{i \in [N]}$, for each $z \in \mathcal{Z}$, the empirical conditional random variable (with uniform distribution) is defined as follows:

$$X_z := \{x_i : (x_i, z_i) \in (X, Z), z_i = z\}.$$

Therefore, on the product data space $\mathcal{X} \times \mathcal{Z}$ with a joint distribution, the law of the random variable or vector X_z is the conditional distribution on $\{Z = z\}$.

The present work often assumes the conditionals $\{\mathcal{L}(X_z)\}_{z \in \mathcal{Z}} \subset \mathcal{P}_{2,ac}(\mathcal{X})$. Here, $\mathcal{P}_{2,ac}(\mathcal{X})$ denotes the set of probability measures on \mathcal{X} that have finite second moments and are absolutely continuous with respect to the Lebesgue measure. The finite second moment assumption guarantees the Wasserstein distance to be well-defined without being infinite. The absolute continuity assumption guarantees the existence of their Wasserstein barycenter (See Definition 2.1.3) and the respective (almost surely invertible) optimal transport maps that map them to the barycenter. The present work denotes the barycenter by $\overline{\mathcal{L}(X_z)}$ or $\overline{\mathcal{L}(X)}$ interchangeably, and denotes the optimal transport map that pushes $\mathcal{L}(X_z)$ to $\overline{\mathcal{L}(X)}$ by T_z or $T(\cdot, z)$.

To simplify notation and proof, we define \overline{X} to be the random variable that satisfies the following: for λ -a.e. $z \in \mathcal{Z}$,

$$(1.4) \quad \overline{X}_z = T_z(X_z).$$

In other words, the couple (X_z, \overline{X}_z) is a coupling of $(\mathcal{L}(X_z), \overline{\mathcal{L}(X)})$ and satisfies:

$$(1.5) \quad \|X_z - \overline{X}_z\|_2^2 = \mathcal{W}_2^2(\mathcal{L}(X_z), \overline{\mathcal{L}(X)})$$

for λ -a.e. $z \in \mathcal{Z}$. We refer interested readers to [50, 51] for more details on the assumption of $\mathcal{P}_{2,ac}(\mathcal{X})$ and the coupling of measures. In the rest of the dissertation, we call \overline{X} the Wasserstein barycenter of $\{X_z\}_z$.

In solving the post-processing characterization, with the assumption of $\mathbb{E}(Y|X, Z)$, one first finds the Wasserstein barycenter of $\{\mathcal{L}(\mathbb{E}(Y|X, Z)_z)\}_z$, denoted by $\overline{\mathcal{L}(\mathbb{E}(Y|X, Z)_z)}$. Here, $\mathbb{E}(Y|X, Z)_z$ denotes the conditional of $\mathbb{E}(Y|X, Z)$ on $\{Z = z\}$ for λ -a.e. $z \in \mathcal{Z}$. Then one applies the optimal transport map $T(\cdot, z) : \mathcal{Y} \rightarrow \mathcal{Y}$ which pushes $\mathbb{E}(Y|X, Z)_z$ forward to $\overline{\mathbb{E}(Y|X, Z)}$ for λ -a.e. $z \in \mathcal{Z}$.

In solving the pre-processing characterization, one has two different optimal transport maps to deform X and Y . For the dependent variable, we define $T_y = \{T_y(\cdot, z)\}_z$, $\mathcal{L}(Y_z)$, and $\mathcal{L}(\tilde{Y})$ analogously, but require merely the agreement of $\mathcal{L}(\mathbb{E}(\tilde{Y}|\tilde{X}, Z)_z)$ for λ -a.e. $z \in \mathcal{Z}$. The λ -a.e. agreement of $\mathcal{L}(\mathbb{E}(\tilde{Y}|\tilde{X}, Z)_z)$ means that the laws of the random variables or vectors $\mathbb{E}(\tilde{Y}|\tilde{X}, Z)_z$ are equal, except for some z on a λ -null set on \mathcal{Z} . In other words, on the Borel measurable space $(\mathcal{Y}, \mathcal{B}_\mathcal{Y})$, for any set B in the Borel sigma-algebra $\mathcal{B}_\mathcal{Y}$, we have $\mathbb{P} \circ [\mathbb{E}(\tilde{Y}|\tilde{X}, Z)_{z_1}]^{-1}(B) = \mathbb{P} \circ [\mathbb{E}(\tilde{Y}|\tilde{X}, Z)_{z_2}]^{-1}(B)$ for all $z_1, z_2 \in \mathcal{Z}$, except on a set $N \subset \mathcal{Z}$ such that $\lambda(N) = 0$.

Therefore, by generating and applying (T_x, T_y) to the data, we achieve $\mathbb{E}(\tilde{Y}|\tilde{X}, Z) \perp Z$, i.e. *statistical parity*, due to the enforced λ -a.e. agreement of $\mathcal{L}(\mathbb{E}(\tilde{Y}|\tilde{X}, Z)_z)$. Combining the application of deformation maps and (1.3), we obtain the fair data representation optimization problem

$$(1.6) \quad \inf_{(\tilde{X}, \tilde{Y}) \in \mathcal{D}} \{ \|Y - \mathbb{E}(\tilde{Y}|\tilde{X})\|_2^2 : \tilde{X}, \mathbb{E}(\tilde{Y}|\tilde{X}, Z) \perp Z \}$$

with the admissible set \mathcal{D} is defined as

$$(1.7) \quad \mathcal{D} := \{(\tilde{X}, \tilde{Y}) : \tilde{X} = T_x(X, Z), \tilde{Y} = T_y(Y, Z)\},$$

Here, $T_x(\cdot, z) : \mathcal{X} \rightarrow \mathcal{X}$ and $T_y(\cdot, z) : \mathcal{Y} \rightarrow \mathcal{Y}$ are Borel measurable maps. We denote the set of admissible \tilde{X} and \tilde{Y} by $\mathcal{D}|_{\mathcal{X}}$ and $\mathcal{D}|_{\mathcal{Y}}$, respectively. The reason underlying the definition of \mathcal{D} is that the fair data should still has its foundation from the real data, albeit suitably “deformed”.

1.6. Dissertation Organization

The rest of the dissertation is organized as follows: Chapter 2 reviews the tools in optimal transport that are needed to derive results in the present work: Wasserstein space, Wasserstein barycenter, and optimal affine transport within a location-scale family. In Chapter 3, we first generalize the current barycenter characterization of optimal regression to optimal L^2 -objective supervised learning, then defines pseudo-barycenter, and proves pseudo-barycenter is the optimal affine estimation

of the true barycenter. Furthermore, we provide the theoretical characterization and an explicit formula of the Pareto frontier on the Wasserstein space. Chapter 4 studies the exact solution to the optimal data representation and the optimal affine estimation of the exact solution. Section 4.6 proposes an algorithm based on the theoretical results in the previous sections. Section 4.7 provides an extensive numerical study regarding the application of the pseudo-barycenter and the optimal affine maps to (1) the estimation of optimal fair learning outcome compared to the known fair machine learning techniques on different learning models; and (2) Pareto frontier estimation for different disparity definitions. Chapter 5 studies the (in)compatibility between the optimal statistical parity solutions and individual fairness constraints. In the case of incompatibility, we identify which portion of the Pareto frontier is compatible with the individual fairness constraints. Chapter 6 provides the characterization of the optimal fair equalized odds machine learning via conditional Wasserstein barycenter, but also discusses the limited practical application use case. Finally, Chapter 7 discusses the remaining challenges in fairness machine learning that we hope to continue working on, as well as our future plans in broader AI ethics directions.

Preliminaries: Optimal Transport

In this chapter, we review the theoretical results on optimal transport and the Wasserstein barycenter that are important for the development of the main theoretical results on efficient algorithm design, Wasserstein geodesic characterization of the Pareto frontier, and the pre-processing approach resulting in the optimal fair data representation. For our purposes, we focus on \mathbb{R}^d . We refer readers who are interested in more generalized versions, e.g. on compact Riemannian manifolds, to for example [38].

2.1. General Distribution Case

Given $\mu, \nu \in \mathcal{P}(\mathbb{R}^d)$, which is the set of all probability measures on \mathbb{R}^d , Monge asked for an optimal transportation map $T_{\mu\nu} : \mathbb{R}^d \rightarrow \mathbb{R}^d$ that solves

$$(2.1) \quad \inf_{T_{\#}\mu=\nu} \left\{ \int_{\mathbb{R}^d} \|x - T(x)\|^2 d\mu \right\}$$

Here, $\|\cdot\|$ denotes the Euclidean norm on \mathbb{R}^d . The problem remained open until Brenier showed that Monge's problem coincides with Kantorovich's relaxed version:

$$(2.2) \quad \inf_{\gamma \in \Pi(\mu, \nu)} \left\{ \int_{\mathbb{R}^d \times \mathbb{R}^d} \|x_1 - x_2\|^2 d\gamma(x_1, x_2) \right\}$$

and admits a unique solution provided $\mu \in \mathcal{P}_{2,ac}(\mathbb{R}^d)$. Here, $\mathcal{P}_{2,ac}(\mathbb{R}^d)$ denotes the space of probability measures on \mathbb{R}^d that have finite first two moments and are absolutely continuous w.r.t. (with respect to) the Lebesgue measure. That is, the optimal solution to (2.2) has the form: $\gamma = (Id, T_{\mu\nu})_{\#}\mu$, where $T_{\mu\nu}$ solves (2.1). Here, $\Pi(\mu, \nu)$ denotes all the probability measures on $(\mathbb{R}^{2d}, \mathcal{B}(\mathbb{R}^d) \otimes \mathcal{B}(\mathbb{R}^d))$ such that the marginals are μ and ν . The relaxed problem is easy to solve due to the weak* compactness of $\Pi(\mu, \nu)$. We refer interested readers to [50, 51] for more detailed existence and uniqueness results.

REMARK 2.1.1. *The uniqueness is in the weak sense for γ and μ -a.e. for $T_{\mu\nu}$.*

Kantorovich's problem provides a certain kind of "distance" on $\mathcal{P}(\mathbb{R}^d)$ except for the possibility of being infinite.

DEFINITION 2.1.1 (Wasserstein distance¹). *Given $\mu, \nu \in \mathcal{P}(\mathbb{R}^d)$,*

$$(2.3) \quad \mathcal{W}_2(\mu, \nu) := \left(\inf_{\gamma \in \Pi(\mu, \nu)} \left\{ \int_{\mathbb{R}^d \times \mathbb{R}^d} \|x_1 - x_2\|^2 d\gamma(x_1, x_2) \right\} \right)^{\frac{1}{2}}$$

It is not hard to verify that the Wasserstein distance defined above satisfies the axioms of a metric except for finiteness of $\mathcal{W}_2(\mu, \nu)$ for arbitrary $\mu, \nu \in \mathcal{P}(\mathbb{R}^d)$. In order to guarantee finiteness, one needs to put more restrictions on the set of all probability measures:

DEFINITION 2.1.2 (Wasserstein space). *Define \mathcal{W}_2 as above and*

$$(2.4) \quad \mathcal{P}_2(\mathbb{R}^d) := \left\{ \mu \in \mathcal{P}(\mathbb{R}^d) : \int_{\mathbb{R}^d} \|x\|^2 d\mu < \infty \right\}.$$

The couple $(\mathcal{P}_2(\mathbb{R}^d), \mathcal{W}_2)$ is called Wasserstein space.

The Wasserstein space has gained increasing popularity in image processing, economics [16, 24], and machine learning in recent years due to its useful properties such as polishness (of the space) and robustness (w.r.t. perturbation on the marginal probability measures and hence on sampling). Since the Wasserstein space is a metric space, the Fréchet mean on the space is well-defined and it is called the Wasserstein barycenter in the optimal transport literature.

DEFINITION 2.1.3 (Wasserstein barycenter [2]). *Given $\{\mu_z\}_{z \in \mathcal{Z}} \subset (\mathcal{P}_2(\mathbb{R}^d), \mathcal{W}_2)$ for some index set \mathcal{Z} , the barycenter of $\{\mu_z\}_z$ is the Fréchet mean of the set on $(\mathcal{P}_2(\mathbb{R}^d), \mathcal{W}_2)$. That is, $\bar{\mu}$ is the solution to*

$$(2.5) \quad \inf_{\mu \in \mathcal{P}_2(\mathbb{R}^d)} \left\{ \int_{\mathcal{Z}} \mathcal{W}_2^2(\mu_z, \mu) d\lambda(z) \right\},$$

where $\bar{\mu}$ denotes the Fréchet mean or barycenter.

Here, for our purpose, we focus on the case where the index set $\mathcal{Z} \in \{[k], \mathbb{N}, [0, 1], \mathbb{R}^n\}$.

¹Throughout this dissertation we work with the Wasserstein-2 distance, and thus simply call it the Wasserstein distance.

Next, we look at optimal transport and the barycenter problem from the perspective of optimal coupling. The goal is to show that the multi-marginal coupling problem is equivalent to the Wasserstein barycenter problem. The equivalence is an essential tool in proving our result in optimal affine transport, the optimality of the pseudo-barycenter, and the geodesic characterization of the Pareto frontier.

First, notice that Kantorovich's problem is in fact a 2-marginal coupling problem: Let X_1, X_2 be the random variable satisfy $\mathcal{L}(X_1) = \mu, \mathcal{L}(X_2) = \nu$, the problem looks for a γ with marginals being μ, ν that minimizes $\mathbb{E}_\gamma \|X_1 - X_2\|^2$. It follows naturally by the existence and uniqueness result of the optimal transport map (also known as Brenier's map) [12], that the Wasserstein distance admits the form in the classic probability language:

$$(2.6) \quad \mathcal{W}_2(\mu, \nu) = (\mathbb{E}_\mu \|X_1 - T(X_1)\|^2)^{\frac{1}{2}},$$

where T is the optimal transport map that pushes $\mu = \mathcal{L}(X_1)$ forward to $\nu = \mathcal{L}(X_2)$.

More recent work in mathematics [38, 44] and economics [16, 24] has generalized the Kantorovich problem to the multi-marginal coupling problem:

$$(2.7) \quad \inf_{\gamma \in \Pi(\{\mu_z\}_{z \in \mathcal{Z}})} \left\{ \mathbb{E}_\gamma \left(\int_{\mathcal{Z}^2} \|X_{z_1} - X_{z_2}\|^2 d\lambda(z_1) d\lambda(z_2) \right) \right\},$$

where $\Pi(\{\mu_z\}_{z \in \mathcal{Z}})$ denotes all the Borel probability measures on $(\mathbb{R}^d)^{|\mathcal{Z}|}$ with marginals being $\mu_z = \mathcal{L}(X_z) \in \mathcal{P}(\mathbb{R}^d)$ λ -a.e.. Hence, one can consider $\lambda \in \mathcal{P}(\mathcal{P}(\mathbb{R}^d))$. It can be shown that the above is equivalent to the following:

$$(2.8) \quad \sup_{\gamma \in \Pi(\{\mu_z\}_{z \in \mathcal{Z}})} \left\{ \mathbb{E}_\gamma \left(\left\| \int_{\mathcal{Z}} X_z d\lambda(z) \right\|^2 \right) \right\}$$

REMARK 2.1.2 (Justification for the name of marginals). *Since $\{X_z\}_z$ are the marginals for the admissible couplings in (2.7), with the equivalence between the multi-marginal coupling and Wasserstein barycenter (see Remark 2.1.3 below) in mind, we often call $\{X_z\}_z$ and $\{\mathcal{L}(X_z)\}_z$ the sensitive marginals, even though they are also the conditional random variables and distributions constructed by disintegration.*

Intuitively, (2.8) tends to find a family of random variables parametrized by z with fixed marginals μ_z such that the variance of the matched (by γ) group average is maximized. For readers who are more familiar with stochastic processes, consider $z = t$ as a time variable, then X_t is a stochastic process with fixed time marginals, and (2.8) tends to find a way (γ) to group the fixed marginals into trajectories so that the variance of the trajectory-wise (sample path) average is maximized. (Hence, the expected variance within a randomly chosen sample path is minimized.)

As shown in [2, 44], the above multi-marginal problem is equivalent to the barycenter problem:

REMARK 2.1.3 (Equivalence between multi-marginal coupling and barycenter). *Assume $\{\mu_z\}_z$ are absolutely continuous w.r.t. the Lebesgue measure and let γ^* and $\bar{\mu}$ be the solution to (2.8) and (5.8), respectively. It follows that $\bar{\mu} = \gamma^* \circ T^{-1}$ where $T(\{x_z\}_z) := \int_{\mathcal{Z}} x_z d\lambda(z)$.*

The importance of this equivalence is twofold:

- 1 It is the key to proving the non-degenerate Gaussianity of the Wasserstein barycenter of non-degenerate Gaussian marginal distributions;
- 2 It provides technical support for the interpretation (Section ?? point 4) of how the Wasserstein barycenter solves data-related fairness issues on a point-wise scale.

Therefore, we generalize the equivalence to the case where \mathcal{Z} is a Polish space, which is a metric space that is separable and complete. In particular, $[k]^d, [0, l]^d, \mathbb{N}^d, \mathbb{R}^d$ mentioned above are all examples of Polish spaces. This generalization is important for our purpose as it provides a theoretical foundation for removing Z in the form of random vectors.

Now, the following result provides the existence and uniqueness result of the barycenter problem that is suitable for our purpose.

THEOREM 2.1.1 (Existence and uniqueness of barycenter [39](Theorem 2 and Proposition 6)). *Assume that \mathcal{Z} is a Polish space and that $\lambda := \mathbb{P} \circ Z^{-1}$ satisfies $\int_{\mathcal{Z}} \mathcal{W}_2^2(\mu_z, \nu) d\lambda(z) < \infty$ for some $\nu \in \mathcal{P}_2(\mathcal{X})$ (hence, for all $\nu \in \mathcal{P}_2(\mathcal{X})$). Then the following properties hold:*

- 1 *There exists a barycenter of $\{\mu_z\}_{z \in \mathcal{Z}}$ w.r.t. λ .*
- 2 *If, in addition, $\lambda(\{z : \mu_z \in \mathcal{P}_{ac}(\mathcal{X})\}) > 0$, then the barycenter is unique.*

REMARK 2.1.4 (Applicability of assumptions in Theorem 2.1.1). *The assumption that $\int_{\mathcal{Z}} \mathcal{W}_2^2(\mu_z, \nu) d\lambda(z) < \infty$ in the above result is satisfied in our application to the optimal fair learning outcome or data representation: When generating the optimal transport maps $\{T_z\}_z$, the training set has a finite number*

of data and hence finite different values of z in the discrete case or after discretization in the continuous case. Therefore, since $\{\mu_z\}_z \subset \mathcal{P}_2(\mathcal{X})$, pick a value z_0 that is in the training set, we have that $\mathcal{W}_2^2(\mu_z, \mu_{z_0})$ are essentially (w.r.t. λ) uniformly bounded. That implies $\int_{\mathcal{Z}} \mathcal{W}_2^2(\mu_z, \mu_{z_0}) d\lambda(z) < \infty$.

Now, we have the theoretical results that are needed to prove the main results, except for the McCann interpolation, which will be introduced in Section 4. The next step is to develop a computationally efficient method to compute (an estimation of) the Wasserstein barycenter, (the McCann interpolation of) optimal transport maps, and thereby the optimal fair model and Pareto frontier. More specifically, we focus on positive definite affine optimal transport maps.

2.1.1. Rigid Translation. Before deriving our main result on optimal positive definite affine maps, we first study the case where admissible maps are restricted to the set of rigid translations. The following property of rigid translations makes our results on the optimal affine maps simpler: we can assume, without loss of generality, that the first moments of the marginal measures are zero: $m_{X_z} := \mathbb{E}(X_z) = 0$ and $m_{Y_z} := \mathbb{E}(Y_z) = 0$.

LEMMA 2.1.1. *Let $\mu, \nu \in \mathcal{P}_2$, $m_\mu := \int x d\mu(x)$, and $m_\nu := \int x d\nu(x)$. Also, let μ', ν' be the centered versions of μ, ν , respectively. It follows that*

$$(2.9) \quad \mathcal{W}_2^2(\mu, \nu) = \mathcal{W}_2^2(\mu', \nu') + \|m_\mu - m_\nu\|^2.$$

PROOF.

$$\begin{aligned} \mathcal{W}_2^2(\mu, \nu) &= \int \|x - y\|^2 d\gamma^*(x, y) \\ &= \int \|((x - m_\mu) - (y - m_\nu)) + (m_\mu - m_\nu)\|^2 d\gamma^*(x, y) \\ &= \int \|(x - m_\mu) - (y - m_\nu)\|^2 d\gamma^*(x, y) + \|m_\mu - m_\nu\|^2 \\ &\geq \mathcal{W}_2^2(\mu', \nu') + \|m_\mu - m_\nu\|^2 \\ &= \int \|x - y\|^2 d(\gamma')^*(x, y) + \|m_\mu - m_\nu\|^2 \\ &= \int \|(x + m_\mu) - (y + m_\nu)\|^2 d(\gamma')^*(x, y) \\ &\geq \mathcal{W}_2^2(\mu, \nu) \end{aligned}$$

where γ^* and $(\gamma')^*$ denote the optimal transport plan for (μ, ν) and (μ', ν') , respectively. The first inequality results from the fact that $\gamma'(x, y) := \gamma^*(x - m_\mu, y - m_\nu) \in \Pi(\mu', \nu')$, the second inequality from $\gamma(x, y) := (\gamma')^*(x + m_\mu, y + m_\nu) \in \Pi(\mu, \nu)$, and the equalities from direct expansion. \square

Notice that the above result allows us to assume measures to have vanishing first moments when deriving the optimal transport maps. Indeed, if $T_{\mu'\nu'}$ is the Brenier's map between μ' and ν' , then $T_{\mu\nu} := T_{+m_\nu} \circ T_{\mu'\nu'} \circ T_{-m_\mu}$ is the optimal transport map between μ and ν . Here, $T_{+m_\nu}(x) := x + m_\nu$ and T_{-m_μ} are defined analogously.

In the rest of Section ??, we assume without loss of generality that the first moments of the measures are all equal to zero.

2.2. Location-Scale Case and Optimal Affine Transport

A sufficient condition for Brenier's maps to be positive definite affine is to require a certain “similarity” between the marginal data distributions. One natural choice is to assume $\{Y_z\}_z$ and $\{X_z\}_z$ to be non-degenerate Gaussian vector λ -a.e.. As shown in [4], the assumptions of Gaussian vector can easily be generalized to a *location-scale family*. In the definition below, \mathcal{S}_{++}^d denotes the set of all $d \times d$ positive definite matrices.

The generalization from Gaussian to location-scale families is important for the main result in the next section, where we consider computationally efficient solutions to a relaxation of the Wasserstein barycenter problem in the case of general marginal distributions.

DEFINITION 2.2.1 (Location-Scale Family). *For any $\mathcal{L}(X_0) \in \mathcal{P}(\mathbb{R}^d)$, define*

$$(2.10) \quad \mathcal{F}(\mathcal{L}(X_0)) := \{\mathcal{L}(AX_0 + m) : A \in \mathcal{S}_{++}^d, m \in \mathbb{R}^d\}.$$

The set $\mathcal{F}(\mathcal{L}(X_0))$ is called a location-scale family characterized by $\mathcal{L}(X_0)$.

In other words, under the assumption of vanishing first moments, the random variables that share laws in the same location-scale family can be transformed into each other by a positive definite linear transformation.

In [4] it is shown that Brenier's map between two probability measures, each having a vanishing first moment, within the same location-scale family is linear and has a closed form.

LEMMA 2.2.1 (Optimal affine map). *If $\mu, \nu \in \mathcal{F}(\mathcal{L}(X_0))$ for some X_0 such that $m_\mu = m_\nu = 0$, then the Brenier's map that pushes μ forward to ν is given by:*

$$(2.11) \quad T_{\mu\nu} = \Sigma_\mu^{-\frac{1}{2}} (\Sigma_\mu^{\frac{1}{2}} \Sigma_\nu \Sigma_\mu^{\frac{1}{2}})^{\frac{1}{2}} \Sigma_\mu^{-\frac{1}{2}}$$

where $\Sigma_\mu := \int xx^T d\mu$ and $\Sigma_\nu := \int xx^T d\nu$.

PROOF. See, for example, Theorem 2.3 in [4]. □

REMARK 2.2.1. *The optimal affine map is also the midpoint of the geodesic path joining Σ_μ^{-1} and Σ_ν on the manifold of positive definite matrices. We refer interested readers to, for example, Chapter 6.1 in [9] for more details.*

Now, back to the barycenter problem. It follows from Lemma 2.2.1 that, if one assumes that all the marginals belong to the same location-scale family, then the barycenter also belongs to the family and a nearly closed-form solution to the barycenter is available.

LEMMA 2.2.2 (Barycenter in the location-scale case). *Assume $\{\mu_z\}_z$ belong to the same location-scale family $\mathcal{F}(P_0)$ and satisfy $m_{\mu_z} = 0, \Sigma_{\mu_z} \succ 0, \lambda - a.e.$, then there exists a unique solution, denoted by $\bar{\mu}$, to (5.8). Moreover, $\bar{\mu}$ also belongs to $\mathcal{F}(P_0)$ and is characterized by $m_{\bar{\mu}} = 0$ and $\Sigma_{\bar{\mu}} = \Sigma$ where Σ is the unique solution to the following equation:*

$$(2.12) \quad \int_{\mathcal{Z}} (\Sigma^{\frac{1}{2}} \Sigma_{\mu_z} \Sigma^{\frac{1}{2}})^{\frac{1}{2}} d\lambda(z) = \Sigma,$$

where Σ_{μ_z} is the second moment of $\mu_z, \forall z \in \mathcal{Z}$.

PROOF. Existence and uniqueness follow directly from Theorem 2.1.1. For the equivalent multi-marginal coupling problem, there exists an optimal solution $\gamma^* = \mathcal{L}(\{X_z\}_z)$. It follows from Remark 2.1.3 that $\bar{X} = T(\{X_z\}_z)$ where $\mathcal{L}(\bar{X})$ is the Wasserstein barycenter. Therefore, the Gaussianity of barycenter results from linearity of T in the finite $|\mathcal{Z}|$ case, and the fact that the set of Gaussian distribution is closed in $(\mathcal{P}_{2,ac}, \mathcal{W}_2)$ when $|\mathcal{Z}|$ is infinite. The characterization equation is proved in the case of finite $|\mathcal{Z}|$ in [2]. For infinite $|\mathcal{Z}|$, the equation still holds due to the continuity of the covariance function on $(\mathcal{P}_{2,ac}, \mathcal{W}_2)$. The sufficiency and necessity of the equation follows from the

following characterization of the barycenter via Brenier's maps $\{T_{\bar{X}X_z}\}_z$ derived in [2]:

$$(2.13) \quad \int_{\mathcal{Z}} T_{\bar{X}X_z} d\lambda(z) = Id.$$

It follows from the explicit form of $\{T_{\bar{X}X_z}\}_z$ in Lemma 2.2.1 that

$$\begin{aligned} \int_{\mathcal{Z}} T_{\bar{X}X_z} d\lambda(z) &= \int_{\mathcal{Z}} \Sigma_{\bar{X}}^{-\frac{1}{2}} (\Sigma_{\bar{X}}^{\frac{1}{2}} \Sigma_{X_z} \Sigma_{\bar{X}}^{\frac{1}{2}})^{\frac{1}{2}} \Sigma_{\bar{X}}^{-\frac{1}{2}} d\lambda(z) = Id \\ \iff \Sigma_{\bar{X}}^{\frac{1}{2}} \Sigma_{\bar{X}}^{-\frac{1}{2}} \int_{\mathcal{Z}} (\Sigma_{\bar{X}}^{\frac{1}{2}} \Sigma_{X_z} \Sigma_{\bar{X}}^{\frac{1}{2}})^{\frac{1}{2}} d\lambda(z) \Sigma_{\bar{X}}^{-\frac{1}{2}} \Sigma_{\bar{X}}^{\frac{1}{2}} &= \Sigma_{\bar{X}}^{\frac{1}{2}} Id \Sigma_{\bar{X}}^{\frac{1}{2}} \\ \iff \int_{\mathcal{Z}} (\Sigma_{\bar{X}}^{\frac{1}{2}} \Sigma_{X_z} \Sigma_{\bar{X}}^{\frac{1}{2}})^{\frac{1}{2}} d\lambda(z) &= \Sigma_{\bar{X}}. \end{aligned}$$

□

In the case where $m_{\mu_z} \neq 0$, it follows from Lemma 2.1.1 that

$$\int_{\mathcal{Z}} \mathcal{W}_2^2(\mu_z, \mu) d\lambda(z) = \int_{\mathcal{Z}} \mathcal{W}_2^2(\mu'_z, \mu') d\lambda(z) + \int_{\mathcal{Z}} \|m_{\mu_z} - m_{\mu}\|^2 d\lambda(z)$$

where μ' denotes the centered version of μ . By Lemma 2.2.2, we know the first term on the right is minimized at $\bar{\mu}' \sim \mathcal{N}(0, \Sigma_{\bar{\mu}})$. Also, the second term on the right is minimized at the Fréchet mean with Euclidean metric, which is equal to the expectation. That is, $m_{\bar{\mu}} = \int_{\mathcal{Z}} m_{\mu_z} d\lambda(z)$. As a result, the optimal transport map is

$$(2.14) \quad T_{\mu_z \bar{\mu}} = T_{+m_{\bar{\mu}}} \circ T_{\mu'_z \bar{\mu}'} \circ T_{-m_{\mu_z}}$$

REMARK 2.2.2 (Solution to (2.12)). *The non-linear matrix equation (2.12) has a unique solution that can be approached via the following iterative process:*

$$(2.15) \quad \int_{\mathcal{Z}} (\Sigma_i^{\frac{1}{2}} \Sigma_{\mu_z} \Sigma_i^{\frac{1}{2}})^{\frac{1}{2}} d\lambda(z) \rightarrow \Sigma_{i+1}.$$

We refer interested readers to [4] for more details on the fixed point approach to the Wasserstein barycenter. The present work only applies this fact in the algorithm design in Section 4.6.

Pareto Frontier for L^2 -objective Machine Learning

Optimal transport has been considered an adversarial or constrained optimization problem in its application to machine learning. In particular, some of the most popular unsupervised learning methods, such as K-means and PCA, are specific examples of the Wasserstein barycenter problems when putting restrictions on the admissible transport maps and relaxation on the weak equivalence requirement of the push-forwards w.r.t. test functions. See, for example, [49] for more details. But we apply optimal transport in an opposite direction so that the independence or imperceptibility of the sensitive variable Z becomes theoretically provable.

In this chapter, the primary goal is to develop the optimal affine map and pseudo-barycenter as tools to solve the challenge of the high computational cost of Wasserstein barycenter and optimal transport maps in high-dimensional data space. More specifically, we restrict the admissible transport maps to be merely affine maps while relaxing the fairness constraint to a sufficient and necessary level. The importance of efficiency in computing the barycenter and optimal transport maps will soon be clear in Section 3.3 when we compute the Pareto frontier along the Wasserstein geodesic path. Furthermore, the importance of affinity of transport maps will also be soon clear in Section 4.2 when solving the optimal fair data representation problem (1.3).

The organization of the current chapter is as follows: we first generalize the Wasserstein barycenter characterization of the optimal regression to all L^2 -objective supervised learning models, then apply the optimal affine maps to estimate high-dimension optimal learning outcome.

3.1. Wasserstein Barycenter Characterization of Optimal Fair Learning

Now, we show that the (unique) solution to Problem 1 can be characterized as the Wasserstein barycenter of the conditional expectation sensitive marginals. The barycenter characterization of the optimal fair regression is first proved in [19, 29].

We start with a characterization of the optimal learning outcome of the L^2 -objective supervised learning task. Let $\mathbb{E}(Y|X, Z)_z$ be the sensitive marginals of $(\mathbb{E}(Y|X, Z), Z)$ (or, equivalently, the sensitive conditionals of $\mathbb{E}(Y|X, Z)$ on $\{Z = z\}$ by Remark 2.1.2) for λ -a.e. $z \in \mathcal{Z}$, $\mathcal{L}(\mathbb{E}(Y|X, Z)_z) := \mu_z$, and $\bar{\mu}$ denote the Wasserstein barycenter of $\{\mu_z\}_{z \in \mathcal{Z}}$. Also, let $T(\cdot, z)$ denote the optimal transport map from μ_z to $\bar{\mu}$.

LEMMA 3.1.1 (Optimal fair L^2 -objective supervised learning characterization). *Assume that the conditional expectation marginals $\{\mu_z\}_{z \in \mathcal{Z}} \subset \mathcal{P}_{2,ac}(\mathcal{Y})$, then*

$$(3.1) \quad \overline{\mathbb{E}(Y|X, Z)} = T(\mathbb{E}(Y|X, Z), Z) := \{T(\mathbb{E}(Y|X, Z)_z, z)\}_{z \in \mathcal{Z}}$$

is the unique solution to Problem 1. Furthermore, we have

$$\begin{aligned} \|Y - T(\mathbb{E}(Y|X, Z), Z)\|_2^2 &= \inf_{f \in L^2(\mathcal{X} \times \mathcal{Z}, \mathcal{Y})} \{\|Y - f(X, Z)\|_2^2 : f(X, Z) \perp Z\} \\ &= \|Y - \mathbb{E}(Y|X, Z)\|_2^2 + \int_{\mathcal{Z}} \mathcal{W}_2^2(\mu_z, \bar{\mu}) d\lambda \end{aligned}$$

PROOF. First, notice that the fairness constraint $f(X, Z) \perp Z$ is equivalent to $\mathcal{L}(f(X, Z)_z) = \mu$ λ -a.e. for some $\mu \in \mathcal{P}(\mathcal{Y})$. Now, we prove the lower bound: let $f \in L^2(\mathcal{X} \times \mathcal{Z}, \mathcal{Y})$ satisfies $f(X, Z) \perp Z$, we have

$$\begin{aligned} \|Y - f(X, Z)\|_2^2 &= \|Y - \mathbb{E}(Y|X, Z)\|_2^2 + \|\mathbb{E}(Y|X, Z) - f(X, Z)\|_2^2 \\ &= \|Y - \mathbb{E}(Y|X, Z)\|_2^2 + \int_{\mathcal{Z}} \|\mathbb{E}(Y|X, Z)_z - f(X, Z)_z\|_2^2 d\lambda \\ &\geq \|Y - \mathbb{E}(Y|X, Z)\|_2^2 + \int_{\mathcal{Z}} \mathcal{W}_2^2(\mu_z, \mathcal{L}(f(X, Z)_z)) d\lambda \\ &\geq \|Y - \mathbb{E}(Y|X, Z)\|_2^2 + \int_{\mathcal{Z}} \mathcal{W}_2^2(\mu_z, \bar{\mu}) d\lambda \end{aligned}$$

Here, the first line follows from the L^2 projection characterization of conditional expectation, the second follows from disintegration, the third from the definition of \mathcal{W}_2 , and the fourth from the definition of the Wasserstein barycenter and the fairness restriction $f(X, Z) \perp Z$.

Next, we construct a $f_Y \in L^2(\mathcal{X} \times \mathcal{Z}, \mathcal{Y})$ such that the lower bound is obtained. Let T_z denote the optimal transport map such that $(T_z)_\# \mu_z = \bar{\mu}$ for λ -a.e. $z \in \mathcal{Z}$. Define $T(\cdot, z) := T_z(\cdot)$ and

$$(3.2) \quad f_Y(X, Z) := T(\mathbb{E}(Y|X, Z), Z).$$

Here, $\pi := (Id, T_z)_\# \mu_z d\lambda$ defines $\pi \in \mathcal{P}(\mathcal{Z} \times \mathcal{Y} \times \mathcal{Y})$. Hence, we have $\pi = \pi_{(y,z)} d\lambda_{(\mathbb{E}(Y|X,Z), Z)}$ and $\pi_{(y,z)} = \delta_{T(y,z)} \lambda_{(\mathbb{E}(Y|X,Z), Z)} - a.e.$. Since $(y, z) \rightarrow \pi_{(y,z)} = \delta_{T(y,z)}$ is $\mathcal{Y} \times \mathcal{Z} / \mathcal{P}(\mathcal{Y})$ measurable, we have $(y, z) \rightarrow T(y, z)$ is $\mathcal{Y} \times \mathcal{Z} / \mathcal{Y}$ measurable. It follows from $\mathbb{E}(Y|\cdot, \cdot) \otimes Id|_{\mathcal{Z}}(\cdot, \cdot) : \mathcal{X} \times \mathcal{Z} \rightarrow \mathcal{Y} \times \mathcal{Z}$ being $\mathcal{X} \times \mathcal{Z} / \mathcal{Y} \times \mathcal{Z}$ measurable that $f_Y = T \circ (\mathbb{E}(Y|\cdot, \cdot) \otimes Id|_{\mathcal{Z}}(\cdot, \cdot))$ is $\mathcal{X} \times \mathcal{Z} / \mathcal{Y}$ measurable. Also, $\bar{\mu} \in \mathcal{P}_2(\mathcal{Y}) \implies \|f_Y(X, Z)\|_2 < \infty$. This proves $f_Y \in L^2(\mathcal{X} \times \mathcal{Z}, \mathcal{Y})$. It remains to show that the lower bound is obtained at $f_Y(X, Z)$. Indeed, by construction, we have

$$\begin{aligned} \|\mathbb{E}(Y|X, Z) - f_Y(X, Z)\|_2^2 &= \int_{\mathcal{Z}} \|\mathbb{E}(Y|X, Z)_z - f_Y(X, Z)_z\|_2^2 d\lambda \\ &= \int_{\mathcal{Z}} \|\mathbb{E}(Y|X, Z)_z - T(\mathbb{E}(Y|X, Z)_z, z)\|_2^2 d\lambda \\ &= \int_{\mathcal{Z}} \mathcal{W}_2^2(\mu_z, (T_z)_\# \mu_z) d\lambda \\ &= \int_{\mathcal{Z}} \mathcal{W}_2^2(\mu_z, \bar{\mu}) d\lambda. \end{aligned}$$

It follows from the derivation of the lower bound above that

$$(3.3) \quad \|Y - f_Y(X, Z)\|_2^2 = \inf_{f \in L^2(\mathcal{X} \times \mathcal{Z}, \mathcal{Y})} \{\|Y - f(X, Z)\|_2^2 : f(X, Z) \perp Z\}$$

Uniqueness follows from the uniqueness of $\bar{\mu}$ and the uniqueness of $T(\cdot, z)$. We are done. \square

The above result shows that the minimum L^2 -loss for statistical parity can be nicely decomposed into two parts: (1) an $L^2(\mathcal{X} \times \mathcal{Z}, \mathcal{Y})$ orthogonal projection loss due to the inference capability of (X, Z) w.r.t. Y and (2) an independence projection loss due to the statistical parity constraint.

That is,

$$\underbrace{\inf_f \{\|Y - f(X, Z)\|_2^2 : f(X, Z) \perp Z\}}_{\text{minimum loss for statistical parity}} = \underbrace{\|Y - \mathbb{E}(Y|X, Z)\|_2^2}_{\text{orthogonal projection loss}} + \underbrace{\int_{\mathcal{Z}} \mathcal{W}_2^2(\mu_z, \bar{\mu}) d\lambda}_{\text{independence projection loss}}.$$

Furthermore, to construct the optimal fair L^2 learning outcome, one first performs L^2 orthogonal projection to obtain the conditional expectation $\mathbb{E}(Y|X, Z)$, then outputs the Wasserstein barycenter of the sensitive marginals of $\mathbb{E}(Y|X, Z)$ as the optimal (with respect to L^2 -objective) fair (for statistical parity) result.

Unfortunately, in practice, the characterization suffers from a lack of efficient methods to compute the Wasserstein barycenter and obtain an explicit formula of the optimal transport maps [3]. Current methods restrict the sensitive variable Z to be binary mainly because the computation of a multi-marginal barycenter is expensive. Furthermore, notice the current methods restrict the dependent variable Y to be one-dimensional, because the only well-known exact solution to transport maps is the inverse of cumulative function that merely works for one-dimensional variables.

Therefore, to provide methods using the characterization in high-dimensional dependent variable cases, we introduce the optimal affine map and the associated pseudo-barycenter.

3.2. Optimal Affine Estimation of Barycenter

To solve the challenge of deriving an explicit formula for the Wasserstein barycenter and optimal transport maps, we restrict the admissible transport maps to be affine and show that the estimation of the Wasserstein barycenter via optimal affine maps coincides with the true Wasserstein barycenter in the Gaussian case, and that the estimation error is bounded in the case of general distributions. In other words, we consider the choice of positive definite affine maps under two circumstances:

- 1 We assume the marginals are non-degenerate Gaussian. That is, $\{\mathbb{E}(Y|X, Z)_z\}_z$ are assumed to be non-degenerate Gaussian vectors λ -a.e..
- 2 Instead of making assumptions on the data distribution, we relax the independence constraint to the independence between Z and merely the first two moments of $f(X, Z)$.

From a theoretical perspective, affine maps allow us to derive (nearly) closed-form solutions under either of the assumptions mentioned above. Also, affine maps allow us to develop a pre-processing approach by directly applying the obtained maps to the original data before training, even though such maps are constructed to push the post-training marginals toward their barycenter.

From a practical perspective, the advantage is obvious: the computation of affine maps only uses (sample estimation of) the first two moments of the marginal distributions and hence is highly efficient compared to the computation of general Brenier’s maps, especially in the case of high-dimension data.

Before developing the pseudo-barycenter, the following remarks compare in more detail the exact barycenter with its affine approximation.

REMARK 3.2.1 (Applying pseudo-barycenter vs exact barycenter). *The comparison between the pseudo-barycenter method and the exact barycenter is an analog of the comparison between the linear regression model and the exact conditional expectation: When there is no worry about over-fitting, a practitioner who cares more about the strict goal of minimizing L^2 error (analog: the strict statistical parity guarantee) should always try to find the exact conditional expectation function (analog: the exact barycenter and the corresponding exact transport maps) by using more complicated models. But the simplicity, robustness, and interpretability of linear regression (analog: pseudo-barycenter and optimal affine maps) are often useful in practice.*

We define the pseudo-barycenter, using merely matrix calculations, as follows:

DEFINITION 3.2.1. *The post-processing pseudo-barycenter \hat{Y}^\dagger is given via*

$$(3.4) \quad \hat{Y}^\dagger := T_{\text{affine}}(\hat{Y}, Z),$$

where

$$(3.5) \quad T_{\text{affine}}(\cdot, z) := \Sigma_{\hat{Y}_z}^{-\frac{1}{2}} (\Sigma_{\hat{Y}_z}^{\frac{1}{2}} \Sigma \Sigma_{\hat{Y}_z}^{\frac{1}{2}})^{\frac{1}{2}} \Sigma_{\hat{Y}_z}^{-\frac{1}{2}},$$

and Σ is the unique solution to

$$(3.6) \quad \int_{\mathcal{Z}} (\Sigma^{\frac{1}{2}} \Sigma_{\hat{Y}_z} \Sigma^{\frac{1}{2}})^{\frac{1}{2}} d\lambda(z) = \Sigma.$$

To obtain (an approximation of) the unique solution, we apply the iterative method (2.15) in Remark 2.2.2 when designing our algorithm in Section 4.6.

Now, Lemma 2.2.1 shows that under the assumption of Gaussianity of the learning outcome marginals, the optimal transport map is affine and the pseudo-barycenter is indeed the Wasserstein barycenter. Moreover, Lemma 2.2.2 shows that the barycenter of Gaussian marginals is still Gaussian. Therefore, the optimal maps from the marginals to the barycenter are determined entirely by the first two moments.

LEMMA 3.2.1 (Post-processing pseudo-barycenter in the Gaussian case). *Assume $\hat{Y}_z \sim \mathcal{N}(0, \Sigma_z)$ for λ -a.e. $z \in \mathcal{Z}$, then \hat{Y}^\dagger is the Wasserstein barycenter of $\{\hat{Y}_z\}_z$.*

It follows from Theorem 3.2.1 that, if $\hat{Y} = \mathbb{E}(Y|X, Z)$, then Y^\dagger is the solution to the Wasserstein barycenter characterization of the optimal fair learning outcome.

Finally, we show that the pseudo-barycenter is the optimal affine estimation of the Wasserstein barycenter in the case of general marginal distributions. To do so, we need to first put restrictions on the admissible transport maps. However, such a restriction on admissible maps leads to a necessary relaxation of the fairness constraint. To see the necessity, Lemma 2.2.1 shows positive definite affine maps transform distributions within the same location-scale family. Therefore, given marginals Y_1 and Y_2 from different location-scale families, affine maps are not able to transform them to each other. That implies the non-existence of the barycenter under the original independence restriction. Indeed, if a barycenter of $\{Y_z\}_{z \in \{1,2\}}$ exists under the restriction of positive definite affine maps, then Y_1 and Y_2 belong to the same location-scale family as their barycenter, which contradicts the assumption of general distributions. That is, the Wasserstein barycenter characterization does not have a solution when we admit merely affine transport maps in the general marginal distribution case.

On the other hand, notice that the best affine maps can achieve is to map Y_1 to a Y_2' , which shares the same first two moments with Y_2 within the Y_1 location-scale family. We call such Y_2' a Y_1 location-scale family analog of Y_2 . Therefore, we propose the following relaxation of the fairness constraint that suffices to guarantee the existence of a solution to the relaxed version of (5.5) with merely positive definite affine transport maps:

$$(3.7) \quad m_{f(X,Z)}, \Sigma_{f(X,Z)} \perp Z$$

where $m_{f(X,Z)}$, and $\Sigma_{f(X,Z)}$, denotes respectively the first and second moment of $f(X, Z)$.

REMARK 3.2.2 (Fairness guarantee of the relaxation). *The adversarial task of testing and exploiting probabilistic independence between $f(X, Z)$ and Z is equivalently difficult to enforcing the independence. One common strategy is to explore its equivalence to the independence between all moments of $f(X, Z)$ and Z , provided the boundedness of the two random variables. But the verification or enforcement of independence among higher moments is extremely vulnerable to data noise in practice. Thus, instead of enforcing $f(X, Z) \perp Z$, one could relax the constraint to the independence between Z and some of the moments of $f(X, Z)$. In this section, we focus on the first two moments. That is, $m_{f(X,Z)}, \Sigma_{f(X,Z)}$ where $m_{f(X,Z)} := \mathbb{E}(f(X, Z))$ and*

$\Sigma_{f(X,Z)} := \mathbb{E}((f(X,Z) - \mathbb{E}(f(X,Z)))(f(X,Z) - \mathbb{E}(f(X,Z)))^T)$. It is not hard to notice that the relaxation is already strong enough to result in imperceptibility to any unsupervised learning algorithm that uses merely the mean and covariance of data to extract information, such as K-means and PCA.

Therefore, the optimal affine estimation of the Wasserstein barycenter characterization is given by:

PROBLEM 4 (Optimal affine estimation of barycenter problem).

$$(3.8) \quad \inf_{f \in L^2(\mathcal{X} \times \mathcal{Z}, \mathcal{Y})} \{ \|Y - f(X, Z)\|_2^2 : m_{f(X,Z)}, \Sigma_{f(X,Z)} \perp Z \}.$$

Now, we show that the pseudo-barycenter defined above is indeed the solution to Problem 4 and hence the optimal affine estimate of the optimal fair learning outcome. To prove the main result, we need the following result: given any fixed covariance matrix, the optimal positive definite affine maps result in the lowest Wasserstein distance such that the push-forwards all share the same fixed covariance matrix. To simplify notation, let $\mu_z := \mathcal{L}(\mathbb{E}(Y|X, Z)_z)$. Also, let $m_{Y|X_z}$ and $\Sigma_{Y|X_z}$ denote the mean and covariance matrix of $\mathbb{E}(Y|X, Z)_z$ respectively.

LEMMA 3.2.2 (Projection Lemma). *Assume $\{\mu_z\}_z \subset \mathcal{P}_{2,ac}(\mathcal{Y})$. If $m_{Y|X_z} = 0, \Sigma_{Y|X_z} \succ 0$ λ -a.e., for any $\Sigma \succ 0$,*

$$(3.9) \quad \inf_{\hat{Y} : \Sigma_{\hat{Y}_z} = \Sigma} \int_{\mathcal{Z}} \mathcal{W}_2^2(\mu_z, \mathcal{L}(\hat{Y}_z)) d\lambda(z)$$

admits a unique solution, denoted by \hat{Y}_Σ , that satisfies

$$(3.10) \quad \hat{Y}_{\Sigma,z} := T_\Sigma(\hat{Y}_z, z)$$

where $T_\Sigma(\cdot, z) := \Sigma_{Y|X_z}^{-\frac{1}{2}} (\Sigma_{Y|X_z}^{\frac{1}{2}} \Sigma \Sigma_{Y|X_z}^{\frac{1}{2}})^{\frac{1}{2}} \Sigma_{Y|X_z}^{-\frac{1}{2}}$.

PROOF.

$$\begin{aligned} \int_{\mathcal{Z}} \mathcal{W}_2^2(\mu_z, \mathcal{L}(\hat{Y}_z)) d\lambda(z) &= \int_{\mathcal{Z}} \|\mathbb{E}(Y|X, Z)_z - T_\Sigma(\hat{Y}_z, z)\|_2^2 d\lambda(z) \\ &= \int_{\mathcal{Z}} \inf_{\nu : \Sigma_\nu = \Sigma} \mathcal{W}_2^2(\mu_z, \nu) d\lambda(z) \\ &= \inf_{\nu : \Sigma_{\nu_z} = \Sigma} \int_{\mathcal{Z}} \mathcal{W}_2^2(\mu_z, \nu_z) d\lambda(z), \end{aligned}$$

where the second equality follows from the characterization of Gelbrich's bound, see for example Proposition 2.4 in [22]. Now, let $\hat{Y}' \neq \hat{Y}_\Sigma$ but also satisfy $\Sigma_{\hat{Y}'} = \Sigma$ λ -a.e., then we have

$$\begin{aligned} \int_{\mathcal{Z}} \|\mathbb{E}(Y|X, Z)_z - \hat{Y}_{\Sigma, z}\|_2^2 d\lambda(z) &< \int_{\mathcal{Z}} \mathcal{W}_2^2(\mu_z, \mathcal{L}(\hat{Y}'_z)) d\lambda(z) \\ &\leq \int_{\mathcal{Z}} \|\mathbb{E}(Y|X, Z)_z - \hat{Y}'_z\|_2^2 d\lambda(z), \end{aligned}$$

where the first inequality is strict due to the uniqueness of Brenier's maps $T_\Sigma(\cdot, z)$ and hence of $T_\Sigma(\hat{Y}_z, z)$ λ -a.e.. The proof is complete. \square

REMARK 3.2.3 (Intuition of the Projection Lemma). *Intuitively, for an arbitrary positive definite matrix Σ , one can consider $T_\Sigma(\cdot, z)$ as the projection map (w.r.t. \mathcal{W}_2 distance) onto*

$$(3.11) \quad \{\nu \in \mathcal{P}_2(\mathcal{Y}) : \Sigma_\nu = \Sigma\}$$

which is the set of centered probability measures with fixed covariance matrix Σ in $(\mathcal{P}_2(\mathcal{Y}), \mathcal{W}_2)$. In other words, given a probability measure, the maps $\{T_\Sigma(\cdot, z)\}_z$ finds the closest (w.r.t. the Wasserstein distance) point in the set for each of the marginals.

Finally, we are ready to prove the justification of the pseudo-barycenter in the case of general distributions.

THEOREM 3.2.1 (Optimal affine estimation of \mathcal{W}_2 barycenter: Pseudo-barycenter). $\mathbb{E}(Y|X, Z)^\dagger := \{T_{\text{affine}}(\mathbb{E}(Y|X, Z)_z, z)\}_z$ is the unique solution to Problem 4:

$$(3.12) \quad \inf_{f \in L^2(\mathcal{X} \times \mathcal{Z}, \mathcal{Y})} \{\|Y - f(X, Z)\|_2^2 : m_{f(X, Z)}, \Sigma_{f(X, Z)} \perp Z\},$$

provided $\{\mu_z\}_z \subset \mathcal{P}_{2, \text{ac}}(\mathcal{Y})$.

PROOF. First, we fix $\Sigma \succ 0$ arbitrary and denote $\hat{Y}_{\Sigma, z} := T_\Sigma(\mathbb{E}(Y|X, Z)_z, z)$ for λ -a.e. $z \in \mathcal{Z}$, we have

$$(3.13) \quad \|Y - T_\Sigma(\mathbb{E}(Y|X, Z), Z)\|_2^2 - \|Y - \mathbb{E}(Y|X, Z)\|_2^2 = \int_{\mathcal{Z}} \|\mathbb{E}(Y|X, Z)_z - \hat{Y}_{\Sigma, z}\|_2^2 d\lambda(z)$$

and it follows from Lemma 3.2.2 that

$$\begin{aligned} \int_{\mathcal{Z}} \|\mathbb{E}(Y|X, Z)_z - \hat{Y}_{\Sigma, z}\|_2^2 d\lambda(z) &= \int_{\mathcal{Z}} \mathcal{W}_2^2(\mu_z, \mathcal{L}(T_{\Sigma}(\mathbb{E}(Y|X, Z)_z, z))) d\lambda(z) \\ &= \min_{\nu: \Sigma_{\nu_z} = \Sigma} \int_{\mathcal{Z}} \mathcal{W}_2^2(\mu_z, \nu_z) d\lambda(z). \end{aligned}$$

Therefore, (3.8) boils down to the following:

$$(3.14) \quad \inf_{\Sigma \succ 0} \left\{ \int_{\mathcal{Z}} \|\mathbb{E}(Y|X, Z)_z - T_{\Sigma}(\mathbb{E}(Y|X, Z)_z, z)\|_2^2 d\lambda(z) \right\}.$$

Finally, notice that

$$\begin{aligned} & \int_{\mathcal{Z}} \|\mathbb{E}(Y|X, Z)_z - T_{\Sigma}(\mathbb{E}(Y|X, Z)_z, z)\|_2^2 d\lambda(z) \\ &= \int_{\mathcal{Z}} \|\mathbb{E}(Y|X, Z)_z\|_2^2 + \|T_{\Sigma}(\mathbb{E}(Y|X, Z)_z, z)\|_2^2 - 2\langle \mathbb{E}(Y|X, Z)_z, T_{\Sigma}(\mathbb{E}(Y|X, Z)_z, z) \rangle_2 d\lambda(z) \\ &= \int_{\mathcal{Z}} \text{Trace}(\Sigma_{Y|X_z}) + \text{Trace}(\Sigma) - 2\mathbb{E}(\mathbb{E}(Y|X, Z)_z^T T_{\Sigma}(\mathbb{E}(Y|X, Z)_z, z)) d\lambda(z) \\ &= \int_{\mathcal{Z}} \text{Trace}(\Sigma_{Y|X_z}) + \text{Trace}(\Sigma) - 2\langle T_{\Sigma}, \Sigma_{Y|X_z} \rangle_F d\lambda(z) \\ &= \int_{\mathcal{Z}} \|\mathbb{E}(Y|X, Z)_z' - T_{\Sigma}(\mathbb{E}(Y|X, Z)_z', z)\|_2^2 d\lambda(z), \end{aligned}$$

where $\langle \cdot, \cdot \rangle_F$ denotes the Frobenius inner product and $X' \sim \mathcal{N}(m_X, \Sigma_X)$ denotes the Gaussian analog of X . It follows from definition of $T_{\text{affine}}(\mathbb{E}(Y|X, Z)_z, z)$ with

$$T_{\text{affine}}(\cdot, z) := \Sigma_{Y|X_z}^{-\frac{1}{2}} (\Sigma_{Y|X_z}^{\frac{1}{2}} \Sigma \Sigma_{Y|X_z}^{\frac{1}{2}})^{\frac{1}{2}} \Sigma_{Y|X_z}^{-\frac{1}{2}}$$

and Lemma 2.2.2 that $\int_{\mathcal{Z}} \|\mathbb{E}(Y|X, Z)_z - \mathbb{E}(Y|X, Z)_z^{\dagger}\|_2^2 d\lambda(z)$ is the unique lower bound of the objective function in (3.14). It then follows from the uniqueness of Brenier's map that $\mathbb{E}(Y|X, Z)^{\dagger}$ is the unique solution to (3.8). \square

In this section, we focus on applying the optimal affine transport map and the pseudo-barycenter to find a computationally efficient estimation of the optimal fair learning outcome in high-dimensional space. As we mentioned above, it will soon become clear in the next two sections and numerical experiments that a combination of McCann interpolation and optimal affine maps in matrix form results in not only a mathematically neat solution to estimate the Pareto frontier, which significantly

reduces computational expense in practice, but also a necessary tool to help us circumvent the post-processing nature and solve the optimal fair data representation problem (1.3).

Now, we are ready to address the lack of a precise theoretical characterization of the Pareto frontier between utility and fairness, which turns out to be a natural generalization of the Wasserstein barycenter characterization of the optimal fair L^2 -objective learning outcome.

3.3. Wasserstein Geodesics Characterization of Pareto Frontier

In reality, rather than looking for the optimal fair learning outcome, practitioners may have to choose a middle ground: sacrificing some prediction accuracy while tolerating a certain level of disparity. Therefore, it is tempting to generalize the barycenter characterization of the optimal fair learning outcome to the entire Pareto frontier between prediction error and statistical disparity. In this section, we show that the constant-speed geodesics from the conditional expectation sensitive marginals to their Wasserstein barycenter characterize the Pareto frontier on the Wasserstein space, in which utility loss and statistical disparity are quantified respectively by the L^2 norm and the average pair-wise Wasserstein distance among the sensitive marginals. As a result, given the optimal transport maps, one can derive a closed-form solution to the geodesics and thereby the Pareto frontier using McCann interpolation.

Here, we first provide a post-processing characterization of the Pareto frontier, Theorem 3.3.1, which is of theoretical interest and great generality. Then, we derive a closed-form solution to Problem 2 based on this characterization. The results form a direct generalization of the barycenter characterization, which is Lemma 5.3.1, and practitioners can apply the result together with the pseudo-barycenter and McCann interpolation to obtain the optimal affine estimation to the post-processing Pareto frontier. Later in Section 4.2, we further apply the result to provide a characterization of the exact solution and an optimal affine estimation of the solution to the optimal fair data representation problem (1.3).

Now, we start to characterize the Pareto frontier. In the rest of the section, we denote $\mathcal{L}(\mathbb{E}(Y|X, Z)) =: \mu, \mathcal{L}(\mathbb{E}(Y|X, Z)_z) =: \mu_z$. For utility, given any measurable function $f : \mathcal{X} \times \mathcal{Z} \rightarrow \mathcal{Y}$, we define the increased prediction error by the L^2 -norm of the difference between $f(X, Z)$ and the orthogonal projection $\mathbb{E}(Y|X, Z)$:

$$(3.15) \quad L(f(X, Z)) := \|\mathbb{E}(Y|X, Z) - f(X, Z)\|_2 = \left(\int_{\mathcal{Z}} \|\mathbb{E}(Y|X, Z)_z - f(X, Z)_z\|_2^2 d\lambda(z) \right)^{\frac{1}{2}}.$$

To simplify notation, we also denote

$$(3.16) \quad L(T') := L(T'(\mathbb{E}(Y|X, Z), Z)) = \left(\int_{\mathcal{Z}} \|\mathbb{E}(Y|X, Z)_z - T'_z(\mathbb{E}(Y|X, Z)_z)\|_2^2 d\lambda(z) \right)^{\frac{1}{2}}.$$

for any measurable $T' : \mathcal{Y} \times \mathcal{Z} \rightarrow \mathcal{Y}$.

To relax the hard independence constraint for the Pareto frontier, we quantify the statistical disparity of a given learning outcome or prediction \hat{Y} by the average pairwise Wasserstein distance among its sensitive marginals:

DEFINITION 3.3.1 (Wasserstein disparity).

$$(3.17) \quad D(\hat{Y}, Z) := \left(\int_{\mathcal{Z}^2} \mathcal{W}_2^2(\mathcal{L}(\hat{Y}_{z_1}), \mathcal{L}(\hat{Y}_{z_2})) d\lambda(z_1) d\lambda(z_2) \right)^{\frac{1}{2}}.$$

In our setting, $\hat{Y} = f(X, Z)$ for some $f : \mathcal{X} \times \mathcal{Z} \rightarrow \mathcal{Y}$. Also, to simplify notation, we denote the Wasserstein disparity that remains in the already deformed (by applying T') conditional expectation by

$$(3.18) \quad D(T') := D(T'(\mathbb{E}(Y|X, Z), Z), Z) = \left(\int_{\mathcal{Z}^2} \mathcal{W}_2^2((T'_{z_1})_{\#}\mu_{z_1}, (T'_{z_2})_{\#}\mu_{z_2}) d\lambda(z_1) d\lambda(z_2) \right)^{\frac{1}{2}}$$

for any measurable $T' : \mathcal{Y} \times \mathcal{Z} \rightarrow \mathcal{Y}$. Here, $T_z = T(\cdot, z) : \mathcal{Y} \rightarrow \mathcal{Y}$ for λ -a.e. $z \in \mathcal{Z}$.

We adopt the Wasserstein disparity as a statistical disparity quantification due to the following desirable properties:

- (*Wasserstein disparity characterizes statistical parity, Lemma 3.3.1*)

$$\mathcal{D}(f(X, Z), Z) = 0 \iff f(X, Z) \perp Z$$

- (*Physics interpretation*) Due to the definition based on the Wasserstein distance, Wasserstein disparity can be understood as the expected minimum amount of work that is required to move one randomly chosen marginal to another random chosen one. Therefore, the larger $\mathcal{D}(f(X, Z), Z)$ is, the more necessary work is expected to remove the distributional discrepancy among the sensitive groups on $f(X, Z)$.

LEMMA 3.3.1 (Characterization of statistical parity).

$$D(\hat{Y}, Z) = 0 \iff \hat{Y} \perp Z.$$

PROOF. We first show that (1) $\mathcal{W}_2^2(\mathcal{L}(\hat{Y}_z), \bar{\mu}) = 0$, λ -a.e. $\iff \hat{Y} \perp Z$, then complete the proof by showing (2) $D(\hat{Y}, Z) = 0 \iff \mathcal{W}_2^2(\mathcal{L}(\hat{Y}_z), \bar{\mu}) = 0$, λ -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 $\mathcal{L}(\hat{Y}_z) = \mathcal{L}(\hat{Y})$ λ -a.e. which implies $\mathcal{W}_2^2(\mathcal{L}(\hat{Y}_z), \mathcal{L}(\hat{Y})) = 0$, λ -a.e.. Then, it follows from the uniqueness of $\bar{\mu}$ that $\bar{\mu} = \mathcal{L}(\hat{Y})$ and $\mathcal{W}_2^2(\mathcal{L}(\hat{Y}_z), \bar{\mu}) = 0$, λ -a.e.. For the other direction, assume that $\mathcal{W}_2^2(\mathcal{L}(\hat{Y}_z), \bar{\mu}) = 0$ λ -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 $\mathcal{W}_2^2(\mathcal{L}(\hat{Y}_z), \bar{\mu}) \neq 0$, λ -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 $\mathcal{W}_2^2(\mathcal{L}(\hat{Y}_z), \bar{\mu}) = 0$, λ -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. \square

Now, let $T : \mathcal{Y} \times \mathcal{Z} \rightarrow \mathcal{Y}$ satisfy $T(\cdot, z)$ being the optimal transport maps from $\{\mu_z\}_z$ to their barycenter $\bar{\mu}$ for λ -a.e. $z \in \mathcal{Z}$ (See construction of T in the proof of Lemma 5.3.1), we define

$$(3.19) \quad V := L(T) = \left(\int_{\mathcal{Z}} \|\mathbb{E}(Y|X, Z)_z - T(\mathbb{E}(Y|X, Z)_z, z)\|_2^2 d\lambda(z) \right)^{\frac{1}{2}}$$

$$(3.20) \quad = \left(\int_{\mathcal{Z}} \|\mathbb{E}(Y|X, Z)_z - \overline{\mathbb{E}(Y|X, Z)}_z\|_2^2 d\lambda(z) \right)^{\frac{1}{2}}.$$

As shown in Lemma 5.3.1, V is the minimum increase of L^2 error (or, in physics, the minimum work/energy required) to deform $\mathbb{E}(Y|X, Z)$ to satisfy statistical parity. Before showing the main result, we need to define the geodesic on metric space to show the explicit form of constant speed geodesic on the Wasserstein space, which plays a key role in the proof.

DEFINITION 3.3.2 (Constant-speed geodesic between two points on metric space). *Given a metric space (X, d) and $x, x' \in X$, the constant-speed geodesic between x and x' is a continuously parametrized path $\{x_t\}_{t \in [0, 1]}$ such that $x_0 = x$, $x_1 = x'$, and $d(x_s, x_t) = |t - s|d(x, x')$, $\forall s, t \in [0, 1]$.*

The following lemma, which is well known as the McCann (displacement) interpolation [51, Chapter 7] in the optimal transport literature, shows that a linear interpolation using the optimal transport plan results in the constant-speed geodesic on the Wasserstein space.

LEMMA 3.3.2 (Constant-speed geodesic on Wasserstein space, [41, 51]). *Given $\mu_0, \mu_1 \in (\mathcal{P}_2(\mathbb{R}^d), \mathcal{W}_2)$ and γ the optimal transport plan in between, let $\pi_t(x, y) := (1 - t)x + ty$, then*

$$(3.21) \quad \mu_t := (\pi_t)_\# \gamma, t \in [0, 1]$$

is the constant-speed geodesic between μ_0 and μ_1 .

PROOF. First, it follows from the triangle inequality that

$$\mathcal{W}_2(\mu_0, \mu_1) \leq \mathcal{W}_2(\mu_0, \mu_s) + \mathcal{W}_2(\mu_s, \mu_t) + \mathcal{W}_2(\mu_t, \mu_1)$$

for any $s, t \in [0, 1]$. On the other hand, it follows from the definition of μ_t that for $s, t \in [0, 1]$

$$\begin{aligned}
\mathcal{W}_2^2(\mu_s, \mu_t) &\leq \int_{(\mathbb{R}^d)^2} \|x - y\|^2 d(\pi_s)_\# \gamma(x) \otimes d(\pi_t)_\# \gamma(y) \\
&= \int_{(\mathbb{R}^d)^2} \|\pi_s(x, y) - \pi_t(x, y)\|^2 d\gamma(x, y) \\
&= \int_{(\mathbb{R}^d)^2} \|(1-s)x + sy - (1-t)x - ty\|^2 d\gamma(x, y) \\
&= \int_{(\mathbb{R}^d)^2} \|(t-s)x - (t-s)y\|^2 d\gamma(x, y) \\
&= |t-s|^2 \int_{(\mathbb{R}^d)^2} \|x - y\|^2 d\gamma(x, y) = |t-s|^2 \mathcal{W}_2^2(\mu_0, \mu_1),
\end{aligned}$$

where the first equation results from definition of \mathcal{W}_2 . Given the above two facts, we complete the proof by contradiction. Assume $\exists s, t \in [0, 1]$ such that $\mathcal{W}_2(\mu_s, \mu_t) < |t-s|\mathcal{W}_2(\mu_0, \mu_1)$, then

$$\begin{aligned}
\mathcal{W}_2(\mu_0, \mu_1) &\leq \mathcal{W}_2(\mu_0, \mu_s) + \mathcal{W}_2(\mu_s, \mu_t) + \mathcal{W}_2(\mu_t, \mu_1) \\
&< |s|\mathcal{W}_2(\mu_0, \mu_1) + |t-s|\mathcal{W}_2(\mu_0, \mu_1) + |1-t|\mathcal{W}_2(\mu_t, \mu_1) \\
&= \mathcal{W}_2(\mu_0, \mu_1).
\end{aligned}$$

□

REMARK 3.3.1 (Linear interpolation formula for \mathcal{W}_2 geodesics). *If there exists an optimal transport map T such that $T_\#(\mu_0) = \mu_1$, then the McCann interpolation has the simple form*

$$(3.22) \quad \mu_t = ((1-t)Id + tT)_\# \mu_0, t \in [0, 1].$$

We apply this simple formula to obtain a closed-form estimation of the Pareto frontier in algorithm design, see Section 4.6.

Now, we are ready to establish the main result, which shows that V is a lower bound of $L(f(X, Z)) + \frac{1}{\sqrt{2}}D(f(X, Z), Z)$ for any measurable function $f : \mathcal{X} \times \mathcal{Z} \rightarrow \mathcal{Y}$ and is achieved along the constant-speed geodesics from the sensitive marginals of the conditional expectation to their barycenter on the Wasserstein space.

THEOREM 3.3.1 (\mathcal{W}_2 geodesics characterization of a linear Pareto frontier). *Define L, D, V as above and assume $\mu_z \in \mathcal{P}_{2,ac}(\mathcal{Y}), \lambda - a.e.$. It follows that*

$$(3.23) \quad V \leq L(f(X, Z)) + \frac{1}{\sqrt{2}}D(f(X, Z), Z)$$

for any measurable function $f : \mathcal{X} \times \mathcal{Z} \rightarrow \mathcal{Y}$. Furthermore, define $T(t)$ such that $T(t)(\cdot, z) := (1-t)Id + t(T(\cdot, z)), t \in [0, 1]$ is the linear interpolation between the identity map and the optimal transport map for λ -a.e. $z \in \mathcal{Z}$, then equality holds in (3.23) if and only if $f(X, Z) = T(t)(\mathbb{E}(Y|X, Z), Z), t \in [0, 1]$ as

$$(3.24) \quad L(T(t)) = tL(T(1)) = tV$$

$$(3.25) \quad \frac{1}{\sqrt{2}}D(T(t)) = \frac{1}{\sqrt{2}}(1-t)D(T(0)) = (1-t)V.$$

PROOF. First, we derive the inequality from the triangle inequality and the optimality of $\{T(\cdot, z)\}_z$: Let $f : \mathcal{X} \times \mathcal{Z} \rightarrow \mathcal{Y}$ be an arbitrary measurable function. It follows that

$$\begin{aligned} V &\leq \left(\int_{\mathcal{Z}} \|\mathbb{E}(Y|X, Z)_z - \overline{f(X, Z)}_z\|_2^2 d\lambda(z) \right)^{\frac{1}{2}} \\ &\leq L(f(X, Z)) + \left(\int_{\mathcal{Z}} \|f(X, Z)_z - \overline{f(X, Z)}_z\|_2^2 d\lambda(z) \right)^{\frac{1}{2}} \\ &\leq L(f(X, Z)) + \left(\int_{\mathcal{Z}} \mathcal{W}_2^2(\mathcal{L}(f(X, Z)_z), \overline{\mathcal{L}(f(X, Z)_z)}) d\lambda(z) \right)^{\frac{1}{2}} \\ &= L(f(X, Z)) + \left(\frac{1}{2} \int_{\mathcal{Z}^2} \mathcal{W}_2^2(\mathcal{L}(f(X, Z)_{z_1}), \mathcal{L}(f(X, Z)_{z_2})) d\lambda(z_1) d\lambda(z_2) \right)^{\frac{1}{2}} \\ &= L(f(X, Z)) + \frac{1}{\sqrt{2}}D(f(X, Z)). \end{aligned}$$

Here, the penultimate equation results from the fact that, for any $\{\nu_z\}_z \subset \mathcal{P}_{2,ac}(\mathbb{R}^d)$,

$$(3.26) \quad \int_{\mathcal{Z}^2} \mathcal{W}_2^2(\nu_{z_1}, \nu_{z_2}) d\lambda(z_1) d\lambda(z_2) = 2 \int_{\mathcal{Z}} \mathcal{W}_2^2(\nu_z, \bar{\nu}) d\lambda(z),$$

where $\bar{\nu}$ is the Wasserstein barycenter of $\{\nu_z\}_z$. Now, we show that the lower bound is achieved if and only if $f(X, Z) = T(t)(\mathbb{E}(Y|X, Z), Z), t \in [0, 1]$. Let $t \in [0, 1]$, $T_z := T(\cdot, z)$, and $\mu_z := \mathcal{L}(\mathbb{E}(Y|X, Z)_z)$. It follows from Lemma 3.3.2 and Remark 3.3.1 that:

$$\begin{aligned}
V &= \left(\int_{\mathcal{Z}} \mathcal{W}_2^2(\mu_z, \bar{\mu}) d\lambda(z) \right)^{\frac{1}{2}} \\
&\leq \left(\int_{\mathcal{Z}} \mathcal{W}_2^2(\mu_z, T_z(t) \# \mu_z) d\lambda(z) \right)^{\frac{1}{2}} + \left(\int_{\mathcal{Z}} \mathcal{W}_2^2(T_z(t) \# \mu_z, \bar{\mu}) d\lambda(z) \right)^{\frac{1}{2}} \\
&= (t^2 \int_{\mathcal{Z}} \mathcal{W}_2^2(\mu_z, \bar{\mu}) d\lambda(z))^{\frac{1}{2}} + ((1-t)^2 \int_{\mathcal{Z}} \mathcal{W}_2^2(\mu_z, \bar{\mu}) d\lambda(z))^{\frac{1}{2}} \\
&= tV + (1-t)V = V.
\end{aligned}$$

Therefore, the second inequality is an equality where the first term is $L(T(t))$:

$$\begin{aligned}
L(T(t)) &= \left(\int_{\mathcal{Z}} \|\mathbb{E}(Y|X, Z)_z - T_z(t)(\mathbb{E}(Y|X, Z)_z)\|_2^2 d\lambda(z) \right)^{\frac{1}{2}} \\
&= \left(\int_{\mathcal{Z}} \mathcal{W}_2^2(\mu_z, T_z(t) \# \mu_z) d\lambda(z) \right)^{\frac{1}{2}} \\
&= t \left(\int_{\mathcal{Z}} \mathcal{W}_2^2(\mu_z, \bar{\mu}) d\lambda(z) \right)^{\frac{1}{2}} = tV.
\end{aligned}$$

For the second term, we claim that it equals $\frac{1}{\sqrt{2}}D(T(t))$. To see this, we need to first show $\overline{T_z(t) \# \mu_z} = \bar{\mu}$. Indeed, if not, then $\int_{\mathcal{Z}} \mathcal{W}_2^2(T_z(t) \# \mu_z, \overline{T_z(t) \# \mu_z}) d\lambda(z)$ is strictly less than $\int_{\mathcal{Z}} \mathcal{W}_2^2(T_z(t) \# \mu_z, \bar{\mu}) d\lambda(z)$ by the definition and uniqueness of $\overline{T_z(t) \# \mu_z}$. It follows that

$$\begin{aligned}
&\left(\int_{\mathcal{Z}} \mathcal{W}_2^2(\mu_z, \overline{T_z(t) \# \mu_z}) d\lambda(z) \right)^{\frac{1}{2}} \\
&\leq \left(\int_{\mathcal{Z}} \mathcal{W}_2^2(\mu_z, T_z(t) \# \mu_z) d\lambda(z) \right)^{\frac{1}{2}} + \left(\int_{\mathcal{Z}} \mathcal{W}_2^2(T_z(t) \# \mu_z, \overline{T_z(t) \# \mu_z}) d\lambda(z) \right)^{\frac{1}{2}} \\
&< L(T(t)) + \left(\int_{\mathcal{Z}} \mathcal{W}_2^2(T_z(t) \# \mu_z, \bar{\mu}) d\lambda(z) \right)^{\frac{1}{2}} \\
&= \left(\int_{\mathcal{Z}} \mathcal{W}_2^2(\mu_z, \bar{\mu}) d\lambda(z) \right)^{\frac{1}{2}},
\end{aligned}$$

which contradicts the definition and uniqueness of $\bar{\mu}$. Therefore,

$$\begin{aligned}
D(T(t)) &= \left(\int_{\mathcal{Z}^2} \mathcal{W}_2^2(T_{z_1}(t) \# \mu_{z_1}, T_{z_2}(t) \# \mu_{z_2}) d\lambda(z_1) d\lambda(z_2) \right)^{\frac{1}{2}} \\
&= \left(2 \int_{\mathcal{Z}} \mathcal{W}_2^2(T_z(t) \# \mu_z, \overline{T_z(t) \# \mu_z}) d\lambda(z) \right)^{\frac{1}{2}} \\
&= \sqrt{2} \left(\int_{\mathcal{Z}} \mathcal{W}_2^2(T_z(t) \# \mu_z, \bar{\mu}) d\lambda(z) \right)^{\frac{1}{2}} \\
&= \sqrt{2} ((1-t)^2 \int_{\mathcal{Z}} \mathcal{W}_2^2(\mu_z, \bar{\mu}) d\lambda(z))^{\frac{1}{2}} \\
&= \sqrt{2} (1-t)V.
\end{aligned}$$

That completes the proof. □

REMARK 3.3.2 (Intuition of Theorem 3.3.1: a Euclidean analog). *Here, we provide a Euclidean analog of Theorem 3.3.1. In fact, our proof is based on the observation of the analog and equivalent to it when one considers $x \rightarrow \delta_x$ as an embedding from \mathcal{X} to $\mathcal{P}_2(\mathcal{X})$.*

Let $X := \{x_i\}_{i=1}^N$ be a fixed data set on the Euclidean space \mathcal{X} ($N = 3$ in Figure 3.1), $\tilde{X} := \{\tilde{x}_i\}_{i=1}^N$ be a data set consisting of N arbitrarily chosen data points on \mathcal{X} , and define the following:

- [Euclidean analog of V] $std(X) := \left(\frac{1}{N} \sum_{i=1}^N \|x_i - m_x\|^2 \right)^{\frac{1}{2}}$ with $m_x := \frac{1}{N} \sum_{i=1}^N x_i$,
- [Euclidean analog of L] $solid(\tilde{X}) := \left(\frac{1}{N} \sum_{i=1}^N \|x_i - \tilde{x}_i\|^2 \right)^{\frac{1}{2}}$,
- [Euclidean analog of D] $dotted(\tilde{X}) := \left(\frac{1}{N^2} \sum_{i,j=1}^N \|\tilde{x}_i - \tilde{x}_j\|^2 \right)^{\frac{1}{2}}$.




FIGURE 3.1. In this figure, we have three data points on an Euclidean space traveling along straight lines (Euclidean geodesics) to their average (Euclidean barycenter). Define (1) $\text{std} :=$ the standard deviation of the three points, (2) solid line (loss) $:=$ the average moving (Euclidean) distance away from their original location, and (3) dotted line (disparity) $:=$ the average pairwise (Euclidean) distance among them. One can show that $\text{std} \leq \text{solid} + \frac{1}{\sqrt{2}} \text{dotted}$ where equality holds if and only if the three points travel at constant-speed along straight lines to their average.

It is straight-forward to verify that (1)

$$\text{std}(X) \leq \underbrace{\text{solid}(\tilde{X})}_{\text{utility loss}} + \frac{1}{\sqrt{2}} \underbrace{\text{dotted}(\tilde{X})}_{\text{disparity}},$$

and (2) equality holds if and only if $\tilde{X} = X(t) := \{(1-t)x_i + tm_x\}_{i=1}^N$ for $t \in [0, 1]$ as $\text{loss}(X(t)) = t\text{std}(X)$ and $\frac{1}{\sqrt{2}}\text{disparity}(X(t)) = (1-t)\text{std}(X)$.

Since V (the minimum work or energy required for statistical parity) is fixed for the data (X, Y, Z) when one applies (X, Z) to predict Y , the above theorem implies that the Pareto frontier between the increased prediction error $L(T)$ and the remaining disparity $D(T)$ is a line that results from the constant-speed geodesics from the marginal conditional expectations to their barycenter on the Wasserstein space. In particular, let $T(t)(\mathbb{E}(Y|X, Z), Z) := \{T(t)(\mathbb{E}(Y|X, Z)_z, z)\}_z$, λ -a.e., $t \in [0, 1]$, we arrive at a closed-form solution to Problem 2:

COROLLARY 3.3.1 (Pareto optimal fair L^2 -objective learning). *Given (X, Y, Z) satisfying $\mu_z \in \mathcal{P}_{ac}$, λ -a.e., then*

$$(3.27) \quad f_d(X, Z) := \begin{cases} T(1 - \frac{d}{\sqrt{2V}})(\mathbb{E}(Y|X, Z), Z), & \text{if } d \in [0, \sqrt{2V}] \\ \mathbb{E}(Y|X, Z), & \text{if } d \in (\sqrt{2V}, \infty) \end{cases}$$

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

PROOF. If $d \in (\sqrt{2V}, \infty)$, then it follows from Theorem 3.3.1 that $D(\mathbb{E}(Y|X, Z)) = D(T(0)) = \sqrt{2V} < d$. Hence, Problem 2 reduces to the unconstrained L^2 projection problem and the optimal solution is $\mathbb{E}(Y|X, Z)$. Now, for a fixed $d \in [0, \sqrt{2V}]$, assume for contradiction that $\exists f \in L^2(\mathcal{X} \times \mathcal{Z}, \mathcal{Y})$ such that

$$\|Y - f(X, Z)\|_2^2 < \|Y - T(t)(\mathbb{E}(Y|X, Z), Z)\|_2^2$$

for $t = 1 - \frac{d}{\sqrt{2V}}$. Then, let $\overline{f(X, Z)}$ denote the Wasserstein barycenter of $\{f(X, Z)_z\}_z$, we have

$$\begin{aligned} \|Y - \overline{f(X, Z)}\|_2^2 &\leq \|Y - f(X, Z)\|_2^2 + \|f(X, Z) - \overline{f(X, Z)}\|_2^2 \\ &< \|Y - T(t)(\mathbb{E}(Y|X, Z), Z)\|_2^2 + \|f(X, Z) - \overline{f(X, Z)}\|_2^2 \\ &= \|Y - \mathbb{E}(Y|X, Z)\|_2^2 + L(T(t)) + \frac{1}{\sqrt{2}}D(f(X, Z)) \\ &= \|Y - \mathbb{E}(Y|X, Z)\|_2^2 + (V - \frac{1}{\sqrt{2}}d) + \frac{1}{\sqrt{2}}d \\ &= \|Y - \mathbb{E}(Y|X, Z)\|_2^2 + V \end{aligned}$$

where the second line follows from the assumption, the third from L^2 orthogonal decomposition and Theorem 3.3.1, and the fourth from the assumption and Theorem 3.3.1. The strict inequality above contradicts the optimality of $\overline{\mathbb{E}(Y|X, Z)}$ shown in Lemma 5.3.1. That proves the optimality of $T(1 - \frac{d}{\sqrt{2V}})(\mathbb{E}(Y|X, Z), Z)$ for the fixed d . Uniqueness result follows from the uniqueness of $\overline{\mathbb{E}(Y|X, Z)}$ shown in Lemma 5.3.1. Since the choice of $d \in [0, \sqrt{2V}]$ is arbitrary, we are done. \square

We note that Corollary 3.3.1 together with Lemma 3.3.2 and Remark 3.3.1 provide a post-processing approach to (estimate) the Pareto frontier: applying McCann interpolation to the Brenier's maps between the learning outcome sensitive marginals $\{\mathbb{E}(Y|X, Z)_z\}_z$ and their (pseudo-) barycenter. One can apply Algorithm 1 directly with the learning outcome marginals as inputs.

From a theoretical perspective, various metrics of disparity that differ from D , the Wasserstein disparity (Definition 3.3.1), can be used and the theoretical results derived in this section provide a lower bound estimation for the Pareto frontier that uses other disparity metrics. The quality of the lower bound can be studied using the relationship between the Wasserstein distance and the defined disparity metric. Also, the present work provides a numerical study on the lower bound estimation in Section 6 to which we refer the interested readers for more details.

In practice, various metrics of disparity are adopted, such as the prediction success ratio (difference from 1) in classification [14] and the Kolmogorov-Smirnov distance for 1-dimensional regression [19]. The proposed estimation of the Pareto frontier leaves the choice of α to practitioners who would face specific fairness requirements and disparity metrics.

Fair Data Representation for Conditional Expectation Estimation

In this chapter, we study the optimal fair data representation problem, Problem 3, that is motivated by the current challenges in the pre-processing or synthetic data design approach to fair machine learning. To solve the problem, we first characterize the exact solution using a dependent and independent Wasserstein barycenter pair, see Lemma 4.2.4. Then, we define a dependent and independent pseudo-barycenter pair via optimal affine maps, and prove that the pair is the exact optimal fair data representation with Gaussian marginals, cf. Lemma 4.3.2 and the optimal affine estimate of the representation with general marginals in Theorem 4.4.1.

4.1. Objective & Constraint for Fair Data Representation

In this section, we derive a fairness objective function that is both theoretically tractable and practically appealing. This task is more involved than one initially might expect, and it sheds light on some subtleties of both the post-processing and the pre-processing approaches.

Before proceeding, we need some preparation. Let X , Y , and Z represent respectively the independent, dependent, and sensitive random variable, with the same underlying probability space $(\Omega, \Sigma, \mathbb{P})$. We use the term ‘random variables’ to denote random vectors with an arbitrary but finite dimension. That is, $S : \Omega \rightarrow \mathcal{S}$ where $\mathcal{S} \in \{[k_S]^{d_S}, \mathbb{N}^{d_S}, [0, l_S]^{d_S}, \mathbb{R}^{d_S}\}$ with $k_S \in \mathbb{N}$, $l_S \in \mathbb{R}$ and $d_S < \infty$ for $S \in \{X, Y, Z\}$.

It follows from [19, 29] that the optimal fair regression outcome can be characterized by the Wasserstein barycenter. In Lemma 5.3.1 we will generalize their result from regression to all functions in $L^2(\mathcal{X} \times \mathcal{Z}, \mathcal{Y})$, which shows that the optimal fair L^2 -objective supervised learning outcome can be characterized by solutions to Problem 1:

$$(4.1) \quad \inf_{f \in L^2(\mathcal{X} \times \mathcal{Z}, \mathcal{Y})} \{\|Y - f(X, Z)\|_2^2 : f(X, Z) \perp Z\}$$

The utility loss is quantified by L^2 -norm: $\|Y - f(X, Z)\|_2^2 = \int_{\Omega} \|Y - f(X, Z)\|^2 d\mathbb{P}$, where $\|\cdot\|$ is the Euclidean norm. The constraint $f(X, Z) \perp Z$ guarantees that the final result satisfies statistical parity and, therefore, is fair.

Since it follows from L^2 orthogonal decomposition that

$$(4.2) \quad \|Y - f(X, Z)\|_2^2 = \|Y - \mathbb{E}(Y|X, Z)\|_2^2 + \|\mathbb{E}(Y|X, Z) - f(X, Z)\|_2^2$$

and only the second term on the right hand side depends on the choice of $f \in L^2(\mathcal{X} \times \mathcal{Z}, \mathcal{Y})$, we conclude that (5.5) is equivalent to

$$(4.3) \quad \inf_{f \in L^2(\mathcal{X} \times \mathcal{Z}, \mathcal{Y})} \{\|\mathbb{E}(Y|X, Z) - f(X, Z)\|_2^2 : f(X, Z) \perp Z\}.$$

It turns out—see Lemma 5.3.1—that the solution to (4.3) is exactly the Wasserstein barycenter. Therefore, we say that the optimal fair L^2 -objective supervised learning outcome is characterized by the Wasserstein barycenter. But notice that the Wasserstein barycenter characterization (4.3) assumes knowledge of the learning outcome $\mathbb{E}(Y|X, Z)$. That is, if practitioners apply the characterization to estimate the optimal learning outcome, it is necessary to obtain an estimator of $\mathbb{E}(Y|X, Z)$ via supervised learning before solving the post-processing rescue step (4.3). Therefore, we say that the characterization has a post-processing nature and hence call it a post-processing characterization.

Now, notice that the estimator of $\mathbb{E}(Y|X, Z)$ is obtained via the training process

$$(4.4) \quad \inf_{f \in \mathcal{F}} \{\|Y - f(X, Z)\|_2^2\},$$

where the admissible function set \mathcal{F} depends on the choice of supervised learning models. Denote the estimator by $f'(X, Z)$. Then in practice (4.3) becomes

$$(4.5) \quad \inf_{f \in L^2(\mathcal{X} \times \mathcal{Z}, \mathcal{Y})} \{\|f'(X, Z) - f(X, Z)\|_2^2 : f(X, Z) \perp Z\}.$$

That is, the application of the post-processing characterization is model-dependent. The fundamental reason for model dependence is that (5.5) is optimizing over all L^2 functions while in practice it is necessary to reduce the admissible set from L^2 to some \mathcal{F} which depends on the choice of the model. As a result, the optimizer is necessarily dependent on the choice of the model. Therefore, the constrained optimization (5.5) and its characterization are not suitable for our ultimate goal of

deriving a model-independent pre-processing approach to the optimal fair learning outcome. The present work proposes a different constrained optimization problem that characterizes the optimal fair data representation for all L^2 -objective supervised learning models.

To make a constraint optimization problem suitable for fair data representation design, we require both the objective function and the fairness constraint to be model-independent. Furthermore, the data representation design objective and the training objective given the data representation have to be consistent in the following sense: the better training and testing result on the fair data representation leads to less L^2 -fitting error with respect to the true data.

We now derive an objective function that is suitable for fair data representation design purpose. To start, notice that our goal is to generate a synthetic data representation (\tilde{X}, \tilde{Y}) , a deformation of (X, Y) , via which any L^2 -objective model that is trained by

$$(4.6) \quad \inf_{f \in \mathcal{F}} \|\tilde{Y} - f(\tilde{X})\|_2^2$$

would result in (an estimation of) the optimal fair learning outcome. In the rest of this dissertation, we denote the solution to (4.6) by $f_{\tilde{Y}}$.

Also, because conditional expectation is an orthogonal projection operator on L^2 -space, we obtain the following orthogonal decomposition of the objective in (4.6):

$$(4.7) \quad \|\tilde{Y} - f(\tilde{X})\|_2^2 = \|\tilde{Y} - \mathbb{E}(\tilde{Y}|\tilde{X})\|_2^2 + \|\mathbb{E}(\tilde{Y}|\tilde{X}) - f(\tilde{X})\|_2^2.$$

Only the second term on the right hand side depends on the choice of $f \in \mathcal{F}$, hence the training step objective (4.6) is equivalent to the following:

$$(4.8) \quad \inf_{f \in \mathcal{F}} \|\mathbb{E}(\tilde{Y}|\tilde{X}) - f(\tilde{X})\|_2^2.$$

Thus, the solution to (4.8) is also $f_{\tilde{Y}}$, which depends on the choice of \mathcal{F} .

The key observation is that, given a data representation (\tilde{X}, \tilde{Y}) , (4.8) is the objective that practitioners try to achieve via model selection, modification, and parameter turning. Furthermore, it follows from the triangle or Minkowski inequality that

$$(4.9) \quad \underbrace{\|Y - f_{\tilde{Y}}(\tilde{X})\|_2}_{\text{total utility loss}} \leq \underbrace{\|Y - \mathbb{E}(\tilde{Y}|\tilde{X})\|_2}_{\text{data representation utility loss}} + \underbrace{\|\mathbb{E}(\tilde{Y}|\tilde{X}) - f_{\tilde{Y}}(\tilde{X})\|_2}_{\text{learning utility loss}}.$$

The second term on the right-hand side is the target of a supervised learning task which should be left to practitioners. Thus, the natural choice of the model-independent objective of the optimal fair synthetic data design is to minimize the first term:

$$(4.10) \quad \inf_{(\tilde{X}, \tilde{Y}) \in \mathcal{D}} \|Y - \mathbb{E}(\tilde{Y}|\tilde{X})\|_2,$$

where \mathcal{D} is some admissible set of deformed versions of the original data (X, Y) that we define later. Intuitively, the loss function can be interpreted as the potential utility sacrifice resulting from deforming (X, Y) to (\tilde{X}, \tilde{Y}) for L^2 -objective supervised learning, while leaving the task of minimizing the second term on the right-hand side to practitioners via model selection, modification, or parameter tuning.

Next, we derive a fairness constraint for synthetic data design purposes. That is, the goal is to design (\tilde{X}, \tilde{Y}) such that $f_{\tilde{Y}}(\tilde{X}) \perp Z$ for any admissible function set $\mathcal{F} \subset L^2(\mathcal{X}, \mathcal{Y})$. The flexibility of model choice becomes important due to the increasing complexity of models in practice nowadays, such as neural networks. The key observation here is that, due to the potential dependence of $f_{\tilde{Y}}$ on Z , one needs to look at both models that use merely measurable functions from \mathcal{X} to \mathcal{Y} and more complicated models consisting of Z -dependent measurable functions:

- 1 For measurable functions from \mathcal{X} to \mathcal{Y} , if we require $\tilde{X} \perp Z$, then it follows that for any $f : \mathcal{X} \rightarrow \mathcal{Y}$, it is guaranteed that $f(\tilde{X}) \perp Z$. Hence, we require $\tilde{X} \perp Z$ to prevent models from exploiting sensitive information from the independent variables.
- 2 For advanced or adversarial models that use Z -dependent functions from $\mathcal{X} \times \mathcal{Z}$ to \mathcal{Y} , the trained model f_Y could still depend on Z because Y and Z are not independent. For example, consider the extreme case where $Y = kZ, k \in \mathbb{R}$ and a perfect model results in $\mathbb{E}(kZ|\tilde{X}, Z) = kZ$ which fully depends on Z even if we require $\tilde{X} \perp Z$. Therefore, we also require $f_{\tilde{Y}}(\tilde{X}, Z) \perp Z$ to prevent such a model from exploiting sensitive information from the dependent variables.

But notice that the second requirement leads us back to the post-processing nature of fairness constraints as in (4.5). For fair data representation design purposes, it is necessary to keep the constraint model-independent. Therefore, instead of enforcing $f_{\tilde{Y}}(\tilde{X}, Z) \perp Z$, the present work requires $\mathbb{E}(\tilde{Y}|\tilde{X}, Z) \perp Z$ for the following two reasons: (1) Under the modified constraint $\mathbb{E}(\tilde{Y}|\tilde{X}, Z) \perp Z$, the better $f_{\tilde{Y}}(\tilde{X}, Z)$ estimates $\mathbb{E}(\tilde{Y}|\tilde{X}, Z)$, the more independent of Z becomes $f_{\tilde{Y}}(\tilde{X}, Z)$. Such alignment between training objective and fairness makes the modification a natural choice under

the assumption that the goal of L^2 -objective (adversarial) supervised learning tasks is to minimize $\|\mathbb{E}(\tilde{Y}|\tilde{X}, Z) - f_{\tilde{Y}}(\tilde{X}, Z)\|_2^2$, which is equivalent to minimizing $\|\tilde{Y} - f_{\tilde{Y}}(\tilde{X}, Z)\|_2^2$. (2) Since a supervised learning model with poor prediction accuracy already results in severe unfairness, the dependence on sensitive information is of less concern when designing a fair data representation. Based on the fairness requirement for both measurable functions on merely \mathcal{X} and Z -dependent functions, a natural choice of (pre-processing) statistical parity constraint for data representation has the following form:

$$(4.11) \quad \tilde{X}, \mathbb{E}(\tilde{Y}|\tilde{X}, Z) \perp Z.$$

It guarantees: (1) statistical parity for any model that uses only a deterministic function and any model that results in a perfect estimation of $\mathbb{E}(\tilde{Y}|\tilde{X})$; (2) the better $f_{\tilde{Y}}(\tilde{X}, Z)$ estimates $\mathbb{E}(\tilde{Y}|\tilde{X}, Z)$, the more independent $f_{\tilde{Y}}(\tilde{X}, Z)$ becomes of Z .

While the fairness constraint (4.11) is not the only choice, it does balance utility and fairness. The following remark discusses two alternative fairness constraint choices, which are more polarized in optimizing utility or fairness.

REMARK 4.1.1 (Alternative fair data representation constraints). *There are two alternative choices of fairness constraints that are valuable in practice:*

- 1 $\tilde{X} \perp Z$: *the weaker constraint guarantees any model using merely a deterministic function, even if sub-optimal, to result in statistical parity. But it does not protect Z from advanced models, which exploit the dependence of Y on Z and apply Z -dependent functions. Therefore, $\tilde{X} \perp Z$ provides more utility but less sensitive information protection, compared to our choice.*
- 2 $(\tilde{X}, \tilde{Y}) \perp Z$: *the stronger constraint guarantees statistical parity in the learning outcome of any supervised learning model, even for those that adopt Z -dependent functions and are suboptimal. But it sacrifices more utility. This stronger constraint is particularly useful in practice when one does not know which variables are dependent and which ones are independent.*

Our choice is a compromise of the two alternatives in terms of balancing utility sacrifice and protecting sensitive information. Furthermore, simple modifications of our analysis and algorithm would solve the two alternatives because they are essentially simplified versions of our choice. Hence, the present work targets (4.11).

Finally, combining the objective and constraint for synthetic data design, we aim to solve Problem 3:

$$(4.12) \quad \inf_{(\tilde{X}, \tilde{Y}) \in \mathcal{D}} \{ \|Y - \mathbb{E}(\tilde{Y}|\tilde{X})\|_2^2 : \tilde{X}, \mathbb{E}(\tilde{Y}|\tilde{X}, Z) \perp Z \}.$$

The solution provides a fair data representation via which the trained L^2 -objective supervised learning models become estimations of the optimal fair conditional expectation.

Compared to the original constrained optimization problem (5.5) which results in the post-processing nature of its barycenter characterization (4.3), the proposed constrained optimization problem (1.3) has the following advantages by design:

- It provides a fairness guarantee for arbitrary L^2 -objective models.
- The model-independence together with the alignment between training objective and fairness enables practitioners to enjoy flexibility in model selection, modification, and parameter tuning on the fair data representation.
- The fair data representation approach has more applicable models than the post-processing approach. See Remark 4.1.2 below for a detailed explanation of two different interpretations of L^2 -objective models.

In the following remark, we explain the different interpretations of L^2 -objective models in the post-processing and pre-processing approaches.

REMARK 4.1.2 (Interpretation of L^2 -objective models). *For the post-processing approach, it follows from (4.3) and (4.5) that the barycenter characterization works only if the supervised learning model comes with an objective function in explicit L^2 -form. For the proposed pre-processing approach, the applicable L^2 -objective models include all the models that aim to estimate the conditional expectation. In particular, it follows from (4.9) and (4.10) that the proposed fair data representation works for any supervised learning model that aims to estimate conditional expectation or conditional probability, even though some of them do not come with an explicit objective function in L^2 -form. For example, all classification models share the goal of estimating the conditional probability of $\{Y = 1\}$ given an observation of $\{X = x\}$, which is $\mathbb{E}(1_{Y=1}|X = x)$. Therefore, the resulting synthetic data can be used for any classification model, even models such as logistic regression and random forest that do not have L^2 -based objective functions.*

4.2. Wasserstein Barycenter Pair Characterization

We will prove a characterization of the solutions to Problem 3. To start, notice that since $(\tilde{X}, Z) = T \otimes Id|_{\mathcal{Z}}(X, Z)$ for some measurable map $T \otimes Id|_{\mathcal{Z}} : \mathcal{X} \times \mathcal{Z} \rightarrow \mathcal{X} \times \mathcal{Z}$, we have $\sigma((\tilde{X}, Z)) \subset \sigma((X, Z))$. Also, from $\tilde{X} \perp Z$, we have $\sigma(\tilde{X}) \subset \sigma(\tilde{X}) \otimes \sigma(Z) = \sigma((\tilde{X}, Z))$. Therefore, $\sigma(\tilde{X}) \subset \sigma((X, Z))$ and it follows from L^2 orthogonal decomposition that

$$(4.13) \quad \|Y - \mathbb{E}(\tilde{Y}|\tilde{X})\|_2^2 = \|Y - \mathbb{E}(Y|X, Z)\|_2^2 + \|\mathbb{E}(Y|X, Z) - \mathbb{E}(\tilde{Y}|\tilde{X})\|_2^2.$$

The first term on the right hand side can be interpreted as the minimum loss of information by using (X, Z) to predict Y . Furthermore, one can decompose the second term on the right hand side of (4.13):

$$\begin{aligned} & \|\mathbb{E}(Y|X, Z) - \mathbb{E}(\tilde{Y}|\tilde{X})\|_2^2 \\ &= \|\mathbb{E}(Y|X, Z) - \mathbb{E}(Y|\tilde{X}, Z)\|_2^2 + \|\mathbb{E}(Y|\tilde{X}, Z) - \mathbb{E}(\tilde{Y}|\tilde{X})\|_2^2 \\ &= \|\mathbb{E}(Y|X, Z) - \mathbb{E}(Y|\tilde{X}, Z)\|_2^2 + \int_{\mathcal{Z}} \|\mathbb{E}(Y_z|\tilde{X}) - \mathbb{E}(\tilde{Y}|\tilde{X})_z\|_2^2 d\lambda(z). \end{aligned}$$

Here, the first equality follows from L^2 orthogonal decomposition. The second equality follows from disintegration, the fairness constraint $\tilde{X}, \mathbb{E}(\tilde{Y}|\tilde{X}) \perp Z$, and Lemma 4.2.1 below, which shows that $\tilde{X} \perp Z$ implies $\mathbb{E}(Y_z|\tilde{X}) = \mathbb{E}(Y|\tilde{X}, Z)_z$.

LEMMA 4.2.1. $\tilde{X} \perp Z \implies \mathbb{E}(Y_z|\tilde{X}) = \mathbb{E}(Y|\tilde{X}, Z)_z$

PROOF. Let $\tilde{X} \perp Z$ and assume for contradiction that $\mathbb{E}(Y_z|\tilde{X}) \neq \mathbb{E}(Y|\tilde{X}, Z)_z$. Then, we have

$$\begin{aligned} \|Y - \mathbb{E}(Y|\tilde{X}, Z)\|_2^2 &= \int_{\mathcal{Z}} \|Y_z - f^*(\tilde{X}, Z)_z\|_2^2 d\lambda \\ &= \int_{\mathcal{Z}} \|Y_z - f^*(\tilde{X}, z)\|_2^2 d\lambda \\ &> \int_{\mathcal{Z}} \|Y_z - \mathbb{E}(Y_z|\tilde{X})\|_2^2 d\lambda \\ &= \int_{\mathcal{Z}} \|Y_z - \tilde{f}_z(\tilde{X})\|_2^2 d\lambda \end{aligned}$$

where the first line follows from disintegration and the fact that there exists a measurable function $f^* : \mathcal{X} \times \mathcal{Z} \rightarrow \mathcal{Y}$ such that $f^*(\tilde{X}, Z) = \mathbb{E}(Y|\tilde{X}, Z)$, the second from $\tilde{X} \perp Z$, the third line follows from orthogonal projection property of conditional expectation and the assumption, and the fourth from the fact that there exists a measurable function $\tilde{f}_z : \mathcal{X} \rightarrow \mathcal{Y}$ such that $\tilde{f}_z(\tilde{X}) = \mathbb{E}(Y_z|\tilde{X})$. Now, define $\tilde{f} : \mathcal{X} \times \mathcal{Z} \rightarrow \mathcal{Y}$ by $\tilde{f}(\cdot, z) := \tilde{f}_z$ for λ -a.e. $z \in \mathcal{Z}$. It follows that

$$\begin{aligned} \|Y - \mathbb{E}(Y|\tilde{X}, Z)\|_2^2 &> \int_{\mathcal{Z}} \|Y_z - \tilde{f}_z(\tilde{X})\|_2^2 d\lambda \\ &= \int_{\mathcal{Z}} \|Y_z - \tilde{f}(\tilde{X}, z)\|_2^2 d\lambda \\ &= \|Y - \tilde{f}(\tilde{X}, Z)\|_2^2 \\ &= \|Y - \mathbb{E}(Y|\tilde{X}, Z)\|_2^2 + \|\mathbb{E}(Y|\tilde{X}, Z) - \tilde{f}(\tilde{X}, Z)\|_2^2. \end{aligned}$$

That implies $\|\mathbb{E}(Y|\tilde{X}, Z) - \tilde{f}(\tilde{X}, Z)\|_2^2 < 0$, a contradiction. This completes the proof. \square

Now, the key observation is that, given a fixed $\tilde{X} \perp Z$, the choice of \tilde{Y} depends only on the second term on the right, which forms a Wasserstein barycenter problem with marginals being $\{\mathbb{E}(Y_z|\tilde{X})\}_z$. Hence, the optimal choice of \tilde{Y} is the one which satisfies $\mathbb{E}(\tilde{Y}|\tilde{X}) = \overline{\mathbb{E}(Y|\tilde{X}, Z)}$, where $\overline{\mathbb{E}(Y|\tilde{X}, Z)}$ is the Wasserstein barycenter of $\{\mathbb{E}(Y_z|\tilde{X})\}_z$. Therefore, we denote the optimal choice of \tilde{Y} to be \bar{Y} which satisfies $\mathbb{E}(\bar{Y}|\tilde{X}) = \overline{\mathbb{E}(Y|\tilde{X}, Z)}$.

It remains to find the optimal choice of \tilde{X} . The following result shows that the optimal choice is the one admissible \tilde{X} which generates the finest sigma-algebra.

LEMMA 4.2.2 (Finer sigma-algebra, more accurate optimal fair learning). *Let $\tilde{X}, \tilde{X}' \in \{\tilde{X} \in \mathcal{D}_{\mathcal{X}} : \tilde{X} \perp Z\}$. If $\sigma(\tilde{X}') \subset \sigma(\tilde{X})$, then*

$$(4.14) \quad \|\mathbb{E}(Y|X, Z) - \mathbb{E}(\bar{Y}|\tilde{X})\|_2^2 \leq \|\mathbb{E}(Y|X, Z) - \mathbb{E}(\bar{Y}'|\tilde{X}')\|_2^2$$

where \bar{Y} and \bar{Y}' satisfy $\mathbb{E}(\bar{Y}|\tilde{X}) = \overline{\mathbb{E}(Y|\tilde{X}, Z)}$ and $\mathbb{E}(\bar{Y}'|\tilde{X}') = \overline{\mathbb{E}(Y'|\tilde{X}', Z)}$.

PROOF. Let $\tilde{X}, \tilde{X}' \in \{\tilde{X} \in \mathcal{D}_{\mathcal{X}} : \tilde{X} \perp Z\}$ satisfy $\sigma(\tilde{X}') \subset \sigma(\tilde{X})$. We have

$$\begin{aligned} &\|\mathbb{E}(Y|X, Z) - \mathbb{E}(\bar{Y}|\tilde{X}, Z)\|_2^2 - \|\mathbb{E}(Y|X, Z) - \mathbb{E}(\bar{Y}'|\tilde{X}', Z)\|_2^2 \\ &= \|\mathbb{E}(Y|X, Z) - \overline{\mathbb{E}(Y|\tilde{X}, Z)}\|_2^2 - \|\mathbb{E}(Y|X, Z) - \overline{\mathbb{E}(Y|\tilde{X}', Z)}\|_2^2 \end{aligned}$$

Notice that

$$\|\mathbb{E}(Y|X, Z) - \overline{\mathbb{E}(Y|\tilde{X}, Z)}\|_2^2 = \|\mathbb{E}(Y|X, Z) - \mathbb{E}(Y|\tilde{X}, Z)\|_2^2 + \int_{\mathcal{Z}} \mathcal{W}_2^2(\mu_z, \bar{\mu}) d\lambda$$

where $\mu_z := \mathcal{L}(\mathbb{E}(Y|\tilde{X}, Z)_z)$ and $\bar{\mu} := \overline{\mathcal{L}(\mathbb{E}(Y|\tilde{X}, Z))}$. Also, we define μ'_z and $\bar{\mu}'$ analogously to have

$$\begin{aligned} & \|\mathbb{E}(Y|X, Z) - \overline{\mathbb{E}(Y|\tilde{X}', Z)}\|_2^2 \\ &= \|\mathbb{E}(Y|X, Z) - \mathbb{E}(Y|\tilde{X}', Z)\|_2^2 + \int_{\mathcal{Z}} \mathcal{W}_2^2(\mu'_z, \bar{\mu}') d\lambda \\ &= \|\mathbb{E}(Y|X, Z) - \mathbb{E}(Y|\tilde{X}, Z)\|_2^2 + \|\mathbb{E}(Y|\tilde{X}, Z) - \mathbb{E}(Y|\tilde{X}', Z)\|_2^2 + \int_{\mathcal{Z}} \mathcal{W}_2^2(\mu'_z, \bar{\mu}') d\lambda. \end{aligned}$$

Combining the above, we have

$$\begin{aligned} & \|\mathbb{E}(Y|X, Z) - \mathbb{E}(\bar{Y}|\tilde{X}, Z)\|_2^2 - \|\mathbb{E}(Y|X, Z) - \mathbb{E}(\bar{Y}'|\tilde{X}', Z)\|_2^2 \\ &= \int_{\mathcal{Z}} \mathcal{W}_2^2(\mu_z, \bar{\mu}) d\lambda - \int_{\mathcal{Z}} \mathcal{W}_2^2(\mu'_z, \bar{\mu}') d\lambda - \|\mathbb{E}(Y|\tilde{X}, Z) - \mathbb{E}(Y|\tilde{X}', Z)\|_2^2. \end{aligned}$$

It remains to show that $\int_{\mathcal{Z}} \mathcal{W}_2^2(\mu_z, \bar{\mu}) d\lambda < \int_{\mathcal{Z}} \mathcal{W}_2^2(\mu'_z, \bar{\mu}') d\lambda + \|\mathbb{E}(Y|\tilde{X}, Z) - \mathbb{E}(Y|\tilde{X}', Z)\|_2^2$. Indeed, assume for contradiction that $\int_{\mathcal{Z}} \mathcal{W}_2^2(\mu'_z, \bar{\mu}') d\lambda + \|\mathbb{E}(Y|\tilde{X}, Z) - \mathbb{E}(Y|\tilde{X}', Z)\|_2^2 \leq \int_{\mathcal{Z}} \mathcal{W}_2^2(\mu_z, \bar{\mu}) d\lambda$, then we have

$$\begin{aligned} \int_{\mathcal{Z}} \mathcal{W}_2^2(\mu_z, \bar{\mu}') d\lambda &\leq \|\mathbb{E}(Y|\tilde{X}, Z) - \mathbb{E}(Y|\tilde{X}', Z)\|_2^2 + \int_{\mathcal{Z}} \mathcal{W}_2^2(\mu'_z, \bar{\mu}') d\lambda \\ &\leq \int_{\mathcal{Z}} \mathcal{W}_2^2(\mu_z, \bar{\mu}) d\lambda. \end{aligned}$$

This contradicts the optimality and uniqueness of $\bar{\mu}$ by Lemma 5.3.1. Therefore, we prove by contradiction that $\int_{\mathcal{Z}} \mathcal{W}_2^2(\mu_z, \bar{\mu}) d\lambda < \int_{\mathcal{Z}} \mathcal{W}_2^2(\mu'_z, \bar{\mu}') d\lambda + \|\mathbb{E}(Y|\tilde{X}, Z) - \mathbb{E}(Y|\tilde{X}', Z)\|_2^2$ and, hence,

$$\|\mathbb{E}(Y|X, Z) - \mathbb{E}(\bar{Y}|\tilde{X}, Z)\|_2^2 - \|\mathbb{E}(Y|X, Z) - \mathbb{E}(\bar{Y}'|\tilde{X}', Z)\|_2^2 < 0.$$

That completes the proof. □

Therefore, it is clear that our optimal choice of \tilde{X} is the one that generates the finest sigma-algebra while satisfying $\tilde{X} \perp Z$. The following technical lemma shows that the barycenter of $\{X_z\}_{z \in \mathcal{Z}}$ is one of the optimal choices.

LEMMA 4.2.3 (\bar{X} generates the finest sigma-algebra among admissible). *If $\{\mathcal{L}(X_z)\}_z \subset \mathcal{P}_{2,ac}(\mathcal{X})$ λ -a.e., then $\sigma((\bar{X}, Z)) = \sigma((X, Z))$. In addition, $\sigma(\tilde{X}) \subset \sigma(\bar{X})$ for all $\tilde{X} \in \{\tilde{X} \in \mathcal{D}|\mathcal{X} : \tilde{X} \perp Z\}$.*

PROOF. We first prove $\sigma((\bar{X}, Z)) = \sigma((X, Z))$. Since $\mathcal{L}(X_z) \subset \mathcal{P}_{2,ac}$, it follows from Lemma 5.3.1 that there exists a measurable map $T : \mathcal{X} \times \mathcal{Z} \rightarrow \mathcal{X}$ such that $T(X_z, z) = \bar{X}_z$ λ -a.e., where \bar{X} denotes the Wasserstein barycenter of $\{X_z\}_z$. Define $T \otimes Id|_{\mathcal{Z}} : \mathcal{X} \times \mathcal{Z} \rightarrow \mathcal{X} \times \mathcal{Z}$, we have $T \otimes Id|_{\mathcal{Z}}$ is $\mathcal{X} \times \mathcal{Z} / \mathcal{X} \times \mathcal{Z}$ -measurable and satisfies $T \otimes Id|_{\mathcal{Z}}((X, Z)) = (\bar{X}, Z)$. That implies $\sigma((\bar{X}, Z)) \subset \sigma((X, Z))$. Furthermore, since $\mathcal{L}(\bar{X}) \in \mathcal{P}_{2,ac}$, it follows from Brenier's theorem [12] that there exists $T^{-1}(\cdot, z)$ such that $T^{-1}(\bar{X}_z, z) = X_z$. Therefore, we have $(T \otimes Id|_{\mathcal{Z}})^{-1} = T^{-1} \otimes Id|_{\mathcal{Z}}$ is $\mathcal{X} \times \mathcal{Z} / \mathcal{X} \times \mathcal{Z}$ -measurable and satisfies $(T \otimes Id|_{\mathcal{Z}})^{-1}((\bar{X}, Z)) = (X, Z)$. That implies $\sigma((X, Z)) \subset \sigma((\bar{X}, Z))$. That completes the proof of $\sigma((\bar{X}, Z)) = \sigma((X, Z))$. Now, we show $\sigma(\tilde{X}) \subset \sigma(\bar{X})$. From the construction of \tilde{X} , we have $\sigma((\tilde{X}, Z)) \subset \sigma((\bar{X}, Z)) = \sigma((X, Z))$. But $\tilde{X} \perp Z$ implies that, for any $B_X \in \mathcal{B}_{\mathcal{X}}$, we can construct $B_X \times \mathcal{Z} \in \mathcal{B}_{\mathcal{X}} \otimes \mathcal{B}_{\mathcal{Z}}$. In addition, due to $\sigma((\tilde{X}, Z)) \subset \sigma((\bar{X}, Z))$, there exists $B'_{XZ} \in \mathcal{B}_{\mathcal{X}} \otimes \mathcal{B}_{\mathcal{Z}}$ such that $(\bar{X}, Z)^{-1}(B'_{XZ}) = (X, Z)^{-1}(B_X \times \mathcal{Z})$. Lastly, $\bar{X} \perp Z$ also implies that there exists $B'_X \in \mathcal{B}_{\mathcal{X}}$ satisfying $B'_{XZ} = B'_X \times \mathcal{Z}$. It follows that

$$(4.15) \quad \tilde{X}^{-1}(B_X) = (\tilde{X}, Z)^{-1}(B_X \times \mathcal{Z}) = (X, Z)^{-1}(B'_X \times \mathcal{Z}) = X^{-1}(B'_X)$$

Since our choice of $B_X \in \mathcal{B}_{\mathcal{X}}$ is arbitrary, it follows that $\sigma(\tilde{X}) \subset \sigma(\bar{X})$. Finally, since our choice of $\tilde{X} \in \{\tilde{X} \in \mathcal{D}|\mathcal{X} : \tilde{X} \perp Z\}$ is arbitrary, we are done. \square

Therefore, Lemma 4.2.2, Lemma 4.2.3, and the choice of \bar{Y} above together provide a characterization of the solution to Problem 3.

LEMMA 4.2.4 (Characterization of optimal fair data representation). *Let \bar{X} and $\overline{\mathbb{E}(Y|\bar{X}, Z)}$ denote the respective Wasserstein barycenter of $\{X_z\}_z$ and $\{\mathbb{E}(Y_z|\bar{X})\}_z$. If $\{\mathcal{L}(X_z)\}_z \subset \mathcal{P}_{2,ac}(\mathcal{X})$ and $\{\mathcal{L}(\mathbb{E}(Y|\bar{X}, Z)_z)\}_z \subset \mathcal{P}_{2,ac}(\mathcal{Y})$, then the following are equivalent:*

- $(\tilde{X}, \tilde{Y}) \in \arg \min_{(\tilde{X}, \tilde{Y}) \in \mathcal{D}} \{\|Y - \mathbb{E}(\tilde{Y}|\tilde{X})\|_2^2 : \tilde{X}, \mathbb{E}(\tilde{Y}|\tilde{X}, Z) \perp Z\}$.
- $(\tilde{X}, \tilde{Y}) \in \{(\tilde{X}, \tilde{Y}) \in \mathcal{D} : \sigma(\tilde{X}) = \sigma(\bar{X}), \mathbb{E}(\tilde{Y}|\bar{X}) = \overline{\mathbb{E}(Y|\bar{X}, Z)}\}$.

In Lemma 4.2.4, the choice of \bar{X} is not unique. In fact, any random variable \tilde{X} that satisfies $\sigma(\tilde{X}) = \sigma(\bar{X})$ can be our choice according to Lemma 4.2.2 and Lemma 4.2.3. This is because any

\tilde{X} that satisfies the above conditions gives $\mathbb{E}(Y|\tilde{X}) = \mathbb{E}(Y|\bar{X})$. For both theoretical and computational convenience, we fix our choice to be \bar{X} from now on.

REMARK 4.2.1 (Application of the optimal fair representation characterization to algorithm design). *In theory, we should always take \bar{X} because we prove that \bar{X} generates the finest sigma-algebra among all the admissible \tilde{X} that is independent of Z . Especially when working with data sets with clear high-dimensional structure such as image data, one should apply more complicated models to estimate the optimal transport map instead of using affine maps. But when working with data with less high-dimensional structure such as tabular data, we hope to take advantage of the simplicity, robustness, and interpretability of linear maps in practice and hence restrict the admissible transport maps to be affine, as mentioned in Remark 3.2.1. Therefore, we showed that the pseudo-barycenter X^\dagger , which is equal to \bar{X} in the Gaussian case and solves a relaxed version of the barycenter problem in the general distribution case, can be achieved using optimal affine maps. As a result, we apply X^\dagger in the algorithm design and experiments. Still, if there is no concern about over-fitting or computational cost, it is recommended for strict statistical parity guarantee purposes to compute \bar{X} to improve the result.*

Now, it remains to find \bar{Y} to obtain the optimal fair data representation characterized by Lemma 4.2.4. In general, it is difficult to find $\overline{\mathbb{E}(Y|\bar{X}, Z)}$, not to mention find a \tilde{Y} satisfying $\mathbb{E}(\tilde{Y}|\bar{X}) = \overline{\mathbb{E}(Y|\bar{X}, Z)}$. The key observation here is that if the Brenier's maps $\{T_{y|\bar{X}}(\cdot, z)\}_z$ that push $\{\mathbb{E}(Y_z|\bar{X})\}_z$ forward to $\overline{\mathbb{E}(Y|\bar{X}, Z)}$ are affine, then a straight-forward choice in \bar{Y} is $\{T_{y|\bar{X}}(Y_z, z)\}_{z \in \mathcal{Z}} = T_{y|\bar{X}}(Y, Z)$. This step is the key to circumvent the post-processing nature. Therefore, following the same derivation of (3.8) from (5.5) in Section 3.1 to guarantee feasibility of affine maps, we relax the fairness constraint to the first two moments in Problem 3, and show a pseudo-barycenter pair provides us an exact solution to Problem 3 in the Gaussian marginal case and the optimal affine estimation in the general marginal case.

4.3. Gaussian Marginals: Exact Solution

Assume $\{(X_z, Y_z)\}_z$ to be non-degenerate Gaussian vectors λ -a.e. and define the following:

DEFINITION 4.3.1 (Independent pseudo-barycenter: X^\dagger).

$$(4.16) \quad X^\dagger := T_x(X, Z),$$

where

$$(4.17) \quad T_x(\cdot, z) := \Sigma_{X_z}^{-\frac{1}{2}} (\Sigma_{X_z}^{\frac{1}{2}} \Sigma \Sigma_{X_z}^{\frac{1}{2}})^{\frac{1}{2}} \Sigma_{X_z}^{-\frac{1}{2}}$$

and Σ is the unique solution to

$$(4.18) \quad \int_{\mathcal{Z}} (\Sigma^{\frac{1}{2}} \Sigma_{X_z} \Sigma^{\frac{1}{2}})^{\frac{1}{2}} d\lambda(z) = \Sigma.$$

DEFINITION 4.3.2 (Dependent pseudo-barycenter: Y^\dagger).

$$(4.19) \quad Y^\dagger := T_{y|X^\dagger}(Y, Z)$$

where

$$(4.20) \quad T_{y|X^\dagger}(\cdot, z) := \Sigma_{Y_z|X^\dagger}^{-\frac{1}{2}} (\Sigma_{Y_z|X^\dagger}^{\frac{1}{2}} \Sigma \Sigma_{Y_z|X^\dagger}^{\frac{1}{2}})^{\frac{1}{2}} \Sigma_{Y_z|X^\dagger}^{-\frac{1}{2}}$$

with $\Sigma_{Y_z|X^\dagger} := \Sigma_{Y_z X^\dagger} \Sigma_{X^\dagger}^{-1} \Sigma_{Y_z X^\dagger}^T$, and Σ is the unique solution to

$$(4.21) \quad \int_{\mathcal{Z}} (\Sigma^{\frac{1}{2}} \Sigma_{Y_z|X^\dagger} \Sigma^{\frac{1}{2}})^{\frac{1}{2}} d\lambda(z) = \Sigma$$

Here, to obtain (an estimation of) the solution to equations (4.21) and (4.18), we apply the iterative method (2.15) in Remark 2.2.2 when designing our algorithm in Section 4.6.

Since it is a direct result of Lemma 2.2.2 that $X^\dagger = \overline{X}$, the goal is now to show that

$$(4.22) \quad \mathbb{E}(Y^\dagger | \overline{X}) = \overline{\mathbb{E}(Y | \overline{X}, Z)},$$

and therefore by Lemma 4.2.4 to conclude $\mathbb{E}(Y^\dagger | X^\dagger) = \mathbb{E}(Y^\dagger | \overline{X})$ indeed minimizes the estimation error while staying independent of Z .

To prove the above equation and justify the definition of the pseudo-barycenter, we need the following results: (1) existence and uniqueness of both \overline{X} and $\overline{\mathbb{E}(Y|\overline{X}, Z)}$; (2) affinity of the corresponding Brenier's maps $T_x(\cdot, z)$ and $T_{y|X^\dagger}(\cdot, z)$. By assumption, we have $\{\mathcal{L}(X_z)\}_z \subset \mathcal{P}_{2,ac}(\mathcal{X})$, and $\{\mathcal{L}(\mathbb{E}(Y_z|\overline{X}))\}_z \subset \mathcal{P}_{2,ac}(\mathcal{Y})$. The existence and uniqueness then follow directly from Lemma 2.1.1. It remains to show that the corresponding Brenier's maps are affine. But by Lemma 2.2.2, if $\{X_z\}_z$ and $\{\mathbb{E}(Y_z|\overline{X})\}_z$ both are from some location-scale family, then the barycenters are also from the corresponding location-scale family and the Brenier's maps are affine.

The following result shows that if $\{Y_z\}_z$ come from the same location-scale family, then $\{\mathbb{E}(Y_z|\overline{X})\}_z$ also belongs to the same location-scale family.

LEMMA 4.3.1 (Conditional expectation preserves location-scale family). *Assume that $\{Y_z\}_z \subset \mathcal{F}(P_0)$ for some P_0 , then $\{\mathbb{E}(Y_z|\overline{X})\}_z \subset \mathcal{F}(\mathcal{L}(\mathbb{E}(Y_z|\overline{X})))$ for any z .*

PROOF. This follows immediately from the existence of positive definite affine transformations among $\{Y_z\}_z$, Lemma 2.2.1, and the linearity of conditional expectation. \square

Therefore, given $\{(X_z, Y_z)\}_z$ being Gaussian vectors, we have $\{(\overline{X}, Y_z)\}$ being Gaussian vectors, which further implies that $\{\mathbb{E}(Y_z|\overline{X})\}_z$ are Gaussian vectors by Lemma 4.3.1. (We note that it is not necessary to apply Lemma 4.3.1 to show $\{\mathbb{E}(Y_z|\overline{X})\}_z$ are Gaussian because it is a well-known result in probability theory, but the lemma becomes necessary later in the case of general marginal distributions.)

LEMMA 4.3.2 (Solution to the optimal fair data representation in the Gaussian case). *Let $\{(X_z, Y_z)\}_z$ be Gaussian vectors satisfying $\Sigma_z \succ 0$ λ -a.e., then there exists a unique barycenter pair $(\overline{X}, \overline{\mathbb{E}(Y|\overline{X}, Z)})$ which are Gaussian vectors characterized by the covariance matrix being the unique solution to*

$$(4.23) \quad \int_{\mathcal{Z}} (\Sigma^{\frac{1}{2}} S \Sigma^{\frac{1}{2}})^{\frac{1}{2}} d\lambda(z) = \Sigma$$

for $S \in \{\Sigma_{X_z}, \Sigma_{Y_z|X^\dagger}\}$ respectively, where $\Sigma_{Y_z|X^\dagger} = \Sigma_{Y_z X^\dagger} \Sigma_{X^\dagger}^{-1} \Sigma_{Y_z X^\dagger}^T$. Moreover, $\{T_x(\cdot, z)\}_z$ and $\{T_{y|X^\dagger}(\cdot, z)\}_z$ which push X_z and $\mathbb{E}(Y_z|\overline{X})$ respectively to \overline{X} and $\overline{\mathbb{E}(Y|\overline{X}, Z)}$ are affine with closed-form (4.17) and (4.20). As a result, for λ -a.e. $z \in \mathcal{Z}$, we have

$$(4.24) \quad \overline{\mathbb{E}(Y|\overline{X}, Z)}_z = T_{y|X^\dagger}(\mathbb{E}(Y_z|T_x(X_z, z)), z) = \mathbb{E}(T_{y|X^\dagger}(Y_z, z)|T_x(X_z, z))$$

PROOF. The existence, uniqueness, and Gaussianity of the barycenter follow from Lemma 2.2.2, whereas the affinity of corresponding Brenier's maps results from Lemmas 4.3.1 and 2.2.1. \square

The above result provides us a theoretical foundation to apply the affine maps $\{T_x(\cdot, z)\}_z$ and $\{T_{y|X^\dagger}(\cdot, z)\}_z$ to $\{X_z\}_z$ and $\{Y_z\}_z$ respectively as a pre-processing step before the training step.

Furthermore, notice that although $T_{y|X^\dagger}(\mathbb{E}(Y_z|\bar{X}), z) = \overline{\mathbb{E}(Y_z|\bar{X}, Z)}_z$ λ -a.e. by construction, $\{T_{y|X^\dagger}(Y_z, z)\}_z$ does not agree in general: for $z_1 \neq z_2$,

$$(4.25) \quad T_{y|X^\dagger}(Y_{z_1}, z_1) \neq T_{y|X^\dagger}(Y_{z_2}, z_2).$$

The pseudo-barycenter solves the disagreement by merging them directly. Despite of the differences among $\{T_{y|X^\dagger}(Y_z, z)\}_z$, the L^2 projections of them on $\sigma(\bar{X})$ agree. Therefore, a direct merging of $\{T_{y|X^\dagger}(Y_z, z)\}_z$ is simply: $T_{y|X^\dagger}(Y, Z) = Y^\dagger$. It follows:

$$\begin{aligned} \mathbb{E}(Y^\dagger|X^\dagger) &= \mathbb{E}(Y^\dagger|\bar{X}) = \mathbb{E}(T_{y|X^\dagger}(Y, Z)|\bar{X}) \\ &= \int_{\mathcal{Z}} \mathbb{E}(T_{y|X^\dagger}(Y_z, z)|\bar{X}) d\lambda(z) \\ &= \int_{\mathcal{Z}} T_{y|X^\dagger}(\mathbb{E}(Y_z|\bar{X}), z) d\lambda(z) \\ &= \int_{\mathcal{Z}} T_{y|X^\dagger}(\mathbb{E}(Y|\bar{X}, Z)_z, z) d\lambda(z) \\ &= \int_{\mathcal{Z}} \overline{\mathbb{E}(Y|\bar{X}, Z)}_z d\lambda(z) = \overline{\mathbb{E}(Y|\bar{X}, Z)}, \end{aligned}$$

where the second equality follows from disintegration, the third from linearity of $T_{y|\bar{X}}$, and the fourth from $\mathbb{E}(Y_z|\bar{X}) = \mathbb{E}(Y|\bar{X}, Z)_z$. Therefore, we have proved a result that justifies the definition of the pseudo-barycenter:

THEOREM 4.3.1 (Justification of Y^\dagger in Gaussian case). (X^\dagger, Y^\dagger) is a solution to Problem 3

$$(4.26) \quad \inf_{(\tilde{X}, \tilde{Y}) \in \mathcal{D}} \{ \|Y - \mathbb{E}(\tilde{Y}|\tilde{X})\|_2^2 : \tilde{X}, \mathbb{E}(\tilde{Y}|\tilde{X}, Z) \perp Z \},$$

if $\{(X_z, Y_z)\}_z$ are non-degenerate Gaussian vectors.

4.4. General Distribution: Optimal Affine Estimation

In practice, one should not always expect the sensitive marginal data distributions to be Gaussian, and the results we derived under the assumption of Gaussianity may not apply to the general marginal distribution case. Instead, we solve the following relaxed optimal fair data representation problem:

$$(4.27) \quad \inf_{(\tilde{X}, \tilde{Y}) \in \mathcal{D}} \{ \|Y - \mathbb{E}(\tilde{Y}|\tilde{X})\|_2^2 : m_{\tilde{X}}, m_{\tilde{Y}|\tilde{X}}, \Sigma_{\tilde{X}}, \Sigma_{\tilde{Y}|\tilde{X}} \perp Z \},$$

where $m_{\tilde{Y}|\tilde{X}} := \mathbb{E}(\mathbb{E}(\tilde{Y}|\tilde{X}, Z))$ and similarly for $\Sigma_{\tilde{Y}|\tilde{X}}$, to find the optimal affine estimation of the true solution to the original Problem 3. The fairness guarantee of the affine estimation is the same as mentioned in Remark 3.2.2.

Now, we justify the pseudo-barycenter pair (X^\dagger, Y^\dagger) in the case of general distributions by proving it is a solution to the relaxed optimal fair L^2 -objective supervised learning problem (4.27). To start, notice that $(X^\dagger, Y^\dagger) \in \mathcal{D}$ and satisfies $m_{X^\dagger}, m_{Y^\dagger|X^\dagger}, \Sigma_{X^\dagger}, \Sigma_{Y^\dagger|X^\dagger} \perp Z$ by construction and therefore is admissible.

REMARK 4.4.1 (Finest sigma-algebra vs. most variance). *Due to the relaxation, the admissible $\tilde{X} \in \mathcal{D}|_{\mathcal{X}}$ are no longer required to be independent of Z . Furthermore, without the assumption of Gaussianity, X^\dagger is no longer equal to \bar{X} . As a result, although one can still prove $\sigma((X, Z)) = \sigma((X^\dagger, Z))$ by following the same argument in the proof of Lemma 4.2.3 as in the Gaussian case, but this fact now cannot imply $\sigma(\tilde{X}) \subset \sigma(X^\dagger)$ due to the lack of independence condition. Instead, the present work shows that $\text{Var}(\tilde{X}) \leq \text{Var}(X^\dagger)$ for all admissible $\tilde{X} \in \mathcal{D}|_{\mathcal{X}}$, which in general implies $\sigma(\tilde{X}) \subset \sigma(X^\dagger)$. For example, whenever set inclusion forms an order between $\sigma(\tilde{X})$ and $\sigma(X^\dagger)$, then it is true that $\text{Var}(\tilde{X}) \leq \text{Var}(X^\dagger)$ implies $\sigma(\tilde{X}) \subset \sigma(X^\dagger)$. As a result, we still fix X^\dagger as our optimal choice among all the admissible $\tilde{X} \in \mathcal{D}|_{\mathcal{X}}$.*

In addition, for any $\Sigma \succ 0$, define

$$(4.28) \quad T_{\Sigma, x} := \Sigma_{X_z}^{-\frac{1}{2}} (\Sigma_{X_z}^{\frac{1}{2}} \Sigma \Sigma_{X_z}^{\frac{1}{2}})^{\frac{1}{2}} \Sigma_{X_z}^{-\frac{1}{2}}$$

$$(4.29) \quad T_\Sigma := \Sigma_{Y_z|X_z^\dagger}^{-\frac{1}{2}} (\Sigma_{Y_z|X_z^\dagger}^{\frac{1}{2}} \Sigma_{Y_z|X_z^\dagger}^{\frac{1}{2}})^{\frac{1}{2}} \Sigma_{Y_z|X_z^\dagger}^{-\frac{1}{2}}$$

where $\Sigma_{Y_z|X_z^\dagger} := \mathbb{E}((\mathbb{E}(Y_z|X_z^\dagger) - m_{Y_z})(\mathbb{E}(Y_z|X_z^\dagger) - m_{Y_z})^T)$ and $\mathbb{E}(Y_z|X_z^\dagger) := \mathbb{E}(Y|X^\dagger, Z)_z$. Now, the goal is to show (X^\dagger, Y^\dagger) is indeed a solution to the relaxed problem (4.27), under the following two assumptions:

- 1 Set inclusion forms an order between X^\dagger and all $\tilde{X} \in \{\tilde{X} \in \mathcal{D}|_{\mathcal{X}} : m_{\tilde{X}}, \Sigma_{\tilde{X}} \perp Z\}$.
- 2 $\Sigma_{Y_z|X_z^\dagger} = \Sigma_{Y_z X_z^\dagger} \Sigma_{X_z^\dagger}^{-1} \Sigma_{Y_z X_z^\dagger}^T$.

REMARK 4.4.2 (Applicability of the assumptions). *For the first assumption, Lemma 4.4.1 below guarantees that X^\dagger generates the finest sigma-algebra among all the admissible sigma-algebras. In other words, for any admissible \tilde{X} , either it generates a coarser sigma-algebra than $\sigma(X^\dagger)$ or the two sigma-algebras do not contain each other. In other words, there is no admissible \tilde{X} such that $\sigma(X^\dagger) \subset \sigma(\tilde{X})$.*

The second assumption allows us to directly compute the covariance matrix of $\mathbb{E}(Y_z|X_z^\dagger)$ from $\Sigma_{Y_z X_z^\dagger}$ and $\Sigma_{X_z^\dagger}$. The second assumption is necessary to keep our pre-processing approach. In general, $\mathbb{E}(Y_z|X_z^\dagger)$ is not a linear function of X_z^\dagger as in the Gaussian case. When the second assumption is not true, our pre-processing approach uses $\Sigma_{Y_z X_z^\dagger} \Sigma_{X_z^\dagger}^{-1} \Sigma_{Y_z X_z^\dagger}^T$ as our best affine estimate of $\Sigma_{Y_z|X_z^\dagger}$.

To that end, we need the following result on the relationship among the variance of the original distribution, the variance of the barycenter, and the Wasserstein distance.

LEMMA 4.4.1 (Variance reduction of Wasserstein barycenter [49]). *Given X satisfies $\{\mathcal{L}(X_z)\}_z \subset \mathcal{P}_{2,ac}(\mathcal{X})$ and \bar{X} satisfies $\mathcal{L}(\bar{X})$ being the Wasserstein barycenter of $\{\mathcal{L}(X_z)\}$, it follows that*

$$(4.30) \quad \|X - \mathbb{E}(X)\|_2^2 - \|\bar{X} - \mathbb{E}(\bar{X})\|_2^2 = \int_{\mathcal{Z}} \mathcal{W}_2^2(\mathcal{L}(X_z), \mathcal{L}(\bar{X})) d\lambda(z)$$

As a result, we obtain the following:

LEMMA 4.4.2 (X^\dagger Contains the largest variance among admissible). *X^\dagger is the unique solution to*

$$(4.31) \quad \sup_{\tilde{X} \in \mathcal{D}|_{\mathcal{X}}} \{\text{Var}(\tilde{X}) : m_{\tilde{X}}, \Sigma_{\tilde{X}} \perp Z\}.$$

PROOF. To simplify notation, by the invariance of variance under translation and Lemma 2.1.1, we can assume without loss of generality that $m_{X_z} = 0$ $\lambda - a.e.$ in the rest of the proof, which only deal with variance and Wasserstein distance. Now, for $\lambda - a.e.$ $z \in \mathcal{Z}$, we have

$$\begin{aligned} \|X_z - T_{\Sigma, x}(X_z, z)\|_2^2 &= \|X_z\|_2^2 + \|T_{\Sigma, x}(X_z, z)\|_2^2 - 2\langle X_z, T_{\Sigma, x}(X_z, z) \rangle_2 \\ &= \text{Trace}(\Sigma_{X_z}) + \text{Trace}(\Sigma) - 2\mathbb{E}(X_z^T T_{\Sigma, x}(X_z, z)) \\ &= \text{Trace}(\Sigma_{X_z}) + \text{Trace}(\Sigma) - 2\langle T_{\Sigma, x}, \Sigma_{X_z} \rangle_F \\ &= \text{Trace}(\Sigma_{X'_z}) + \text{Trace}(\Sigma) - 2\langle T_{\Sigma, x}, \Sigma_{X'_z} \rangle_F \\ &= \|X'_z - T_{\Sigma, x}(X'_z, z)\|_2^2 \\ &= \mathcal{W}_2^2(\mathcal{L}(X'_z), \mathcal{L}(T_{\Sigma, x}(X'_z))) \end{aligned}$$

where $X' \sim \mathcal{N}(m_X, \Sigma_X)$ is the Gaussian analog of X and $\langle \cdot, \cdot \rangle_F$ is the Frobenius inner product. Similarly, by the disintegration theorem, we also have for $S \in \{X, X^\dagger\}$

$$(4.32) \quad \text{Var}(S) = \|S\|_2^2 = \int_{\mathcal{Z}} \|S_z\|_2^2 d\lambda = \int_{\mathcal{Z}} \text{Trace}(\Sigma_{S_z}) d\lambda.$$

Therefore, it follows from Lemma 4.4.1 that

$$\begin{aligned} \text{Var}(X) - \text{Var}(X^\dagger) &= \text{Var}(X') - \text{Var}((X')^\dagger) \\ &= \text{Var}(X') - \text{Var}(\overline{X'}) \\ &= \int_{\mathcal{Z}} \mathcal{W}_2^2(\mathcal{L}(X'_z), \mathcal{L}(\overline{X'})) d\lambda(z). \end{aligned}$$

Finally, assume there exists a $\tilde{X} \in \mathcal{D}|_{\mathcal{X}}$ such that $\text{Var}(X^\dagger) \leq \text{Var}(\tilde{X})$. It follows $\text{Var}(X') - \text{Var}(\tilde{X}') \leq \text{Var}(X') - \text{Var}((X')^\dagger) = \text{Var}(X') - \text{Var}(\overline{X'})$. But since $m_{\tilde{X}'}, \Sigma_{\tilde{X}'} \perp Z$, we have $\tilde{X}' \perp Z$ as \tilde{X}' is Gaussian by construction. In other words, there exists a $\tilde{X}' \perp Z$ such that

$$(4.33) \quad \int_{\mathcal{Z}} \mathcal{W}_2^2(\mathcal{L}(X'_z), \mathcal{L}(\tilde{X}')) d\lambda(z) \leq \int_{\mathcal{Z}} \mathcal{W}_2^2(\mathcal{L}(X'_z), \mathcal{L}(\overline{X}')) d\lambda(z)$$

which contradicts the uniqueness of \overline{X}' . \square

The above lemma shows that $\text{Var}(\tilde{X}) \leq \text{Var}(X^\dagger)$ for all admissible $\tilde{X} \in \mathcal{D}|_{\mathcal{X}}$ satisfies $m_{\tilde{X}}, \Sigma_{\tilde{X}} \perp Z$, which together with the first assumption imply $\sigma(\tilde{X}) \subset \sigma(\overline{X})$ in practice. Therefore, from now on, we fix the choice of \tilde{X} to be X^\dagger and prove the general characterization result based on the two assumptions listed above.

It remains to justify the choice of Y^\dagger . To do so, we need the following lemma, which provides a multi-marginal characterization of the optimal affine map.

LEMMA 4.4.3 (Projection Lemma for conditional expectations). *Given $m_{Y_z|X_z^\dagger} = 0$ and $\Sigma_{Y_z|X_z^\dagger} \succ 0$ λ -a.e., for any $\Sigma \succ 0$,*

$$(4.34) \quad \inf_{\mathbb{E}(\tilde{Y}|X^\dagger): \Sigma_{\tilde{Y}_z|X_z^\dagger} = \Sigma} \int_{\mathcal{Z}} \mathcal{W}_2^2(\mathcal{L}(\mathbb{E}(Y_z|X_z^\dagger)), \mathcal{L}(\mathbb{E}(\tilde{Y}_z|X_z^\dagger))) d\lambda(z)$$

admits a unique solution, denoted by Y_Σ^\dagger , that has the form

$$(4.35) \quad Y_\Sigma^\dagger := T_\Sigma(Y, Z)$$

where $T_\Sigma(\cdot, z) := \Sigma_{\tilde{Y}_z|X_z^\dagger}^{-\frac{1}{2}} (\Sigma_{\tilde{Y}_z|X_z^\dagger}^{\frac{1}{2}} \Sigma_{Y_z|X_z^\dagger} \Sigma_{\tilde{Y}_z|X_z^\dagger}^{\frac{1}{2}})^{\frac{1}{2}} \Sigma_{\tilde{Y}_z|X_z^\dagger}^{-\frac{1}{2}}$

PROOF. This is a direct corollary from Lemma 3.2.2. \square

Finally, we are ready to prove the justification of the pseudo-barycenter in the case of general distributions.

THEOREM 4.4.1 (Justification of (X^\dagger, Y^\dagger) in general distribution case). $\mathbb{E}(Y^\dagger|X^\dagger)$ is a solution to

$$(4.36) \quad \inf_{(\tilde{X}, \tilde{Y}) \in \mathcal{D}} \{ \|Y - \mathbb{E}(\tilde{Y}|\tilde{X})\|_2^2 : m_{\tilde{X}}, m_{\tilde{Y}|\tilde{X}}, \Sigma_{\tilde{X}}, \Sigma_{\tilde{Y}|\tilde{X}} \perp Z \}$$

under the assumptions: (1) set inclusion forms an order between X^\dagger and all $\tilde{X} \in \{\tilde{X} \in \mathcal{D}|_{\mathcal{X}} : m_{\tilde{X}}, \Sigma_{\tilde{X}} \perp Z\}$; and (2) $\Sigma_{Y_z|X_z^\dagger} = \Sigma_{Y_z X_z^\dagger} \Sigma_{X_z^\dagger}^{-1} \Sigma_{Y_z X_z^\dagger}^T$.

PROOF. The choice of X^\dagger follows from the first assumption and Lemma 4.4.2. It remains to show that Y^\dagger is a solution to

$$(4.37) \quad \inf_{\tilde{Y} \in \mathcal{D}|_{\mathcal{Y}}} \{ \|Y - \mathbb{E}(\tilde{Y}|X^\dagger)\|_2^2 : m_{\tilde{Y}|X^\dagger}, \Sigma_{\tilde{Y}|X^\dagger} \perp Z \}$$

Fix $\Sigma \succ 0$ arbitrary, we have

$$(4.38) \quad \|Y - \mathbb{E}(Y_\Sigma^\dagger|X^\dagger)\|_2^2 - \|Y - \mathbb{E}(Y|X^\dagger)\|_2^2 = \int_{\mathcal{Z}} \|\mathbb{E}(Y_z - Y_{\Sigma,z}^\dagger|X_z^\dagger)\|_2^2 d\lambda(z)$$

and it follows from Lemma 4.4.3 that

$$\begin{aligned} \int_{\mathcal{Z}} \|\mathbb{E}(Y_z - Y_{\Sigma,z}^\dagger|X_z^\dagger)\|_2^2 d\lambda(z) &= \int_{\mathcal{Z}} \mathcal{W}_2^2(\mathcal{L}(\mathbb{E}(Y_z|X_z^\dagger)), \mathcal{L}(T_\Sigma(\mathbb{E}(Y_z|X_z^\dagger), z))) d\lambda(z) \\ &= \min_{\nu: \Sigma_{\nu_z} = \Sigma} \int_{\mathcal{Z}} \mathcal{W}_2^2(\mathcal{L}(\mathbb{E}(Y_z|X_z^\dagger)), \nu_z) d\lambda(z) \end{aligned}$$

Therefore, (4.27) boils down to the following:

$$(4.39) \quad \inf_{\Sigma \succ 0} \left\{ \int_{\mathcal{Z}} \|\mathbb{E}(Y_z - Y_{\Sigma,z}^\dagger|X_z^\dagger)\|_2^2 d\lambda(z) \right\}.$$

Finally, notice that

$$\begin{aligned}
& \int_{\mathcal{Z}} \|\mathbb{E}(Y_z - Y_{\Sigma, z}^\dagger | X_z^\dagger)\|_2^2 d\lambda(z) \\
&= \int_{\mathcal{Z}} \|\mathbb{E}(Y_z | X_z^\dagger) - T_\Sigma(\mathbb{E}(Y_z | X_z^\dagger), z)\|_2^2 d\lambda(z) \\
&= \int_{\mathcal{Z}} \|\mathbb{E}(Y_z | X_z^\dagger)\|_2^2 + \|T_\Sigma(\mathbb{E}(Y_z | X_z^\dagger), z)\|_2^2 - 2\langle \mathbb{E}(Y_z | X_z^\dagger), T_\Sigma(\mathbb{E}(Y_z | X_z^\dagger), z) \rangle_2 d\lambda(z) \\
&= \int_{\mathcal{Z}} \text{Trace}(\Sigma_{Y_z | X_z^\dagger}) + \text{Trace}(\Sigma) - 2\mathbb{E}(\mathbb{E}(Y_z | X_z^\dagger)^T T_\Sigma(\mathbb{E}(Y_z | X_z^\dagger), z)) d\lambda(z) \\
&= \int_{\mathcal{Z}} \text{Trace}(\Sigma_{Y_z | X_z^\dagger}) + \text{Trace}(\Sigma) - 2\langle T_\Sigma, \Sigma_{Y_z | X_z^\dagger} \rangle_F d\lambda(z) \\
&= \int_{\mathcal{Z}} \|\mathbb{E}(Y_z | X_z^\dagger)' - T_\Sigma(\mathbb{E}(Y_z | X_z^\dagger)', z)\|_2^2 d\lambda(z)
\end{aligned}$$

where $\langle \cdot, \cdot \rangle_F$ denotes the Frobenius inner product and $X' \sim \mathcal{N}(m_X, \Sigma_X)$ denotes the Gaussian analog of X . It follows from the definition of Y^\dagger and Lemma 2.2.2 that $\int_{\mathcal{Z}} \|\mathbb{E}(Y_z - Y_z^\dagger | X^\dagger)\|_2^2 d\lambda(z)$ is the lower bound of (4.39). The proof is complete. \square

To conclude, given an arbitrary L^2 -objective supervised learning model that aims to estimate conditional expectation, the training via (X^\dagger, Y^\dagger) results in an estimate of $\overline{\mathbb{E}(Y | \overline{X}, Z)}$. In other words, any supervised learning model trained via (X^\dagger, Y^\dagger) is guaranteed to be independent of Z in the location-scale family marginal case (or, to have first two moments independent of Z in the general marginal case), while resulting in the minimum prediction error among all the admissible functions of some specific model due to the training step. Here, the assumption is that the test sample distribution is the same as the training sample distribution, which is a ubiquitous assumption for machine learning.

4.5. Optimal Fair Data Representation at the Pareto Frontier

Finally, we extend the pseudo-barycenter pair, which is the solution to the optimal fair data representation, to the fair data representation at the Pareto frontier using McCann interpolation via a similar approach as we derived the post-processing Pareto frontier in Section 3.3. But notice a direct application of Theorem 3.3.1 does not work here because there is no direct interpolation between $E(Y|X, Z)$ and $\overline{\mathbb{E}(Y|\overline{X}, Z)}$ due to the change of the underlying sigma-algebra. Therefore, we apply a diagonal argument, Remark 4.5.1, to estimate the interpolation between $E(Y|X, Z)$ and $\overline{\mathbb{E}(Y|\overline{X}, Z)}$ and thus the fair data representation at the Pareto frontier.

To start, we derive the following post-processing optimal trade-off result directly from Theorem 3.3.1 for a fixed choice of $\tilde{X} \in \{\tilde{X} \in \mathcal{D}_{|\mathcal{X}} : \tilde{X} \perp Z\}$. For any $f : \mathcal{X} \times \mathcal{Z} \rightarrow \mathcal{Y}$, define $L_{y|\tilde{X}}$, $D_{y|\tilde{X}}$, and $V_{y|\tilde{X}}$ as follows:

$$(4.40) \quad L_{y|\tilde{X}}(f(\tilde{X}, Z)) := \left(\int_{\mathcal{Z}} \|\mathbb{E}(Y_z|\tilde{X}) - f(\tilde{X}, Z)_z\|_2^2 d\lambda(z) \right)^{\frac{1}{2}}$$

$$(4.41) \quad D_{y|\tilde{X}}(f(\tilde{X}, Z)) := \left(\int_{\mathcal{Z}^2} \mathcal{W}_2^2(f(\tilde{X}, Z)_{z_1}, f(\tilde{X}, Z)_{z_2}) d\lambda(z_1) d\lambda(z_2) \right)^{\frac{1}{2}}.$$

To simplify notation, for any $T' : \mathcal{Y} \times \mathcal{Z} \rightarrow \mathcal{Y}$, we also define the following:

$$(4.42) \quad L_{y|\tilde{X}}(T') := \left(\int_{\mathcal{Z}} \|\mathbb{E}(Y_z|\tilde{X}) - T'_z(\mathbb{E}(Y_z|\tilde{X}))\|_2^2 d\lambda(z) \right)^{\frac{1}{2}}$$

$$(4.43) \quad D_{y|\tilde{X}}(T') := \left(\int_{\mathcal{Z}^2} \mathcal{W}_2^2((T'_{z_1})_{\#} \mathcal{L}(\mathbb{E}(Y_{z_1}|\tilde{X})), (T'_{z_2})_{\#} \mathcal{L}(\mathbb{E}(Y_{z_2}|\tilde{X}))) d\lambda(z_1) d\lambda(z_2) \right)^{\frac{1}{2}}.$$

Also, let T denote the optimal transport map from $\{\mathbb{E}(Y_z|\tilde{X})\}_z$ to the barycenter $\overline{\mathbb{E}(Y|\tilde{X}, Z)}$, let $T(t), t \in [0, 1]$ be the McCann interpolation, and define

$$(4.44) \quad V_{y|\tilde{X}} := L_{y|\tilde{X}}(T) = \left(\int_{\mathcal{Z}} \|\mathbb{E}(Y_z|\tilde{X}) - T_z(\mathbb{E}(Y_z|\tilde{X}))\|_2^2 d\lambda(z) \right)^{\frac{1}{2}}$$

$$(4.45) \quad = \left(\int_{\mathcal{Z}} \|\mathbb{E}(Y_z|\tilde{X}) - \overline{\mathbb{E}(Y|\tilde{X}, Z)}\|_2^2 d\lambda(z) \right)^{\frac{1}{2}}.$$

Then the result below follows directly similar to the proof of Theorem 3.3.1.

COROLLARY 4.5.1 (Pareto frontier for conditional expectation on fixed sigma-algebra). *Given $L_{y|\tilde{X}}$, $D_{y|\tilde{X}}$, and $V_{y|\tilde{X}}$ defined above, we have*

$$(4.46) \quad V_{y|\tilde{X}} \leq L_{y|\tilde{X}}(f(\tilde{X}, Z)) + \frac{1}{\sqrt{2}} D_{y|\tilde{X}}(f(\tilde{X}, Z))$$

where equality holds if and only if $f(\tilde{X}, z) = T(t)(\mathbb{E}(Y_z|\tilde{X}), z)$ λ -a.e. for $t \in [0, 1]$ as

$$(4.47) \quad L_{y|\tilde{X}}(T(t)) = tL_{y|\tilde{X}}(T(0)) = tV_{y|\tilde{X}},$$

$$(4.48) \quad \frac{1}{\sqrt{2}} D_{y|\tilde{X}}(T(t)) = \frac{1}{\sqrt{2}} (1-t) D_{y|\tilde{X}}(T(0)) = (1-t) V_{y|\tilde{X}}.$$

The above result shows that by fixing $\tilde{X} \in \{\tilde{X} \in \mathcal{D}|_{\mathcal{X}} : \tilde{X} \perp Z\}$, the McCann interpolation between Id and $T_{y|\tilde{X}}$ yields the Pareto frontier from $\mathbb{E}(Y|\tilde{X}, Z)$ to $\overline{\mathbb{E}(Y|\tilde{X}, Z)}$, which is a weak version of the true frontier from $\mathbb{E}(Y|X, Z)$ to $\overline{\mathbb{E}(Y|X, Z)}$. The only difficulty remaining is to coarsen the underlying sigma-algebra from $\sigma(X, Z)$ to $\sigma(\bar{X})$. But by Remark 4.4.1, we know that one can coarsen the sigma-algebra by reducing the variance. Therefore, we apply a diagonal argument to estimate the McCann interpolation between (X, Y) and (\bar{X}, \bar{Y}) .

REMARK 4.5.1 (Diagonal estimate of the post-processing Pareto frontier). *The key observation is that the optimal affine transport map that pushes (X, Y) forward to (X^\dagger, Y^\dagger) is the pair $(T_x, T_{y|\bar{X}})$. Therefore, McCann interpolation between Id and T_x can optimally reduce variance and thereby coarsen $\sigma((X, Z))$ to $\sigma(X^\dagger)$, whereas the interpolation between Id and $T_{y|\bar{X}}$ forms an estimation of the geodesic path between Y and Y^\dagger . Therefore, the present work matches the two interpolations diagonally*

$$(T_x(t), T_{y|\bar{X}}(t)) := ((1-t)Id_x + tT_x, (1-t)Id_y + tT_{y|\bar{X}}),$$

to estimate the true optimal fair data representation at the Pareto frontier.

Finally, since X^\dagger and $\mathbb{E}(Y^\dagger|X^\dagger)$ are the estimation of \bar{X} and $\overline{\mathbb{E}(Y|\bar{X}, Z)}$, respectively, as shown in the last section, it follows from Corollary 4.5.1 and Remark 4.5.1 that

$$(4.49) \quad \mathbb{E}(T_{y|\bar{X}}(t)(Y, Z)|T_x(t)(X, Z)), t \in [0, 1]$$

provides a pre-processing estimate of the Pareto frontier from $\mathbb{E}(Y|X, Z)$ to $\overline{\mathbb{E}(Y|\bar{X}, Z)}$ that is characterized by Theorem 3.3.1.

4.6. Algorithm Design

In this section, we propose two algorithms based on the theoretical results above. Algorithm 2 is designed for the fair learning outcome in the post-processing approach and for the dependent variable in fair data representation, whereas Algorithm 1 is designed for the independent variable in fair data representation.

1. For practitioners who want to generate fair learning outcomes along the Pareto frontier, Algorithm 2 takes the learning outcomes marginals $\{f(X, Z)_z\}_z$ as input and outputs the learning outcomes at (the optimal affine estimation of) the post-processing estimation of the Pareto

frontier: $\{f(X, Z)(t)\}_{t \in [0,1]}$, which is the Wasserstein geodesic paths from the original learning outcome, $f(X, Z)(0)$, to the estimate of the optimal fair learning outcome, $f(X, Z)(1)$. Here, $f(X, Z)(1)$ is the best estimate of the optima fair learning outcome based on the provided learning outcome $\{f(X, Z)_z\}_z$.

2. For practitioners who want to generate a fair data representation, Algorithm 1 and Algorithm 2 take in respectively the marginal independent and dependent data: $\{X_z\}_z$ and $\{Y_z\}_z$, then outputs respectively the independent and dependent data representations along the Wasserstein geodesics from the marginals to their pseudo-barycenter: $\{(X^\dagger(t), Y^\dagger(t))\}_{t \in [0,1]}$. So that any conditional expectation estimation supervised learning model trained via $\{(X^\dagger(t), Y^\dagger(t))\}_{t \in [0,1]}$ results in (an diagonal affine estimation of) the learning outcome at the Pareto frontier.

Algorithm 1 Pseudo-Barycenter Geodesics for Independent Variable

Input: marginal data sets $\{X_z\}_z$, stop criterion ϵ ;

Step 1: Find the optimal barycenter covariance

Initialization: $\delta = \infty$, $\Sigma = rand$ or Id

while $\delta > \epsilon$ **do**

$$\left| \begin{array}{l} \Sigma_{new} = \frac{1}{|X|} \sum_z |X_z| (\Sigma^{\frac{1}{2}} \Sigma_{X_z} \Sigma^{\frac{1}{2}})^{\frac{1}{2}}; \\ \delta = \|\Sigma - \Sigma_{new}\|_F \quad \Sigma = \Sigma_{new} \end{array} \right. \quad // \quad (2.15)$$

end

Step 2: Find the optimal affine transport maps $T_z = \Sigma_{X_z}^{-\frac{1}{2}} (\Sigma_{X_z}^{\frac{1}{2}} \Sigma \Sigma_{X_z}^{\frac{1}{2}})^{\frac{1}{2}} \Sigma_{X_z}^{-\frac{1}{2}}$; // (4.17)

Step 3: Find the geodesic path to independent pseudo-barycenter

$$X_z^\dagger(t) = T_z(t)(X_z - m_{X_z}) + m_X; \quad // \quad (4.16)$$

where $T_z(t) := (1 - t)Id + tT_z$, $t \in [0, 1]$; // (3.22)

Step 4 (optional): For binary rows $X_{i \in I}$, reshape $(X^\dagger(t))_i$ to binary by randomized rounding

$$\text{for all } i \in I \text{ For all } X_i \text{ binary: } p(t) = \frac{(X_z^\dagger(t))_i}{\max((X_z^\dagger(t))_i) - \min((X_z^\dagger(t))_i)}, (X_z^\dagger(t))_i \sim \text{Bernoulli}(p(t))$$

Step 5 (optional): If sensitive information needs to be attached, merge the marginals back with

$$\text{mitigating } Z \quad X_z^\dagger(t) = (X_z(t), z(t)) \text{ where } z(t) = (1 - t)(z - m_Z) + m_Z, t \in [0, 1]$$

Output: $\{\{X_z^\dagger(t)\}_{z \in Z}\}_{t \in [0,1]}$

Algorithm 2 Dependent (or Post-processing) Pseudo-Barycenter Geodesics

Input: marginal data sets $\{Y_z\}_z$ (post-processing: $\{f(X, Z)_z\}_z$), stop criterion ϵ ;

Step 1: Find the optimal barycenter covariance

Initialization: $\delta = \infty$, $\Sigma = \text{rand}$ or Id

while $\delta > \epsilon$ **do**

$$\left| \begin{array}{l} \Sigma_{new} = \frac{1}{|Y|} \sum_z |Y_z| (\Sigma^{\frac{1}{2}} \Sigma_{Y_z|X_z^\dagger} \Sigma^{\frac{1}{2}})^{\frac{1}{2}} \quad // \text{ (2.15)} \\ \text{(post-processing: } \Sigma_{new} = \frac{1}{|Y|} \sum_z |f(X, Z)_z| (\Sigma^{\frac{1}{2}} \Sigma_{f(X, Z)_z} \Sigma^{\frac{1}{2}})^{\frac{1}{2}} \text{)} \quad \delta = \|\Sigma - \Sigma_{new}\|_F \quad \Sigma = \Sigma_{new} \end{array} \right.$$

end

Step 2: Find the optimal affine transport maps $T_z = \Sigma_{Y_z|X_z^\dagger}^{-\frac{1}{2}} (\Sigma_{Y_z|X_z^\dagger}^{\frac{1}{2}} \Sigma \Sigma_{Y_z|X_z^\dagger}^{\frac{1}{2}})^{\frac{1}{2}} \Sigma_{Y_z|X_z^\dagger}^{-\frac{1}{2}}$ // (4.20)

(post-processing: $T_z = \Sigma_{f(X, Z)_z}^{-\frac{1}{2}} (\Sigma_{f(X, Z)_z}^{\frac{1}{2}} \Sigma \Sigma_{f(X, Z)_z}^{\frac{1}{2}})^{\frac{1}{2}} \Sigma_{f(X, Z)_z}^{-\frac{1}{2}}$); // (3.5)

Step 3: Find the geodesic path to dependent pseudo-barycenter $Y_z^\dagger(t) = T_z(t)(Y_z - m_{Y_z}) + m_Y$ // (4.19)

where $T_z(t) := (1 - t)Id + tT_z, t \in [0, 1]$ // (3.22)

(post-processing: $f(X, Z)_z(t) = T_z(t)(f(X, Z)_z - m_{f(X, Z)_z}) + m_{f(X, Z)_z}$); // (3.4)

Step 4 (optional): For binary rows $Y_{i \in I}$ (post-processing: $(f(X, Z))_{i \in I}$), reshape $(Y^\dagger(t))_i$

(post-processing: $(f(X, Z)(t))_{i \in I}$) to binary by randomized rounding for all $i \in I$ For all Y_i

binary: $p(t) = \frac{(Y_z^\dagger(t))_i}{\max((Y_z^\dagger(t))_i) - \min((Y_z^\dagger(t))_i)}$, $(Y_z^\dagger(t))_i \sim \text{Bernoulli}(p(t))$

Output: $\{\{Y_z^\dagger(t)\}_{z \in \mathcal{Z}}\}_{t \in [0, 1]}$ (post-processing: $\{\{f(X, Z)_z(t)\}_{z \in \mathcal{Z}}\}_{t \in [0, 1]}$)

The choice of the Frobenius norm in Step 1 is due to computational efficiency. Any matrix norm would work.

REMARK 4.6.1 (Solution to alternative fair data representation constraint). *In Section ??, the present work shows two alternative fair data representation constraints: (1) $(\tilde{X}, \tilde{Y}) \perp Z$ and (2) $\tilde{X} \perp Z$, which offer different trade-offs between fairness protection and utility. If a practitioner applies the alternative constraint, the proposed algorithms can be applied to generate (the optimal affine estimation of) corresponding fair data representation as the following:*

- 1 For $(\tilde{X}, \tilde{Y}) \perp Z$, one applies Algorithm 1 to both $\{(X_z, Y_z)\}_z$. This alternative is especially useful when practitioners or data publishers do not know which features would be chosen as independent or dependent.
- 2 For $\tilde{X} \perp Z$, one applies Algorithm 1 to $\{X_z\}_z$ and leaves $\{Y_z\}$ untouched.

Now, we are ready to test the proposed methods and algorithms on real-world data.

4.7. Empirical Study: Fair Supervised Learning

In this section, we present numerical experiments with the proposed Algorithms 1 and 2 from Section 4.6. The proposed fair data representation method is bench-marked against two baselines:

1. the prediction model trained via the original data (denoted by “supervised learning name” in the experiment result figure below): supervised learning models trained via data including the sensitive variable provide an estimation of statistical disparity resulting from both disparate treatment and impact.
2. the prediction model trained via data excluding the sensitive variable (denoted by “supervised learning name + Excluding Z”): supervised learning models trained via data excluding the sensitive variable provide an estimation of statistical disparity resulting from only disparate impact.

4.7.1. Benchmark Data and Comparison Methods. For comparison, we implement the following known methods for different types of supervised learning tests:

1. For classification test, the present work compares the current state-of-the-art pre-processing methods [14, 56] (“supervised learning name + Calmon or Zemel”, the later is also known as “Learning Fair Representation”) with the proposed fair data representation methods (“supervised learning name + pre-proc. Pareto frontier Est. or Pseudo-barycenter”).
2. For uni-variate regression test, we compare the post-processing Wasserstein barycenter based fair regression [19] (“supervised learning name + Chzhen”) with the proposed post-processing pseudo-barycenter methods (“supervised learning name + post-proc. Pareto frontier Est. or Pseudo-barycenter”) and the fair data representation methods.
3. For multi-variate supervised learning test, we compare the post-processing pseudo-barycenter methods with the fair data representation methods.

The reasons for this choice are as follows: (1) the known attempts via the pre-processing approach are only available for fair classification; (2) the post-processing Wasserstein barycenter based methods on fair classification are analogous to the one on fair regression, which is shown to outperform other in-processing or post-processing methods in reducing discrimination while preserving accuracy; (3) there exists no practical attempt along the Wasserstein characterization approach to

multi-dimensional supervised learning due to the computational complexity of finding the barycenter and the optimal transport maps.

We adopt the following metrics of accuracy and discrimination that are frequently used in fair machine learning experiments on various data sets: (1) For fair classification, the prediction accuracy, and statistical disparity are quantified respectively by AUC (area under the Receiver Operator Characteristic curve) and

DEFINITION 4.7.1 (Classification discrimination).

$$Discrimination = \max_{z, z' \in \mathcal{Z}} \left| \frac{\mathbb{P}(\hat{Y}_z = 1)}{\mathbb{P}(\hat{Y}_{z'} = 1)} - 1 \right|$$

as defined in [14]. (2) For univariate supervised learning, the prediction error and statistical disparity are quantified respectively by MSE (mean squared error, equivalent to the squared L^2 norm on sample probability space) and KS (Kolmogorov-Smirnov) distance as in [19] for indirect comparison purpose. So that readers can compare the proposed methods indirectly with other methods that are tested in [14, 19, 56] and their references. (3) For univariate and multivariate supervised learning, the prediction error and statistical disparity are quantified respectively by L^2 and \mathcal{W}_2 (Wasserstein) distances, which are the quantification the current work adopts to prove the Pareto frontier in the above sections.

In addition, we perform tests on four benchmark data sets: CRIME, LSAC, Adult, COMPAS, which are also frequently used in fair learning experiments. A brief summary is given below. For all the test results, we apply 5-fold cross-validation with 50% training and 50% testing split, except for 90% training and 10% testing split in the linear regression test on LSAC due to the high computational cost of the post-processing Wasserstein barycenter method [19]. Therefore, interested readers can also compare the pseudo-barycenter test results indirectly to other methods tested in [14, 19].

- Communities and Crime Data Set (CRIME) contains the social, economic, law executive, and judicial data of communities in the United States with 1994 examples [45]. The task of univariate learning is to predict the number of crimes per 10^5 population using the rest of the information on the data set. Here, race is the sensitive information and, for (indirect) comparison purposes,

Data set	Tests	Data size	$\dim(X)$	$\dim(Y)$
UCI Adult	logit regression, random forest	162805	16	1
COMPAS	logit regression, random forest	26390	7	1
LSAC	linear regression, ANN	20454	9	1
CRIME	linear regression, ANN	1994	97	1
CRIME	linear regression, ANN	1994	87	11

we made race a binary categorical variable of whether the percentage of the African American population (`racepctblack`) is greater than 30%.

In multivariate supervised learning on CRIME, we keep the same sensitive variable. But the learning task is to predict the following vector that represents the local housing and rental market information: (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, number of crimes).

- LSAC National Longitudinal Bar Passage Study data set (LSAC) contains social, economic, and personal data of law school students with 20454 examples [52]. The goal of univariate models is to predict the students' GPA using other information on the data set. Here, race is the sensitive variable and, for (indirect) comparison purposes, we make it a binary variable on whether the student is non-white.
- UCI Adult Data Set (Adult) contains the 1994 Census data with 162805 examples [7]. The goal is to predict the binary categorization (whether gross annual income is greater than 50k) using age, education years, and gender, where gender is the sensitive information.
- Correctional Offender Management Profiling for Alternative Sanctions (COMPAS) is a benchmark set of data from Broward County, Florida for algorithmic bias studies [5]. Following [14], the goal here is to predict whether an individual would commit any violent crime while race is the sensitive binary variable (African-American and Caucasian).

4.7.2. Numerical Result. In this subsection, we summarize the experimental results¹.

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

The classification test result is summarized in Figure 4.1 below. Here, the vertical and horizontal axes are AUC and Discrimination defined in Definition 4.7.1. That is, the more upper-left, the better the result. The first row of Figure 4.1 shows the results of logistic regression (left) and random forest (right) on Adult whereas the second shows the corresponding results on COMPAS.

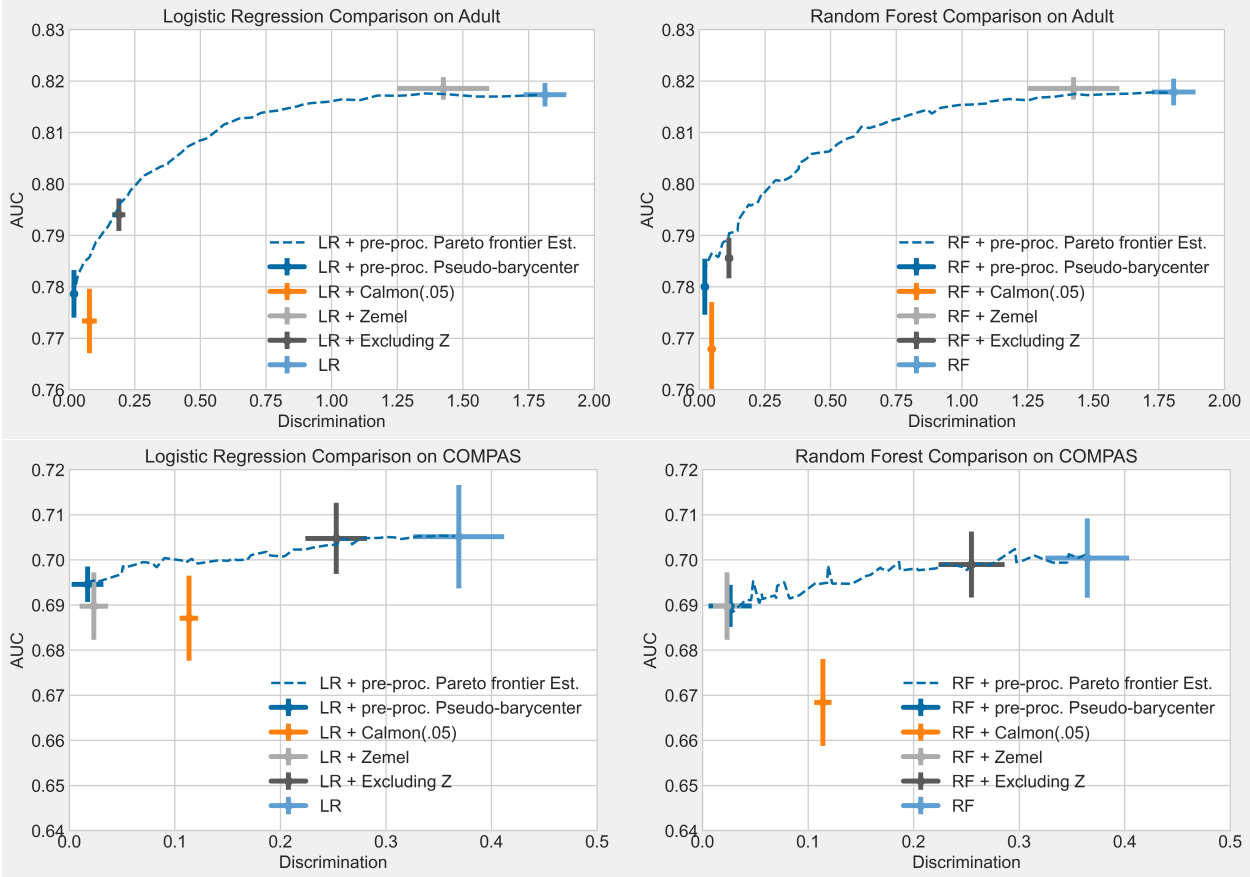


FIGURE 4.1. As shown in the classification test above, the proposed fair data representation method (+ Pre-proc. Pareto frontier Est. or Pseudo-barycenter) outperforms the other methods (+ Zemel or + Calmon) in estimating the optimal fair learning outcome. It reduces the Discrimination metric to nearly zero while keeping the relatively high level of AUC with both logistic regression (LR) and random forest (RF) on both Adult and COMPAS. Furthermore, fair data representation method offers flexibility in choosing the desired trade-off while other methods only estimate a random point near the Pareto frontier.

We note that there exists a large disparate impact in the learning outcome on COMPAS due to the relatively small difference between the “Discrimination” of learning outcome on the original data (LR and RF) and the outcome on the data excluding Z (LR and RF + Excluding Z). Therefore, a further reduction of statistical disparity is needed. In contrast, the relatively large difference in

the Adult data set implies a small disparate impact. That is, a simple exclusion of the sensitive variable Z results in a significant improvement in fairness.

For further reduction of statistical disparity, it is clear from the experiment results on both COMPAS and Adult that the estimation via the Wasserstein geodesics to Pseudo-barycenter (LR and RF + Pseudo-barycenter) consistently outperforms LR and RF + Calmon by obtaining lower Discrimination with higher AUC.

In addition, although “LR and RF + Zemel” achieves a point near the Pareto frontier estimated by the proposed Pseudo-barycenter methods, the point estimation is rather random. Hence, “+ Zemel” is not consistent in estimating the optimal fair learning outcome (the end point of the Pareto curve). Practitioners cannot know which point on the Pareto frontier is estimated by “+ Zemel”. In comparison, the pseudo-barycenter methods are consistent in estimating the optimal fair learning outcome. In addition, they provide the entire Pareto frontier, and hence offer practitioners the flexibility to choose the desired trade-off. Moreover, the proposed method works for any model that aims to estimate conditional expectation, including classification and regression, while “+ Zemel” only works for classification.

The univariate regression test result on the LSAC and the one on CRIME are shown respectively in Figure 4.2 and 4.3 below. Here, the vertical and horizontal axes in the first rows are MSE and KS distance. The corresponding axes in the second row are the L^2 -quantified test error and the \mathcal{W}_2 distance that quantifies the remaining statistical disparity among sensitive groups. Therefore, the more lower-left, the better is the result in both rows. The two supervised learning methods we use are linear regression and artificial neural networks (ANN with 4 linearly stacked layers where each of the first three layers has 32 units all with ReLu activation while the last has 1 unit with linear activation).

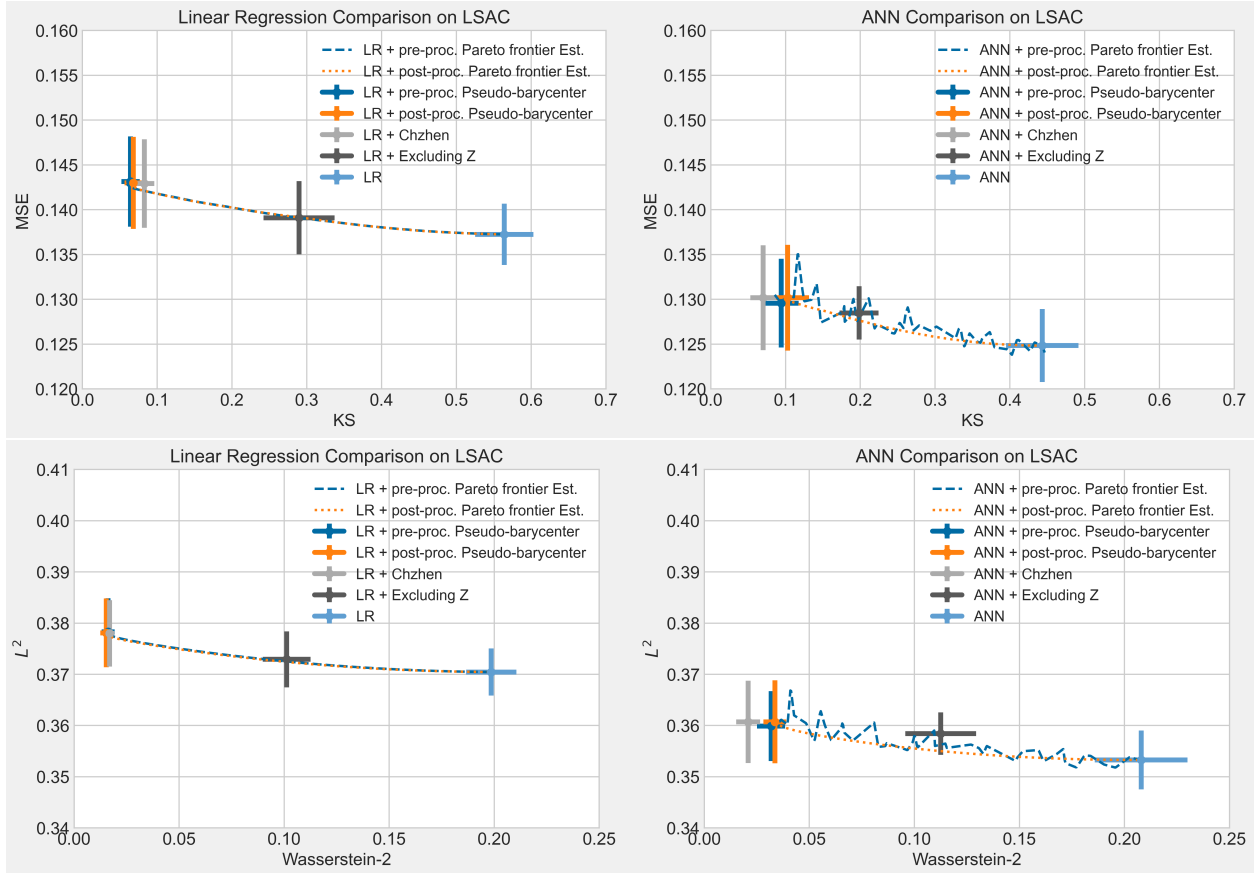


FIGURE 4.2. As shown in the univariate regression test on LSAC above, the proposed fair data representation method (+ pre-proc. Pareto frontier Est. or Pseudo-barycenter) and the post-processing pseudo-barycenter geodesics method (+ post-proc. Pareto frontier Est. or Pseudo-barycenter) achieved similar performance as the exact barycenter method (+ Chzhen). The proposed methods outperformed “+ Chzhen” with linear regression and were exceeded with the artificial neural network, both by a narrow margin. But the performance of the proposed methods is achieved at 0.0128% of the time costs “+ Chzhen” (see Figure 4.5 below). In addition, the proposed methods offer the flexibility of choosing the desired (optimal) trade-off between utility loss (MSE or L^2 -loss) and statistical disparity (KS or \mathcal{W}_2 distance), whereas “+ Chzhen” only estimate the end point of the Pareto curve.

In the regression tests, post-processing Pareto frontier estimation via ANN is smooth while the pre-processing estimation is not. Here, the smoothness is due to the McCann interpolation between the identity matrix and the optimal transport map in the post-processing approach. The non-smoothness is due to the randomness in training the neural network. When testing fair data representations via ANN, one has to train the neural network for the data representation at every

time $t \in [50]$. Hence, the randomness in ANN training results in the non-smoothness in the Pareto frontier estimation via fair data representations.

On the LSAC data set, the proposed methods (+ pre-proc. Pseudo-barycenter and + post-proc. Pseudo-barycenter) obtains a similar performance as the post-processing exact Wasserstein barycenter method (+ Chzhen): the proposed methods outperformed the exact method in the linear regression test and were outperformed by the exact method in the non-linear artificial neural network tests, which is consistent with our theoretical results. But the performance of the proposed methods is achieved at 0.81 seconds on average, whereas the average time cost of “+ Chzhen” is 6365.98 seconds (see Figure 4.5 below). In addition, we gained the flexibility in choosing the desired trade-off, computational efficiency, model selection, parameter tuning, and composition.

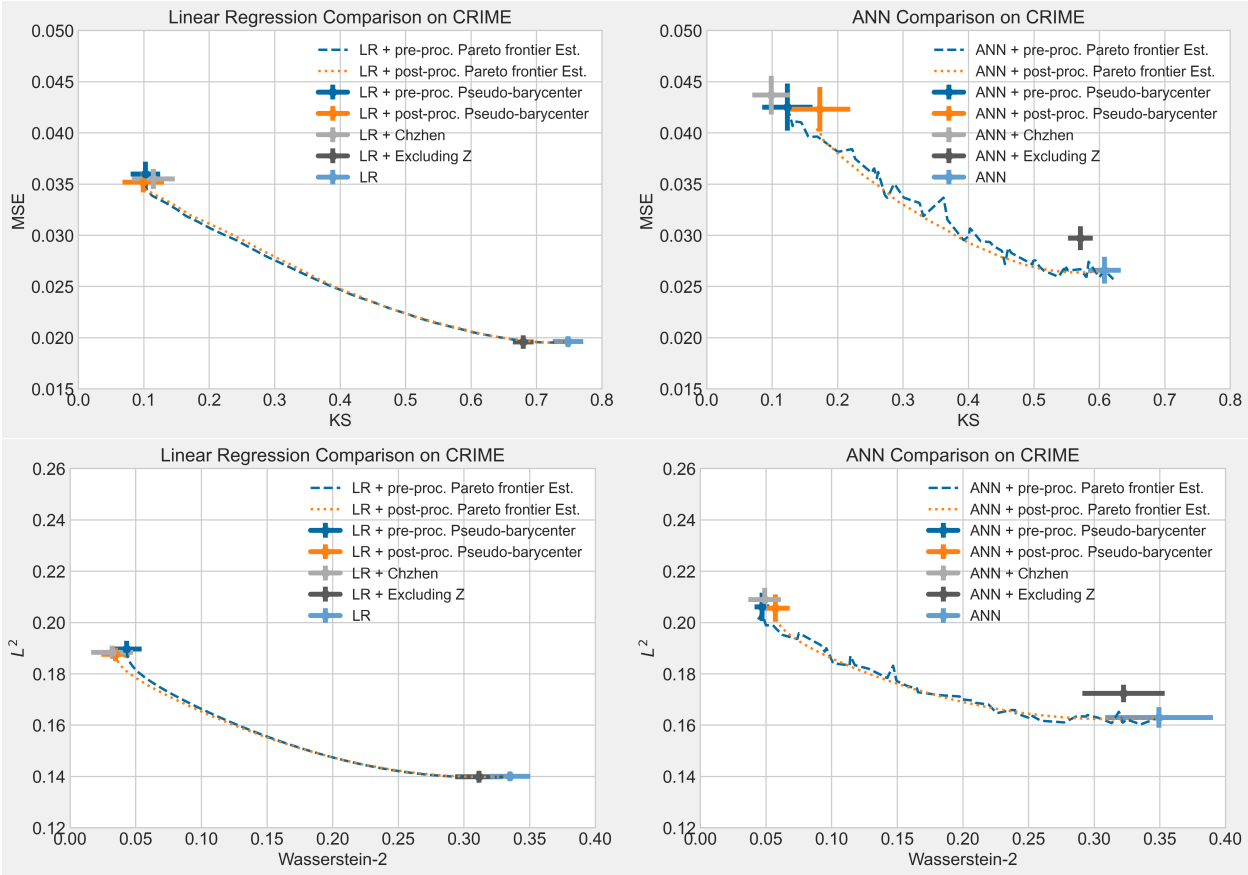


FIGURE 4.3. As shown above, the fair data representation method (+ pre-proc. Pareto frontier Est. or Pseudo-barycenter) achieved the same, if not better, performance as the exact barycenter method (+ Chzhen) in estimating the optimal learning outcome. In addition, the fair data representations method offers flexibility in choosing a desired (optimal) trade-off between utility and fairness.

For CRIME data, the small difference between the KS of learning outcome on the original data (LR and ANN) and the one on the data excluding the sensitive variable (LR and ANN + Excluding Z) implies a significant disparate impact. This observation and the multi-dimensional test below agree with the following statement in [18]: “Simply removing the ‘protected attribute’ is insufficient. As long as the model takes in features that are correlated with, say, gender or race, avoiding explicitly mentioning it will do little good.”

In Figure 4.3, it is clear that the fair data representation methods (+ pre-proc. Pareto frontier Est. or Pseudo-barycenter) achieved the same, if not better, performance as the comparison method (+ Chzhen): the proposed method was outperformed by “+ Chzhen” with linear regression and outperformed “+ Chzhen” with artificial neural network, both by a narrow margin. But the performance of the fair data representation method is achieved at 4.735% of the time costs “+ Chzhen.” In addition, the fair data representation method provides (an estimation of) the entire Pareto frontier and works for multivariate supervised learning (see Figure 4.4 below), whereas “+ Chzhen” only estimates the end point of the Pareto frontier and only works in the univariate learning.

REMARK 4.7.1. *One possible explanation for the proposed method to outperform the exact post-processing Wasserstein barycenter method (“+ Chzhen”) is the following: Although [19] is designed specifically for univariate learning and the KS distance by matching the sensitive marginal cumulative distribution functions, such matching on training data can lead to over-fitting. Therefore, the resulting optimal transport map fits the training data too well to be optimal for the test data.*

Next, we show the multivariate supervised learning on CRIME data to provide a high-dimensional baseline, to which later proposed machine learning fairness methods on high-dimensional data can compare. The vertical and horizontal axes are the L^2 test error and the \mathcal{W}_2 distance among sensitive groups. Hence, the more lower-left, the better the result.

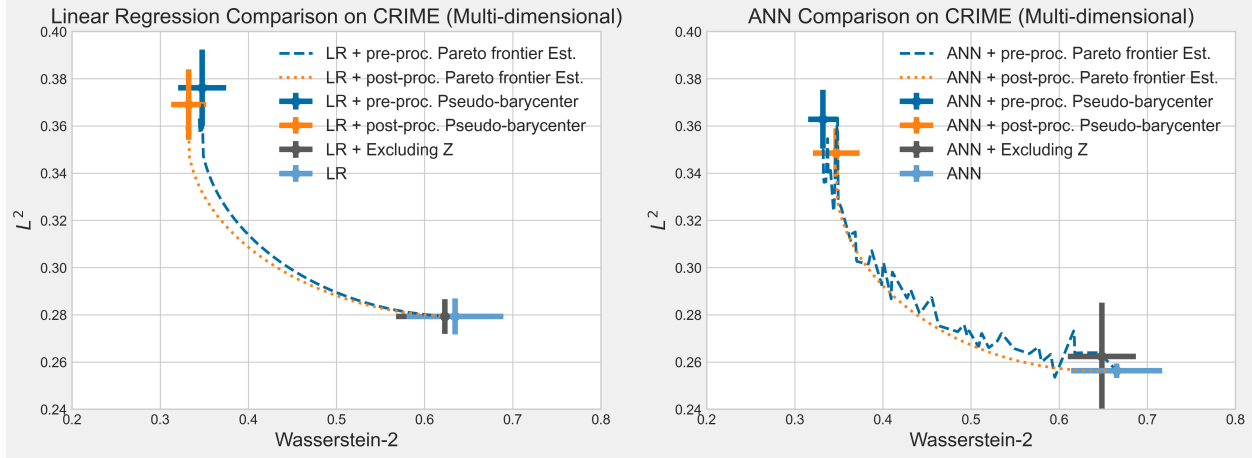


FIGURE 4.4. As shown above, the fair data representation method (+ pre-proc. Pareto frontier Est. or Pseudo-barycenter) achieves similar performance to the post-processing pseudo-barycenter method (+ post-proc. Pareto frontier Est. or Pseudo-barycenter).

Due to the relatively high dimensionality of X (87-dimensional) and Y (11-dimensional), the probabilistic dependence and correlation between the learning outcome and the sensitive variable Z becomes more difficult to remove. It is clear that (LR or ANN + Excluding Z) now removes almost none of the statistical disparity compared to the learning outcome on the original data.

To show the difference in practical computational cost among the comparison methods, we include the following processing time table, where the unit of time is second, and the simulations were run on a 2019 Macbook pro with Intel i9 processor.

Processing Time Table (in seconds)

Test+Method	Data	Train Size	Test Size	Pre-processing	In-processing	Post-processing	Total
LR	CRIME	997	997	0.0	0.0	0.0	0.0
LR+Chzhen	CRIME	997	997	0.0	0.0	78.21	78.21
LR+Pseudo_bary	CRIME	997	997	3.7	0.0	0.0	3.71
ANN	CRIME	997	997	0.0	6.57	0.0	6.57
ANN+Chzhen	CRIME	997	997	0.0	6.57	78.32	84.89
ANN+Pseudo_bary	CRIME	997	997	3.7	6.63	0.0	10.28
LR	LSAC	18408	2046	0.0	0.0	0.0	0.0
LR+Chzhen	LSAC	18408	2046	0.0	0.0	6380.61	6380.61
LR+Pseudo_bary	LSAC	18408	2046	0.81	0.0	0.0	0.82
ANN	LSAC	18408	2046	0.0	105.74	0.0	105.74
ANN+Chzhen	LSAC	18408	2046	0.0	105.74	6351.36	6457.1
ANN+Pseudo_bary	LSAC	18408	2046	0.81	104.2	0.0	106.55

FIGURE 4.5. As shown in the table above, the computational cost of the pseudo-barycenter method is significantly lower than the cost of the known post-processing methods: on average 7836 times faster on LSAC and 21 times faster on CRIME in a single train-test cycle for a single supervised learning model. Furthermore, in model selection or composition, the pre-processing time is a fixed one-time cost while the post-processing time is additive. (See point 4 below for a more detailed explanation)

Now, we show the major advantages of the proposed method compared to the post-processing ones, such as [19, 29, 34]:

- (1) Flexibility in Trade-off: The pre-processing method provides an estimation for the entire Pareto frontier and thereby allows practitioners to balance between prediction error and disparity. In contrast, the known post-processing method merely estimates the starting (left) point of the frontier.
- (2) Sensitive data privacy protection: The geodesics to the pseudo-barycenter allow practitioners to suppress the sensitive information remaining in the data to the desired level. That is, given the resulting suppressed data, anyone who has leaked data from the training or decision stage can merely extract the level of sensitive information up to the pre-determined remaining level. For example, if one chooses to suppress as much sensitive information as possible by setting $t = 1$, then it follows from the construction of dependent and independent pseudobarycenter, it is guaranteed that any unsupervised learning method that uses only the first two moments of the sample data distribution, such as the K-means and PCA, would be unable to extract any information about Z from X^\dagger or $f_{Y^\dagger}(X^\dagger)$.
- (3) Computational efficiency in high-dimensional learning: As summarized in Figure 4.5, the computation of the pseudo-barycenter estimation of the optimal fair learning outcome is significantly

faster than the computation of the exact barycenter via the post-processing matching cdf approach, especially on the LSAC data which has a larger sample size.

- (4) Flexibility in model selection, modification, and composition: in practice, one needs to repeat the training process multiple times to compare different supervised learning algorithms or parameters. The proposed fair data representation method has a fixed pre-processing time while the processing time of post-processing methods is additive. For example, if a practitioner needs to compare linear regression and ANN on LSAC as shown in Figure 4.5 and repeat the training process N times for parameter tuning or validation purpose, the total processing time for pseudo-barycenter method is $0.81 + N(0.0025 + 104.2)$ while the processing time for the post-processing method is $N(0.003 + 6380.61 + 105.738 + 6351.36)$.

(In)Compatibility between Group and Individual Fairness

In this chapter, we study the (in)compatibility between two fairness concepts which look at fairness from different perspectives:

- (*Group fairness*) aims to enforce the (conditional) learning outcome to be equal in distribution or statistics among sensitive groups.
- (*Individual fairness*) aims to guarantee that individuals who share similar qualification data would receive similar learning outcome.

As mentioned in Chapter 1, 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 1.2.1), 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 in 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 [17, 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.

In this chapter, 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 [23], which highlights the potential conflict between these two fairness concepts and proposes methods to achieve both. Subsequent research has delved into this conflict heuristically [10, 17, 21, 35] and experimentally [58]. Despite these efforts, the question of when we can enjoy the best of both remains open [17, Section 3.1]. To our knowledge, only [23, 40] 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 [32, 37]. In this work, our focus is on group fairness defined by statistical parity.

5.1. Generalized Individual Fairness Definitions

For individual fairness, we consider K-Lipschitz Individual Fairness [23] and (ϵ, δ) Individual Fairness [27]. Both definitions share the same heuristics of treating similar individuals 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

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

However, 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 [19, 29, 53], 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:

$$(5.1) \quad f : \mathcal{X} \times \mathcal{Z} \rightarrow \mathcal{Y}.$$

For more general machine learning than the L^2 -objective ones, the logic underlying fairness through awareness [23] 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 5.1.1 (Fairness through awareness). *“Fairness through Awareness” [23] 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 sensitive information and unfairness to diminish the unfairness in learning outcomes. To that end, the authors in [23] proposed a method applying 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 5.1.1 (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 5.1.2 (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 5.1.2 (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 5.1.2 to define the admissible set of functions or maps in the post-processing step to satisfy (ϵ, δ) -IF:

DEFINITION 5.1.3 ((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 5.4.1, 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.

5.2. Problem Setting

In our setting, we put individual fairness regulations as an additional constraint on Problem 1:

$$(5.2) \quad \inf_{f \in L^2(\mathcal{Y} \times \mathcal{Z}, \mathcal{Y})} \{\|\hat{Y} - f(\hat{Y}, Z)\|_2^2 : f(\hat{Y}, Z) \perp Z\}.$$

As in previous chapters, $L^2(\mathcal{Y} \times \mathcal{Z}, \mathcal{Y})$ is the admissible function set consisting of all the square-integrable measurable functions from $\mathcal{Y} \times \mathcal{Z}$ to \mathcal{Y} (See Remark 1.2.1 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 [23] (See Remark 5.1.1 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 $f(\hat{Y}, Z)$.

In the case where the optimal solution to Problem 1, we relax the statistical parity constraint to different tolerance levels of the Wasserstein disparity, $D(\hat{Y}, Z) < d$ for $d \in [0, \infty)$, and put

individual fairness constraints on Problem 2:

$$(5.3) \quad \inf_{f \in L^2(\mathcal{Y} \times \mathcal{Z}, \mathcal{Y})} \{ \|\hat{Y} - f(\hat{Y}, Z)\|_2^2 : D((f(\hat{Y}, Z), Z)) < d \}.$$

As mentioned earlier, Problem 2 characterizes the Pareto frontier between utility loss (quantified by the L^2 norm) and statistical disparity (quantified by the Wasserstein disparity). That is, 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 .

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 5.1.3. Hence, the compatibility between the optimal group fair L^2 learning and individual fairness can be studied by comparing Problem 1 to

PROBLEM 5 ((ϵ, δ)-IF optimal post-processing statistical parity L^2 -objective learning).

$$(5.4) \quad \inf_{f \in \mathcal{D}_{(\epsilon, \delta)\text{-IF}}} \{ \|\hat{Y} - f(\hat{Y}, Z)\|_2^2 : f(\hat{Y}, Z) \perp Z \}.$$

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

PROBLEM 6 ((ϵ, δ)-IF post-processing L^2 -objective learning Pareto frontier).

$$(5.5) \quad \inf_{f \in \mathcal{D}_{(\epsilon, \delta)\text{-IF}}} \{ \|\hat{Y} - f(\hat{Y}, Z)\|_2^2 : D(f(\hat{Y}, Z), Z) < d \}.$$

More specifically, since Problem 1 and 2 can be considered as the respective relaxation of Problem 5 and 6, 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 5 (respectively 6) 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:

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

$$(5.6) \quad \text{Problem 1} \equiv \text{Problem 5.}$$

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

$$(5.7) \quad \text{Problem 2} \equiv \text{Problem 6 at a fixed } d.$$

REMARK 5.2.1 (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 5.6.1 or 5.6.2. In case (2), we can replace \hat{Y} by the qualification or independent variable X and apply the pre-processing composition result in Theorem 5.6.1 or 5.6.2. We refer interested readers to [53] 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:

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

Section 5.2.1 first shows that utility maximization enhances the original statistical parity definition by overcoming the major insufficiency of it as mentioned in [23,31]. 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 sensitive group) similar individuals similarly. See Remark 5.2.2 below for the comparison between relative similarity and absolute similarity. Section 5.1 introduces the two commonly used individual fairness definitions: K -Lipschitz and (ϵ, δ) individual fairness, then generalizes them due to technical reasons. Finally, Section 5.2.2 explains why we choose individual fairness as an additional constraint. Finally, we provide a road map of the main results in the present work.

5.2.1. Explainability: Relative vs. Absolute Similarity. 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 [23,31]: (1) reduced utility, (2) self-fulfilling prophecy, and (3) subset targeting. But, as shown in [53], the first two of the three insufficiencies can be (partially) fixed by maximizing the utility and adopting the Pareto frontier. We refer interested readers to Remark 1.1.1 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 [28] for the technical details on the equivalence between the Wasserstein barycenter problem and the multidimensional Monge-Kantorovich problem.

REMARK 5.2.2 (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 5.2.1 (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.

5.2.2. 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 [26] 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 5.4.1 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 5.4.1 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 5.5.1 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 5.2.1, what is compatibility guarantee for the pre-processed learning result which composes the (Pareto) optimal pre-processing learning steps with the trained model?)

Theorem 5.6.1 and Theorem 5.6.2 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 chapter is structured as follows: Section 5.3 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 5.4 explores the compatibility between individual fairness (K-Lipschitz-IF and (ϵ, δ) -IF) and optimal statistical parity L^2 learning. In Section 5.5, 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 5.7 presents empirical studies validating the theoretical results.

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

To make the chapter self-contained, we summarize the Wasserstein disparity and the existence and uniqueness results of the optimal solutions to Problem 1 and 2 in Chapter 3 and 4, on which our compatibility result is developed.

5.3.1. Quantification of Statistical Disparity. We start by fixing the notations used in the chapter. 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, $\Pi(\mu, \nu) := \{\pi \in \mathcal{P}((\mathbb{R}^d)^2) : \int_{\mathbb{R}^d} d\pi(\cdot, v) = \mu, \int_{\mathbb{R}^d} d\pi(u, \cdot) = \nu\}$. $(\mathcal{P}_2(\mathbb{R}^d), \mathcal{W}_2)$ is 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

$$(5.8) \quad \bar{\mu} := \operatorname{argmin}_{\mu \in \mathcal{P}_2(\mathbb{R}^d)} \left\{ \int_{\mathcal{Z}} \mathcal{W}_2^2(\mu_z, \mu) d\lambda(z) \right\}.$$

Now, in order to relax the hard statistical parity or independence constraint, we apply the average pairwise Wasserstein (\mathcal{W}_2) distance among (the provided predictions of) sensitive groups to quantify statistical disparity and show the desirable properties of the quantification.

DEFINITION 5.3.1 (Wasserstein disparity).

$$(5.9) \quad D(\hat{Y}, Z) := \left(\int_{\mathcal{Z}^2} \mathcal{W}_2^2(\mu_{z_1}, \mu_{z_2}) d\lambda^{\otimes 2}((z_1, z_2)) \right)^{\frac{1}{2}},$$

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:

- (Characterization of statistical parity, Lemma 3.3.1 below)

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

- (*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} .
- (*Characterization of the Pareto frontier*) As shown in [53] or Chapter 3, 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.

5.3.2. Optimal Statistical Disparity L^2 Learning and the Pareto Frontier. Next, we summarize the existence and uniqueness results of the optimal solutions to Problem 1 and 2 that are needed later in the compatibility study.

In Chapter 3, we show that the optimal statistical parity L^2 learning problem has a 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 [12] such that

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

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 5.3.1 (Optimal fair L^2 learning characterization, Chapter 3). *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*

$$(5.10) \quad f^*(\cdot, z) = T_z(\cdot)$$

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 5.3.1 (Generalization to all measurable functions). *The result in Theorem 5.3.1 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,

$$(5.11) \quad V(\hat{Y}, Z) := \min_{f \in L^2(\mathcal{Y} \times \mathcal{Z}, \mathcal{Y})} \{ \|\hat{Y} - f(\hat{Y}, Z)\|_2 : f(\hat{Y}, Z) \perp Z \}.$$

V is an important quantity when studying the Pareto frontier because $t := 1 - \frac{d}{\sqrt{2V}}$ 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 5.3.1 and let

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

denote the McCann interpolation [41] between the identity map (\cdot) 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 5.3.2 (Pareto optimal fair L^2 -objective learning, Chapter 3). *Given (\hat{Y}, Z) satisfying $\mu_z \in \mathcal{P}_{2,ac}$, λ -a.e. and V the independence projection loss defined in (5.11), then*

$$(5.12) \quad f_d(\hat{Y}, Z) := \begin{cases} f^*(1 - \frac{d}{\sqrt{2V}})(\hat{Y}, Z), & \text{if } d \in [0, \sqrt{2V}] \\ \hat{Y}, & \text{if } d \in (\sqrt{2V}, \infty) \end{cases}$$

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.

5.4. 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.

5.4.1. 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 5.4.1 (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 $(f^*(\cdot, z)) = \mathcal{Y}, \lambda - a.e. z \in \mathcal{Z}$ (the optimal fair L^2 learning is capable of making predictions for all individuals, even unobserved),

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

PROOF. (Optimal L^2 learning) It follows from the assumption $\hat{Y} \not\perp Z$ that there exists an event $A \in \mathcal{B}_Z$ 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 5.3.1, 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:

$$(5.13) \quad f(Y_z, z) = \bar{Y}, \text{ for } \lambda - a.e. z \in \mathcal{Z}.$$

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 $y \in \mathcal{Y} = (f(\cdot, z_1)) \cap (f(\cdot, z_2))$ 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 5.3.2, 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 $y \in \mathcal{Y} = (f(\cdot, z_1)) \cap (f(\cdot, z_2))$ 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. \square

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 5.4.1. *Here, we discuss the practicality of the two assumptions in the Theorem 5.4.1*

- 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.

5.4.2. 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 5.4.1 (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*

$$(5.14) \quad \inf_{f \in \mathcal{D}_{(\epsilon, \delta)\text{-IF}}} \{\|\hat{Y} - f(\hat{Y}, Z)\|_2 : f(\hat{Y}, Z) \perp Z\}.$$

The above result shows that, if the solution of Problem 5.3.1 satisfies the assumption $\sup_{(y,z)} \|f^*(y, z) - y\| \leq \frac{\delta - \epsilon}{2}$, then we have

$$(5.15) \quad \text{Problem 1} \equiv \text{Problem 5}$$

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. \square

REMARK 5.4.2 (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,

$$\{f^*(\cdot, z)\}_z \notin \mathcal{D}_{(\epsilon, \delta)\text{-IF}}$$

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

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

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.

5.5. 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 5.3.1, $L(f^*) := \sup_{(y,z)} \|f^*(y, z) - y\|$, V be the independence projection loss defined in equation (5.11), and f_d be the Pareto optimal solutions for $d \in [0, \infty)$ as defined in Lemma 5.3.2. Then we have the following result:

THEOREM 5.5.1 (Compatible portion of Pareto frontier with (ϵ, δ) -IF). *Suppose that*

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

then

$$(5.16) \quad f_d \text{ is } (\epsilon, \delta)\text{-IF for } \begin{cases} d \in [\sqrt{2}V(1 - \frac{\delta - \epsilon}{2L(f^*)}), \infty), & \text{if } \delta - \epsilon \in [0, 2L(f^*)) \\ d \in [0, \infty), & \text{if } \delta - \epsilon \in [2L(f^*), \infty) \end{cases}$$

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

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

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

Here, the last line is due to McCann interpolation: $(1 - \frac{d}{\sqrt{2}V})f(y, z) + \frac{d}{\sqrt{2}V}y = f(1 - \frac{d}{\sqrt{2}V})(y, z)$. Now, by Lemma 5.3.2, $\{f(1 - \frac{d}{\sqrt{2}V})(\cdot, z)\}_z$ are the Pareto optimal solutions at the tolerance level d , it follows from the proof of Lemma 5.4.1 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 5.4.1 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. \square

In other words, Problem 2 \equiv Problem 6 if

$$(5.17) \quad \frac{d}{\sqrt{2V}} \in \begin{cases} [1 - \frac{\delta - \epsilon}{2L(f^*)}, \infty), & \text{if } \delta - \epsilon \in [0, 2L(f^*)) \\ [0, \infty), & \text{if } \delta - \epsilon \in [2L(f^*), \infty) \end{cases}$$

Here, the assumption is exactly the same as assumption 2 in Theorem 5.4.1. The practicality of the assumption is discussed in Remark 5.4.1. 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 $[\sqrt{2V}(1 - \frac{\delta - \epsilon}{2L(f^*)}), \infty)$ (respectively $[0, \infty)$).

REMARK 5.5.1 (Different Cases of (ϵ, δ) in Theorem 5.5.1). *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*

$$(5.18) \quad \sup_{z_1, z_2} \|f(y, z_1) - f(y, z_2)\| \leq \delta < 0,$$

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 $\sqrt{2V}$ 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.

5.6. 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 5.3.1 and Lemma 5.3.2. 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 5.6.1 (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*

$$(5.19) \quad f \circ g \text{ satisfies uniform } (\epsilon, \delta + 2L(f))\text{-IF, for all } \epsilon > 0.$$

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

$$(5.20) \quad g \circ f \text{ satisfies uniform } (\epsilon - 2L(f), \delta)\text{-IF, for all } \epsilon > 2L(f)$$

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, $\forall \epsilon > 2L(f)$. \square

REMARK 5.6.1 (Composition analysis for pre-processing). *As mentioned in Remark 5.2.1, 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 [53]. Therefore, despite our focus on the post-processing case in the present work, Lemma 5.6.1, 5.6.2 and Theorem 5.6.1, 5.6.2 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 5.6.1 (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

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

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. □

That is, if we require the post-processed learning outcome to satisfy uniform (ϵ, δ) -IF, the above result together with Theorem 5.6.1 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

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

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

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

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 5.6.2 (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*

$$(5.21) \quad f \circ g \text{ satisfies uniform } (\epsilon, K\epsilon + 2L(f))\text{-IF, } \forall \epsilon > 0.$$

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

$$(5.22) \quad g \circ f \text{ satisfies uniform } (\epsilon - 2L(f), K\epsilon)\text{-IF}, \forall \epsilon > 2L(f)$$

PROOF. The proof is similar to the proof of Lemma 5.6.1 above and left to the reader. \square

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 5.6.2 (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

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

PROOF. The proof is similar to the proof of Theorem 5.6.1 above and left to the reader. \square

5.7. Empirical Study: Group and Individual Fair Supervised Learning

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

5.7.1. Benchmark Data and Comparison Methods.

- (*Uni-variate regression*) For uni-variate regression test, we implement the compatibility result in Theorem 5.5.1 via the optimal affine estimation of the Wasserstein barycenter that is proposed in [53]. One advantage of applying affine estimation of the optimal transport maps is an accurate derivation of $L(f_{\text{affine}})$ via the fact that affine transport maps can be decomposed into a rigid

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

translation and a linear map:

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

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 $(\sup_z \|f_{\text{linear}}(\cdot, z) - \text{id}\|_{op})(\sup_{y \in (\bar{\mathcal{Y}})} \|y\|)$ 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 [19] (“supervised learning name + chzhen”).

- (*Multi-variate supervised learning*) 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

- (*CRIME dataset*) The CRIME dataset contains social, economic, law enforcement, and judicial data for U.S. communities in 1994 (1994 examples) [45].

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.

- (*LSAC dataset*) The LSAC dataset comprises social, economic, and personal data of 20,454 law school students [52].

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.

5.7.2. Numerical Result. The univariate supervised learning test results for the LSAC and CRIME datasets are presented in Figure 5.1 and Figure 5.2, 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 5.6.2) to provide individual fairness guarantee for the composed learning outcome.

In Figure 5.1, 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 5.6.2 that the Pareto optimal solutions f_d are guaranteed to satisfy the required (ϵ, δ) -IF constraint for all $d \in [\sqrt{2}V[1 - (\frac{\delta}{2L(f^*)} - \epsilon) \wedge 1], \infty)$ where $V = 0.272$ and 0.287 for linear regression and ANN respectively.

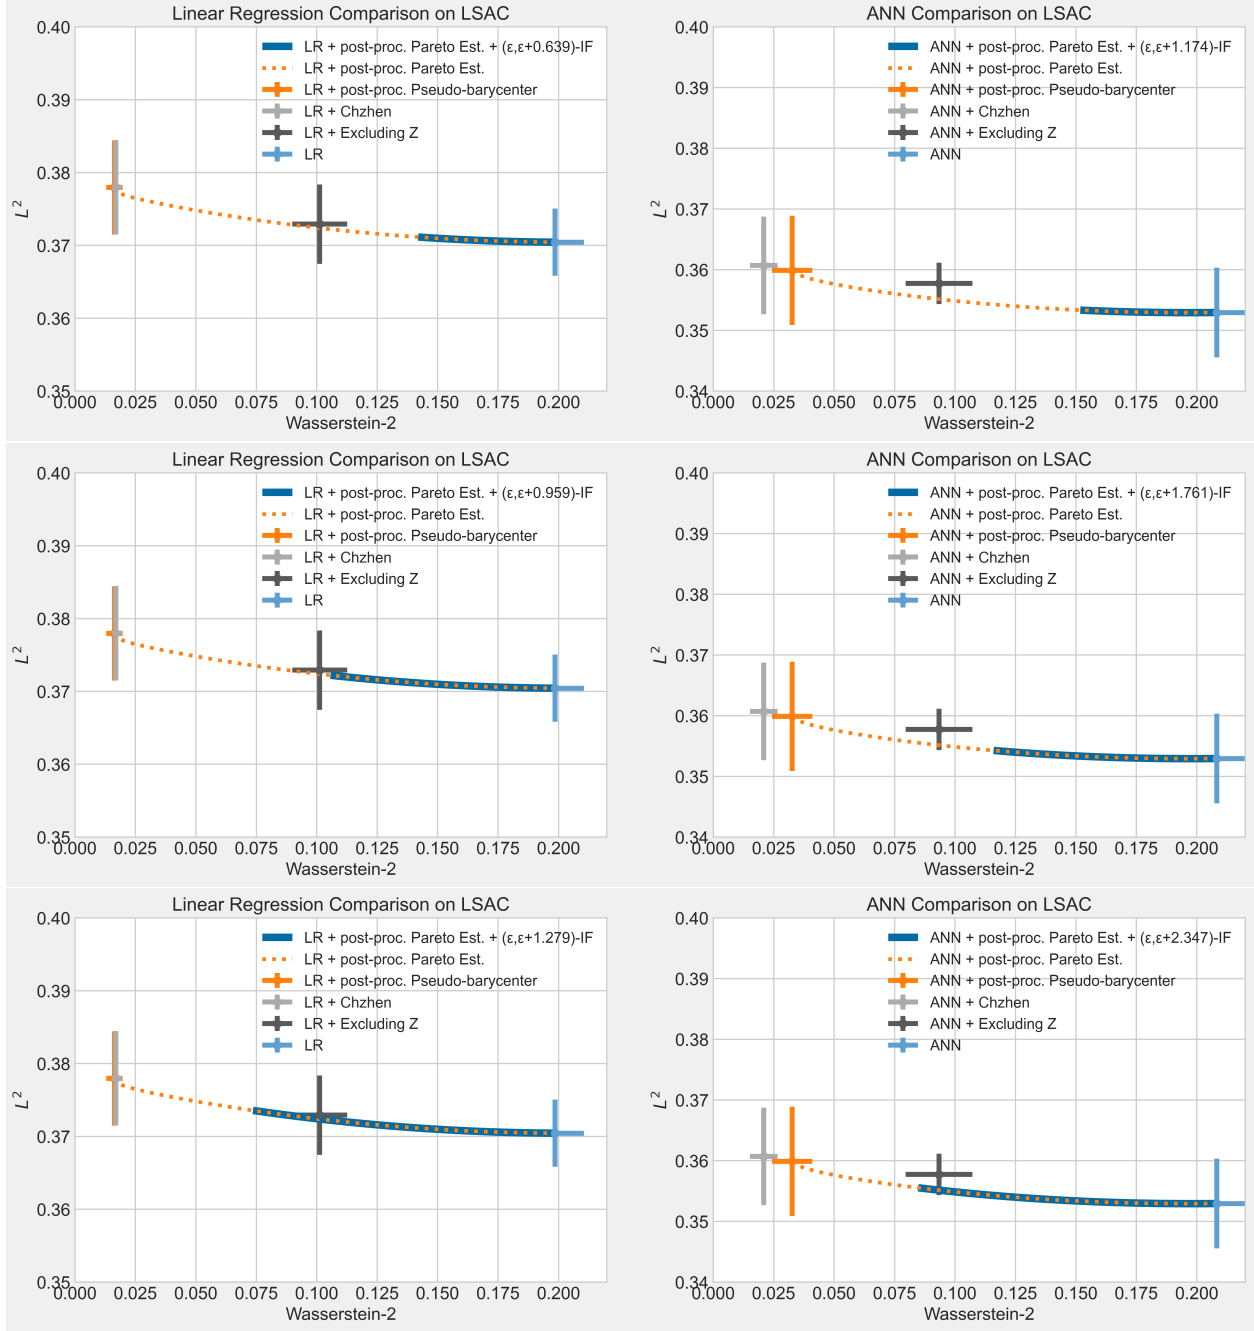


FIGURE 5.1. As shown in the univariate regression test on LSAC above, all three rows consist of the L^2 loss and Wasserstein disparity of the original prediction (LR or ANN), the prediction using data excluding Z (LR or ANN + Excluding Z), the exact post-processing \mathcal{W}_2 barycenter via cumulative 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 5.2, the compatibility result for the CRIME experiment is shown analogously as in Figure 5.1, 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.

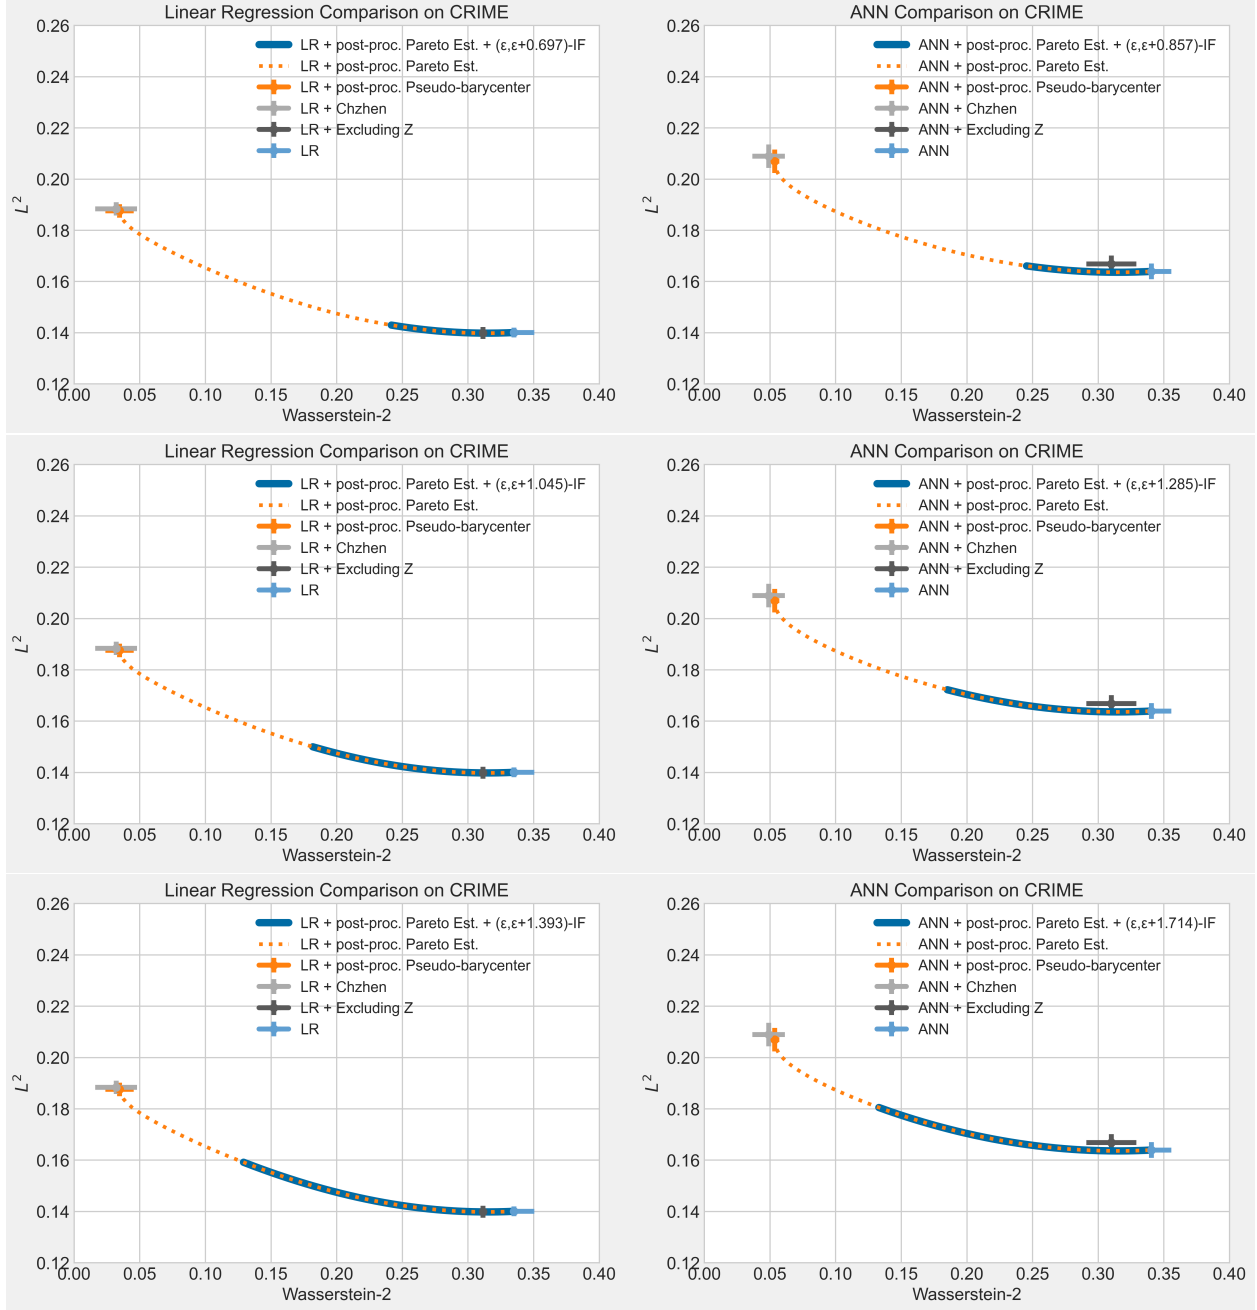


FIGURE 5.2. As shown in the univariate regression test on CRIME above, all three rows consist of the L^2 loss and Wasserstein disparity of the original prediction (LR or ANN), the prediction using data excluding Z (LR or ANN + Excluding Z), the exact post-processing \mathcal{W}_2 barycenter via cdfs 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 5.3. 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.

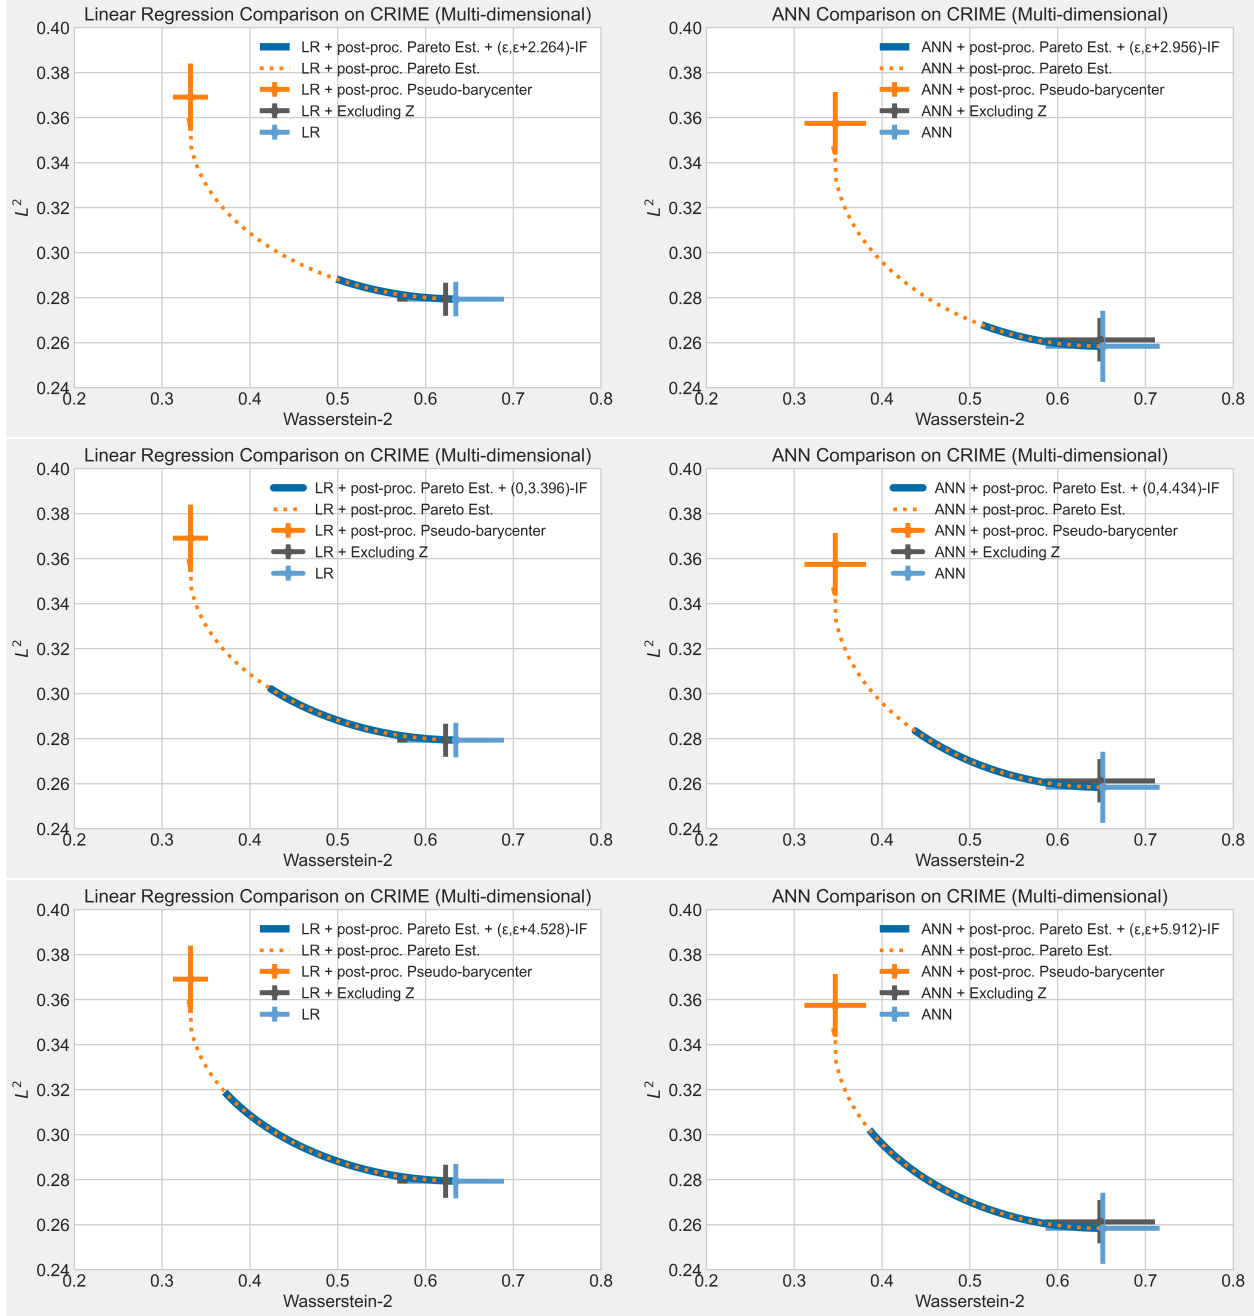


FIGURE 5.3. In the multivariate regression test on CRIME above, all three rows consist of the L^2 loss and Wasserstein disparity of the original prediction (LR or ANN), the prediction using data excluding Z (LR or ANN + Excluding Z), the optimal affine estimation of the post-processing \mathcal{W}_2 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. 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.

CHAPTER 6

Equalized Odds

In this chapter, we target another important group fairness definition: equalized odds. As indicated by its name, the core idea behind equalized odds is that a fair model should provide similar prediction accuracy across all sensitive groups.

6.1. Problem Setting

To start, we define equalized odds in the classic probability language:

DEFINITION 6.1.1 (Equalized odds [31]).

$$(6.1) \quad \hat{Y} \perp Z|Y$$

where \hat{Y} is the prediction generated by the learned model, Z is the sensitive information, and Y is the true label. Therefore, the optimization problem facing us is the following:

PROBLEM 7 (Optimal odds-equalized L^2 -objective learning).

$$(6.2) \quad \inf_{f \in L^2(\mathcal{X} \times \mathcal{Z}, \mathcal{Y})} \{ \|Y - f(X, Z)\|_2^2 : f(X, Z) \perp Z|Y \}$$

By following the same approach as we adopted to find the optimal solution for statistical parity, one might hope to characterize the solution to Problem 7 by applying the conditional (on $Y = y$) Wasserstein barycenter of $\{f(X, Z = z)_y\}_z$. But the conditional barycenter characterization provides a function in

$$\mathcal{X} \times \mathcal{Z} \times \mathcal{Y} \rightarrow \mathcal{Y}.$$

Hence, the characterization is not admissible.

Furthermore, one might also try to circumvent the lack of admissibility by training an unconstrained classifier $\hat{Y} = \hat{Y}(X, Z)$ which is measurable with respect to (X, Z) , then replace Y with \hat{Y} in the characterization. That is, we construct a function $f^*(X, Z) := f(X, Z, \hat{Y})$ by finding the Wasserstein barycenter of $\{f(X, Z = z, \hat{Y} = y)\}_{Z=z}$ for each $y \in \mathcal{Y}$. But we note that this attempt

is a tautology because one needs the true label Y to reduce the disparity in odds discrepancy. In other words, this method requires the perfect solution itself to estimate the solution.

Despite the lack of practicality in solving the post-processing optimal odds-equalized learning (Problem 5), the characterization does provide us practicality via the pre-processing approach. To that end, we first define the optimal odds-equalized data representation problem:

PROBLEM 8 (Optimal odds-equalized data representation for conditional expectation estimation).

$$(6.3) \quad \inf_{X' \in L^2(\mathcal{X} \times \mathcal{Z}, \mathcal{X})} \{ \|Y - \mathbb{E}(Y|X')\|_2^2 : X' \perp Z|Y \}$$

Here, the objective is to maximize the potential utility in X' by minimizing the distance between the dependent variable Y and the best estimator $\mathbb{E}(Y|X')$ in L^2 : $\|Y - \mathbb{E}(Y|X')\|_2^2$. The constraint $X' \perp Z|Y$ guarantees equalized odds by definition.

To make the chapter self-contained, we explain the reason behind the objective and constraint in our problem setting for Problem 8. But the derivation is the same as the optimal fair representation for statistical parity in Chapter 4. Therefore, readers who are familiar with the objective and constraint of Problem 3 can jump to the next section, Section 6.2.

6.1.1. L^2 -objective for Synthetic Data. To start, notice that the L^2 distance between the true Y and the prediction $\hat{Y} = f(X')$ trained via the synthetic data is $\|Y - f(X')\|_2^2$ since learning performance is quantified by L^2 distance in L^2 -objective supervised learning or unsupervised learning. Now, it follows from L^2 orthogonal decomposition that

$$(6.4) \quad \|Y - f(X')\|_2^2 = \|Y - \mathbb{E}(Y|X')\|_2^2 + \|\mathbb{E}(Y|X') - f(X')\|_2^2.$$

$\mathbb{E}(Y|X')$ is the conditional expectation of Y given X' , which is well-known as the L^2 -projection operator in probability or measure theory. Here, the key observation is that, given a fixed synthetic data X' and the true Y , we have the objective for L^2 -supervised learning can be written as the following

$$(6.5) \quad \inf_{f \in \mathcal{F}} \|Y - f(X')\|_2^2$$

\mathcal{F} is the admissible function set that depends on the model choice in the learning task. For example, linear or SVD regression is the usual model choice allowed when building L^2 learning model in finance due to the requirement of interpretability of parameters corresponding to the feature variables. In that case, \mathcal{F} is the family of linear functions parametrized by the linear coefficients.

Therefore, combining equation (6.4) and the training L^2 objective function, we have

$$(6.6) \quad \inf_{f \in \mathcal{F}} \|Y - f(X')\|_2^2 = \underbrace{\|Y - \mathbb{E}(Y|X')\|_2^2}_{\text{data representation utility loss}} + \inf_{f \in \mathcal{F}} \underbrace{\|\mathbb{E}(Y|X') - f(X')\|_2^2}_{\text{training utility loss}}.$$

The first error term only depends on the choice of X' and does not depend on the choice of \mathcal{F} . On the other hand, given any choice of X' , the second error term reduces to zero if the chosen model \mathcal{F} include the conditional expectation: $\mathbb{E}(Y|X') \in \{f(X') : f \in \mathcal{F}\}$. That is, the L^2 decomposition decompose the L^2 learning objective into the following two part:

- (*Training utility loss*) Provided the synthetic data (X', Y) , practitioners would try to minimize

$$\inf_{f \in \mathcal{F}} \|\mathbb{E}(Y|X') - f(X')\|_2^2,$$

where \mathcal{F} denotes the admissible functions that depend on the model choice.

- (*Data representation utility loss*) Since the error term $\|Y - \mathbb{E}(Y|X')\|_2^2$ does not depend on the training process and only depends on the synthetic data (X', Y) , the utility objective of synthetic data design is the following:

$$\inf_{X' \in \mathcal{D}} \|Y - \mathbb{E}(Y|X')\|_2^2,$$

where \mathcal{D} is some admissible set for synthetic data that we define in the next subsection.

6.1.2. Admissible Synthetic Format. It remains to determine the admissible synthetic data form: \mathcal{D} . The present work focuses on $\mathcal{D} = L^2(\mathcal{X} \times \mathcal{Z}, \mathcal{X})$, the all the square integrable measurable maps from $\mathcal{X} \times \mathcal{Z}$ to \mathcal{X} , due to the following considerations:

- (*Measurable*) We require the admissible synthetic data X' to be measurable with respect to the true (X, Z) because the synthetic data should root from the true data: $X' = X'(X, Z)$. More specifically, we keep the sigma-algebra generated by the synthetic data (X', Z) coarser than the

sigma-algebra generated by the original available data (X, Z) :

$$\sigma((X', Z)) \subset \sigma((X, Z)).$$

- (*Square-integrable*) For technical reason, we also require the admissible synthetic data X' to be square integrable such that the composition $f(X') \in L^2(\mathcal{X} \times \mathcal{Z}, \mathcal{Y})$ provided $f \in L^2(\mathcal{X}, \mathcal{Y})$. But our proof shows the solution does not change if one allows all measurable functions.
- (*Product space: $\mathcal{X} \times \mathcal{Z}$*) In order to design the synthetic data that does not contain the sensitive information, one has to access the sensitive variable and have (X, Z) . This is well-known as fairness through awareness.

Finally, putting the equalized odds constraint and the utility objective for synthetic data together, we obtain the optimal odds-equalized synthetic data problem:

$$(6.7) \quad \inf_{X' \in L^2(\mathcal{X} \times \mathcal{Z}, \mathcal{X})} \{ \|Y - \mathbb{E}(Y|X')\|_2^2 : f(X') \perp Z|Y, \forall f \in L^2(\mathcal{X}, \mathcal{Y}) \}.$$

It then follows from the fact

$$(6.8) \quad X' \perp Z|Y \iff f(X') \perp Z|Y \text{ for all } \mathcal{X}/\mathcal{Y}\text{-measurable } f$$

that the problem above is equivalent to Problem 8:

$$(6.9) \quad \inf_{X' \in L^2(\mathcal{X} \times \mathcal{Z}, \mathcal{X})} \{ \|Y - \mathbb{E}(Y|X')\|_2^2 : X' \perp Z|Y \}$$

That completes our problem setting.

6.2. Solution via Conditional Wasserstein Barycenter

Our approach here follows from the Wasserstein barycenter characterization of the optimal fair data representation for statistical parity derived in Chapter 4, which shows that the (almost surely) unique solution to

$$(6.10) \quad \inf_{X' \in L^2(\mathcal{X} \times \mathcal{Z}, \mathcal{X})} \{ \|Y - \mathbb{E}(Y|X')\|_2^2 : X' \perp Z \}$$

can be characterized by the Wasserstein barycenter of the marginal distributions of (X', Z) with respect to $z \in \mathcal{Z}$. We follow the notation in Chapter 4 and denote the solution to (6.10) by \bar{X} to indicate it is the Wasserstein barycenter of $\{X_z\}_{z \in \mathcal{Z}}$. In short, Lemma 4.2.3 shows that the \bar{X}

generates the finest sigma-algebra among all admissible X' satisfying $X' \perp Z$:

$$(6.11) \quad \sigma(\bar{X}) \supset \sigma(X'), \forall X' \in \{X' \in L^2(\mathcal{X} \times \mathcal{Z}, \mathcal{X}) : X' \perp Z\}.$$

It then follows from the fact that

$$(6.12) \quad \sigma(X') \subset \sigma(\bar{X}) \implies \|Y - \mathbb{E}(Y|\bar{X})\|_2^2 \leq \|Y - \mathbb{E}(Y|X)\|_2^2$$

that \bar{X} is the solution to (6.10).

Now, we apply the same technique but now conditioned on $\{Y = y\}$ for almost every $y \in \mathcal{Y}$ to derive a characterization of the solution to Problem 8. To start, we fix notations in the rest of the chapter. Let X_y denote the marginal of X with respect to $Y = y$, $(X_y)_z$ denote the marginal of X_y with respect to $Z = z$, and $\overline{(X_y)}$ the Wasserstein barycenter of $\{(X_y)_z\}_z$, for almost every $y \in \mathcal{Y}$. Also, we denote \bar{X} the random variable satisfying that $\bar{X}_y = \overline{(X_y)}$. Now, we are ready to state our main result, which characterizes the solution to Problem 8 under the assumption that Y is measurable with respect to (X, Z) :

THEOREM 6.2.1 (Characterization of the optimal odds-equalized data representation). *Assume that Y is measurable with respect to (X, Z) and $\{\mathcal{L}((X_y)_z)\}_{z \in \mathcal{Z}} \subset \mathcal{P}_{2,ac}(\mathcal{X})$ for almost every $y \in \mathcal{Y}$, then \bar{X} is admissible: $\bar{X} \in L^2(\mathcal{X} \times \mathcal{Z}, \mathcal{X})$ and $\bar{X} \perp Z|Y$. Furthermore, if set inclusion forms an order between $\{\bar{X}_y\}_y$ and \tilde{X} for all $\tilde{X} \in L^2(\mathcal{X} \times \mathcal{Z}, \mathcal{X})$, then \bar{X} is a solution to Problem 8. If, in addition, the measurable function $f' : \mathcal{X} \times \mathcal{Z} \rightarrow \mathcal{Y}$ satisfying $f'(X, Z) = Y$ is unique, then \bar{X} is the unique solution.*

PROOF. (Admissibility) Here, the proof of measurability in fact helps construct the representation and thereby design an algorithm. To start, by the assumption that Y is measurable with respect to (X, Z) , we have there exists a measurable function

$$f' : \mathcal{X} \times \mathcal{Z} \rightarrow \mathcal{Y}$$

such that $f'(X, Z) = Y$. As a result, we now obtain the data set

$$(X, Y, Z) = (X, f(X, Z), Z).$$

By construction, the map

$$T_1 := (Id|_{\mathcal{X}}, f', Id|_{\mathcal{Z}}) : \mathcal{X} \times \mathcal{Z} \rightarrow \mathcal{X} \times \mathcal{Y} \times \mathcal{Z}$$

is a measurable map. Thereafter, we apply the disintegration theorem once to obtain the conditional data set $\{(X_y, Z_y)\}_{y \in \mathcal{Y}}$ for each value of $y \in \mathcal{Y}$ and apply the theorem twice to obtain the conditional marginals $\{(X_y)_z\}_{z \in \mathcal{Z}}$. Now, for each value of $y \in \mathcal{Y}$, we find the Wasserstein-2 barycenter $\overline{(X_y)}$ via the optimal transport maps $T_y(\cdot, \cdot) : \mathcal{X} \times \mathcal{Z} \rightarrow \mathcal{X}$. By construction, the map

$$T_2 := \{T_y(\cdot, \cdot)\}_y : \mathcal{X} \times \mathcal{Y} \times \mathcal{Z} \rightarrow \mathcal{X} \times \mathcal{Y}$$

is also measurable. Since T_1 and T_2 are both measurable, $T_2 \circ T_1$ is measurable. Hence,

$$\overline{X} := Id|_{\mathcal{X}} \circ T_2 \circ T_1(X, Z)$$

is measurable with respect to (X, Z) . Also, since $\overline{X}_y = \overline{(X_y)} \perp Z$ for $\mathbb{P} \circ Y^{-1}$ -a.e. $y \in \mathcal{Y}$ by construction, \overline{X} satisfies $\overline{X} \perp Z|Y$. Also, it follows from the assumption $\{\mathcal{L}((X_y)_z)\}_{z \in \mathcal{Z}} \subset \mathcal{P}_2(\mathcal{X})$ for almost every $y \in \mathcal{Y}$ that $\|\overline{X}\|_2 < \infty$ which implies $\overline{X} \in L^2(\mathcal{X} \times \mathcal{Z}, \mathcal{X})$.

(Optimality) Now, we prove optimality. To start, it follows from Lemma 4.2.3 and the construction of $\sigma(\overline{X}_y)$ that $\sigma(\overline{X}_y) \supset \sigma(\tilde{X}_y)$ for all $\tilde{X}_y \in \mathcal{D}_y$. Here, the admissible set is defined as

$$\mathcal{D}_y := \{\tilde{X}_y : \tilde{X}_y = T_y(X_y, Z_y) \perp Z_y\}$$

and $T_y(\cdot, \cdot) : \mathcal{X} \times \mathcal{Z} \rightarrow \mathcal{X}$ are Borel measurable maps. Since $\sigma(\tilde{X}_y) \subset \sigma(\overline{X}_y) \implies \text{Var}(\tilde{X}_y) \leq \text{Var}(\overline{X}_y)$ for $\mathbb{P} \circ Y^{-1}$ -a.e. $y \in \mathcal{Y}$, we have

$$\begin{aligned} \text{Var}(\tilde{X}) &= \int_{\mathcal{Y}} \text{Var}(\{\tilde{X}_y\}_y) d\mathbb{P} \circ Y^{-1} \\ &\leq \int_{\mathcal{Y}} \text{Var}(\{\overline{X}_y\}_y) d\mathbb{P} \circ Y^{-1} \\ &= \text{Var}(\overline{X}) \end{aligned}$$

for all $\tilde{X} = Id|_{\mathcal{X}}(\tilde{X}_y \otimes Y)$ with $\tilde{X}_y \in \mathcal{D}_y$. That is, $\text{Var}(\tilde{X}) \leq \text{Var}(\overline{X})$ for all $\tilde{X} \in Id|_{\mathcal{X}}(\mathcal{D}_y \otimes Y) := \{Id|_{\mathcal{X}}(\tilde{X}_y \otimes Y) : \tilde{X}_y \in \mathcal{D}_y\}$. Finally, since for all $\tilde{X} \in L^2(\mathcal{X} \times \mathcal{Z}, \mathcal{X})$, we have $(\tilde{X}, Y) = \tilde{X}_y \otimes Y \in \mathcal{D}_y \otimes Y$. That implies $L^2(\mathcal{X} \times \mathcal{Z}, \mathcal{X}) \subset Id|_{\mathcal{X}}(\mathcal{D}_y \otimes Y)$ and hence $\text{Var}(\tilde{X}) \leq \text{Var}(\overline{X})$ for all

$\tilde{X} \in L^2(\mathcal{X} \times \mathcal{Z}, \mathcal{X})$. If set inclusion forms an order between $\{\bar{X}_y\}_y$ and \tilde{X} for all $\tilde{X} \in L^2(\mathcal{X} \times \mathcal{Z}, \mathcal{X})$, it follows from Remark 4.4.1 that

$$\sigma(\tilde{X}) \subset \sigma(\bar{X})$$

for all $\tilde{X} \in L^2(\mathcal{X} \times \mathcal{Z}, \mathcal{X})$, which further implies that

$$\|Y - \mathbb{E}(Y|\bar{X})\|_2^2 \leq \|Y - \mathbb{E}(Y|\tilde{X})\|_2^2$$

That proves the optimality of \bar{X} .

(Uniqueness) Finally, we assume f' is unique. Then T_1 is unique by construction and the assumption. Also, it follows from the assumption $\{\mathcal{L}((X_y)_z)\}_{z \in \mathcal{Z}} \subset \mathcal{P}_{2,ac}(\mathcal{X})$ that T_2 is unique almost surely. Therefore, the uniqueness of \bar{X} follows from the uniqueness of T_1 and T_2 . That completes the proof. \square

Provided the characterization result, two questions follow naturally:

- (*Applicability of the measurability assumption*) The theoretical result above assumes Y to be measurable with respect to (X, Z) , which is not always true in practice.
- (*Practicality of the construction*) In the proof, it requires the measurable function $f(X, Z) = Y$ to actually construct the exact solution to Problem 8: \bar{X} . But one does not know the explicit form of the measurable function.

To check the measurability assumption and estimate the explicit form of the measurable function for algorithm design purposes, one practical approach is to train the best-supervised learning model with (X, Z) as the independent variable and Y as the dependent variable and then check the estimation error on the training data (or the testing data when assuming training and testing data share the same distribution). If the training error is relatively small compared to the data variance, then the measurability assumption can be considered as satisfied and we can then use the trained model which achieves the relatively small error as a good estimation of the explicit form of the measurable function:

$$\hat{f}(X, Z) = \hat{Y} \approx Y.$$

That is, the triple $(X, \hat{Y}, Z) \approx (X, Y, Z)$ can be considered as a good estimation and therefore the measurability assumption should be considered to be satisfied.

(Application) Once the measurability assumption is considered to be satisfied, one can apply the following steps to generate an estimation of the optimal equalized odds data representation:

- 1 Apply the triple $(X, \hat{Y}, Z) = (X, \hat{f}(X, Z), Z)$ to find an estimation of the conditionals:

$$\{(X_y, Z_y)\}_{y \in \mathcal{Y}} := \{(X|_{\hat{Y}=y}, Z|_{\hat{Y}=y})\}_{y \in \mathcal{Y}}.$$

- 2 Use the obtained conditionals to find the conditional marginals:

$$\{(X_y)_z\}_{z \in \mathcal{Z}} := \{(X_y)|_{Z=z}\}_{z \in \mathcal{Z}}$$

for each value of $y \in \text{range}(\hat{Y})$.

- 3 Apply algorithms that estimate the Wasserstein-2 barycenter, such as Algorithm 1, on the conditional marginals $\{(X_y)_z\}_{z \in \mathcal{Z}}$ to obtain the conditional barycenters:

$$\bar{X}_y := \overline{(X_y)} = \overline{\{(X_y)_z\}_z}$$

for each value of y in the range of \hat{Y} to construct the couple (\bar{X}, \hat{Y}) .

- 4 Output \bar{X} by deleting the \hat{Y} in the couple (\bar{X}, \hat{Y}) :

$$\bar{X} := Id|_{\mathcal{X}}((\bar{X}, \hat{Y})) = Id|_{\mathcal{X}}(\{\bar{X}_y\}_{y \in \text{range}(\hat{Y})}).$$

CHAPTER 7

Future Plan

In this dissertation, we demonstrated how to apply optimal transport to achieve the provable Pareto frontier for fair machine learning, defined by statistical parity, through both post-processing and pre-processing approaches. We also analyzed the incompatibility between the Pareto frontier and individual fairness, and applied conditional optimal transport to deliver odds-equalized machine learning outcomes under certain assumptions.

7.1. Future Plan on Machine Learning Fairness

Despite the promising results, several challenging questions remain unanswered regarding the fairness of machine learning.

- **Optimal Fair Machine Learning in Practice:** How can we achieve theoretically provable optimal fair machine learning characterized by the Wasserstein barycenter in practice? This challenge involves identifying both the barycenter and the optimal transport maps in general marginal distribution cases. In this dissertation, we estimated the optimal solution using affine transport maps. However, relying on affine maps imposes a restrictive constraint, limiting the accuracy of our estimation of general optimal transport maps. Additionally, finding the true Wasserstein barycenter remains a significant challenge in optimal transport research. Although we demonstrated that the optimal fair learning outcome is characterized by the barycenter, further research is needed to develop better methods for estimating general transport maps and the barycenter. This will enhance our ability to achieve more accurate and provable optimal fair learning outcomes.
- **Optimal Odds-Equalized Machine Learning Outcomes in General Cases:** How can we achieve optimal odds-equalized machine learning outcomes in more general cases? In this dissertation, we demonstrated that, under the assumption of measurability of the independent

variable with respect to the dependent and sensitive variables, the optimal odds-equalized learning outcome can be characterized by the conditional Wasserstein barycenter. However, in scenarios where measurability is lacking, characterizing the optimal solution remains a significant challenge.

- **Optimal Fair Learning Characterization for General Objective Functions:** How can we generalize the optimal trade-off characterization to cost functions beyond the quadratic cost? Specifically, for a general cost function in the form of $f(x - y)$ for some bounded measurable function $f : \mathcal{X} \rightarrow \mathbb{R}$, we propose leveraging the fact that polynomials are dense in bounded measurable functions. By applying (1) the characterization of the L^p Wasserstein barycenter of the optimal L^p fair learning and (2) the linearity of cost functions in Kantorovich’s formulation of optimal transport to estimate the optimal fair learning for general cost functions in the form of $f(x - y)$.

7.2. Future Plan in Broader Directions

With the rapid development of AI and machine learning applications, concerns about AI ethics have grown. This was a primary motivation for our study of machine learning fairness. However, AI ethics encompasses much more than fairness. Therefore, we aim to expand our study to include other perspectives within AI ethics.

- **Privacy:** Our approach to machine learning fairness essentially removes sensitive information from datasets. In privacy contexts, sensitive information includes individual identities or personal details. We plan to explore the application of optimal transport to protect (partial or sensitive) individual information.
- **Analysis of Covariance (ANCOVA):** Optimal transport can be used to test or remove statistical dependence between variables, making it a more general version of covariance analysis, which merely captures the linear statistical dependence. We aim to connect current ANCOVA techniques with the fairness techniques we studied and develop systematic methods for analyzing statistical dependence.
- **Robustness:** Robustness addresses the issue of machine learning models being overconfident in their predictions for unseen data. One approach to improve robustness is to generate more data

to help the model discern areas with more or less confident generalization. Adversarial training is a popular method, but it is computationally expensive in high-dimensional data spaces. High-dimensional data often resides on low-dimensional manifolds, suggesting that meaningful adversarial directions are limited. We aim to use optimal transport to create synthetic data, helping models to be appropriately confident about their predictions.

- **Subsampling for Controlled Experiments:** In controlled experiments, it's crucial to partition samples into test and control groups that are similar to avoid confounding variables. By applying unsupervised learning methods to extract potential confounding variables and then using fairness methods to match data, we can form groups that are independent of confounding variables. We aim to combine fairness techniques with unsupervised learning to improve test/control group subsampling processes.
- **Trustworthiness - Conformal Prediction:** Trustworthiness involves providing not just predictions but also confidence intervals around each prediction. Conformal prediction, which estimates conditional distributions, is one approach. We plan to apply optimal transport to estimate conditional distributions for conformal prediction purposes.
- **Generative Model - Science-informed Evolution:** We aim to develop generative models using general optimal transport or stochastic processes to estimate evolution paths between observations. For example, in medical imaging, we might model the progression of a disease. In fairness, we use L^2 Wasserstein interpolation to deform data, but this method is too rigid for modeling the natural evolution. Therefore, we hope to use optimal transport with more general cost functions or the Schrödinger bridge for more flexibility to allow science-informed regularization in evolution path estimation, such as in climate prediction or disease progression modeling.

By exploring these directions, we aim to contribute to a more comprehensive understanding and implementation of AI ethics in machine learning, ensuring that these technologies are developed and applied responsibly.

Bibliography

- [1] B. L. ADAMSON, *Ricci v. DeStefano: Procedural Activism (?)*, National Black Law Journal (University of California, Los Angeles), 24 (2011), pp. 11–01.
- [2] M. AGUEH AND G. CARLIER, *Barycenters in the Wasserstein Space*, SIAM Journal on Mathematical Analysis, 43 (2011), pp. 904–924.
- [3] J. M. ALTSCHULER AND E. BOIX-ADSERA, *Wasserstein Barycenters Are NP-hard to Compute*, SIAM Journal on Mathematics of Data Science, 4 (2022), pp. 179–203.
- [4] P. C. ÁLVAREZ-ESTEBAN, E. DEL BARRIO, J. CUESTA-ALBERTOS, AND C. MATRÁN, *A Fixed-point Approach to Barycenters in Wasserstein Space*, Journal of Mathematical Analysis and Applications, 441 (2016), pp. 744–762.
- [5] J. ANGWIN, J. LARSON, S. MATTU, AND L. KIRCHNER, *Machine Bias*, in Ethics of Data and Analytics, Auerbach Publications, 2022, pp. 254–264.
- [6] J. B. ARISTOTLE ET AL., *The Complete Works of Aristotle*, vol. 2, Princeton University Press Princeton, 1984.
- [7] A. ASUNCION AND D. NEWMAN, *UCI Machine Learning Repository*, 2007.
- [8] 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).
- [9] R. BHATIA, *Positive Definite Matrices*, Princeton University Press, 2009.
- [10] R. BINNS, *On the Apparent Conflict between Individual and Group Fairness*, in Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency, 2020, pp. 514–524.
- [11] A. W. BLUMROSEN, *Strangers in Paradise: Griggs v. Duke Power Co. and the Concept of Employment Discrimination*, Mich. L. Rev., 71 (1972), p. 59.
- [12] Y. BRENIER, *Polar Factorization and Monotone Rearrangement of Vector-Valued Functions*, Communications on Pure and Applied Mathematics, 44 (1991), pp. 375–417.
- [13] T. CALDERS AND I. ŽLIOBAITĖ, *Why Unbiased Computational Processes Can Lead to Discriminative Decision Procedures*, in Discrimination and Privacy in the Information Society: Data Mining and Profiling in Large Databases, Springer, 2013, pp. 43–57.
- [14] 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).
- [15] Y. CAO AND J. YANG, *Towards Making Systems Forget with Machine Unlearning*, in 2015 IEEE Symposium on Security and Privacy, IEEE, 2015, pp. 463–480.
- [16] G. CARLIER AND I. EKELAND, *Matching for Teams*, Economic Theory, 42 (2010), pp. 397–418.

- [17] A. CHOULDECHOVA AND A. ROTH, *The Frontiers of Fairness in Machine Learning*, arXiv preprint arXiv:1810.08810, (2018).
- [18] B. CHRISTIAN, *The Alignment Problem: Machine Learning and Human Values*, WW Norton & Company, 2020.
- [19] E. CHZHEN, C. DENIS, M. HEBIRI, L. ONETO, AND M. PONTIL, *Fair Regression with Wasserstein Barycenters*, Advances in Neural Information Processing Systems, 33 (2020), pp. 7321–7331.
- [20] J. M. COOPER, D. S. HUTCHINSON, ET AL., *Plato: Complete Works*, Hackett Publishing, 1997.
- [21] 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).
- [22] J. A. CUESTA-ALBERTOS, C. MATRÁN-BEA, AND A. TUERO-DIAZ, *On Lower Bounds for the l_2 -Wasserstein Metric in a Hilbert Space*, Journal of Theoretical Probability, 9 (1996), pp. 263–283.
- [23] 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, 2012, pp. 214–226.
- [24] I. EKELAND, *Existence, Uniqueness and Efficiency of Equilibrium in Hedonic Markets with Multidimensional Types*, Economic Theory, 42 (2010), pp. 275–315.
- [25] M. FELDMAN, S. A. FRIEDLER, J. MOELLER, C. SCHEIDEGGER, AND S. VENKATASUBRAMANIAN, *Certifying and Removing Disparate Impact*, in Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 2015, pp. 259–268.
- [26] W. FLEISHER, *What’s Fair about Individual Fairness?*, in Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society, 2021, pp. 480–490.
- [27] 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 (2021), pp. 136–143.
- [28] 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 (1998), pp. 23–45.
- [29] T. L. GOUIC, J.-M. LOUBES, AND P. RIGOLLET, *Projection to Fairness in Statistical Learning*, arXiv preprint arXiv:2005.11720, (2020).
- [30] S. HAJIAN AND J. DOMINGO-FERRER, *A Methodology for Direct and Indirect Discrimination Prevention in Data Mining*, IEEE Transactions on Knowledge and Data Engineering, 25 (2012), pp. 1445–1459.
- [31] M. HARDT, E. PRICE, AND N. SREBRO, *Equality of Opportunity in Supervised Learning*, Advances in Neural Information Processing Systems, 29 (2016).
- [32] U. HÉBERT-JOHNSON, M. KIM, O. REINGOLD, AND G. ROTHBLUM, *Multicalibration: Calibration for the (Computationally-identifiable) Masses*, in International Conference on Machine Learning, PMLR, 2018, pp. 1939–1948.

- [33] L. HU AND Y. CHEN, *A Short-term Intervention for Long-term Fairness in the Labor Market*, in Proceedings of the 2018 World Wide Web Conference, 2018, pp. 1389–1398.
- [34] R. JIANG, A. PACCHIANO, T. STEPLETON, H. JIANG, AND S. CHIAPPA, *Wasserstein Fair Classification*, in Uncertainty in Artificial Intelligence, PMLR, 2020, pp. 862–872.
- [35] 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).
- [36] F. KAMIRAN AND T. CALDERS, *Data Preprocessing Techniques for Classification without Discrimination*, Knowledge and Information Systems, 33 (2012), pp. 1–33.
- [37] 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, PMLR, 2018, pp. 2564–2572.
- [38] Y.-H. KIM AND B. PASS, *Wasserstein Barycenters over Riemannian Manifolds*, Advances in Mathematics, 307 (2017), pp. 640–683.
- [39] T. LE GOUIC AND J.-M. LOUBES, *Existence and Consistency of Wasserstein Barycenters*, Probability Theory and Related Fields, 168 (2017), pp. 901–917.
- [40] 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), IEEE, 2019, pp. 2847–2851.
- [41] R. J. MCCANN, *A Convexity Principle for Interacting Gases*, Advances in Mathematics, 128 (1997), pp. 153–179.
- [42] E. O. OF THE PRESIDENT, *Big Data: Seizing Opportunities, Preserving Values*, President PACT Report, (2014).
- [43] S. PARK, C. YUN, J. LEE, AND J. SHIN, *Minimum Width for Universal Approximation*, arXiv preprint arXiv:2006.08859, (2020).
- [44] B. PASS, *Optimal Transportation with Infinitely Many Marginals*, Journal of Functional Analysis, 264 (2013), pp. 947–963.
- [45] M. REDMOND AND A. BAVEJA, *A Data-driven Software Tool for Enabling Cooperative Information Sharing among Police Departments*, European Journal of Operational Research, 141 (2002), pp. 660–678.
- [46] F. SANTAMBROGIO, *Optimal Transport for Applied Mathematicians*, Birkäuser, NY, 55 (2015), p. 94.
- [47] 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, vol. 34(04), 2020, pp. 3633–3640.
- [48] L. SWEENEY, *Discrimination in Online Ad Delivery: Google Ads, Black Names and White names, Racial Discrimination, and Click Advertising*, Queue, 11 (2013), pp. 10–29.
- [49] E. G. TABAK AND G. TRIGILA, *Explanation of Variability and Removal of Confounding Factors from Data through Optimal Transport*, Communications on Pure and Applied Mathematics, 71 (2018), pp. 163–199.
- [50] C. VILLANI, *Topics in Optimal Transportation*, vol. 58, American Mathematical Soc., 2021.
- [51] C. VILLANI ET AL., *Optimal Transport: Old and New*, vol. 338, Springer, 2009.
- [52] L. F. WIGHTMAN, *LSAC National Longitudinal Bar Passage Study. LSAC Research Report Series.*, (1998).

- [53] S. XU AND T. STROHMER, *Fair Data Representation for Machine Learning at the Pareto Frontier*, Journal of Machine Learning Research, 24 (2023), pp. 1–63.
- [54] S. XU AND T. STROHMER, *On the (In) Compatibility between Group Fairness and Individual Fairness*, SIAM Journal on Mathematics of Data Science (SIMODS), (submitted).
- [55] 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, 2017, pp. 1171–1180.
- [56] R. ZEMEL, Y. WU, K. SWERSKY, T. PITASSI, AND C. DWORK, *Learning Fair Representations*, in International Conference on Machine Learning, PMLR, 2013, pp. 325–333.
- [57] N. ZHOU, Z. ZHANG, V. N. NAIR, H. SINGHAL, J. CHEN, AND A. SUDJANTO, *Bias, Fairness, and Accountability with AI and ML Algorithms*, arXiv preprint arXiv:2105.06558, (2021).
- [58] W. ZHOU, *Group vs. Individual Algorithmic Fairness*, PhD thesis, University of Southampton, 2022.