



ISE

Industrial and
Systems Engineering

The Sharpe Predictor for Fairness in Machine Learning

S. LIU¹ AND L. N. VICENTE¹

¹Lehigh University

ISE Technical Report 21T-019



LEHIGH
UNIVERSITY.

The Sharpe predictor for fairness in machine learning

S. Liu* L. N. Vicente†

August 11, 2021

Abstract

In machine learning (ML) applications, unfair predictions may discriminate against a minority group. Most existing approaches for fair machine learning (FML) treat fairness as a constraint or a penalization term in the optimization of a ML model, which does not lead to the discovery of the complete landscape of the trade-offs among learning accuracy and fairness metrics, and does not integrate fairness in a meaningful way.

Recently, we have introduced a new paradigm for FML based on Stochastic Multi-Objective Optimization (SMOO), where accuracy and fairness metrics stand as conflicting objectives to be optimized simultaneously. The entire trade-offs range is defined as the Pareto front of the SMOO problem, which can then be efficiently computed using stochastic-gradient type algorithms. SMOO also allows defining and computing new meaningful predictors for FML, a novel one being the Sharpe predictor that we introduce and explore in this paper, and which gives the highest ratio of accuracy-to-unfairness. Inspired from SMOO in finance, the Sharpe predictor for FML provides the highest prediction return (accuracy) per unit of prediction risk (unfairness).

1 Introduction

Machine learning algorithms have an increasingly impact on real-life decision-making systems in our society. In many critical applications including credit-scoring, healthcare, hiring, and criminal justice, it is paramount to guarantee that the prediction outcome is both accurate and fair with respect to sensitive attributes such as gender and race. A collection of literature work on algorithm fairness in classification, regression, and clustering problems (see e.g. [4, 5, 7, 11, 13, 17, 20, 22, 36, 37, 38, 39]) focuses on either defining fairness or developing fair learning algorithms. Being of conflicting nature, accuracy and fairness give rise to a set of trade-offs predictors, which can be calculated using stochastic bi-objective optimization methods (penalization [26, 27], epsilon-constrained [37, 38], stochastic bi-gradient [23, 24, 29]). Each of such predictors is non-dominated in the sense that there is no other one simultaneously more accurate and fair. The trade-offs can be typically described by a curve on the accuracy-fair plane, the so-called Pareto front.

*Department of Industrial and Systems Engineering, Lehigh University, 200 West Packer Avenue, Bethlehem, PA 18015-1582, USA (sul217@lehigh.edu).

†Department of Industrial and Systems Engineering, Lehigh University, 200 West Packer Avenue, Bethlehem, PA 18015-1582, USA. Support for this author was partially provided by the Centre for Mathematics of the University of Coimbra under grant FCT/MCTES UIDB/MAT/00324/2020.

In this paper we deal with bi-objective optimization problems of the form $\min F(x)$, where $F : \mathbb{R}^n \rightarrow \mathbb{R}^2$ is a vector function $F = (f_1, f_2)^\top$ defined by two component functions f_1 and f_2 . When f_1 and f_2 are conflicting, no point exists that minimizes both simultaneously. The notion of Pareto dominance allows for a comparison between any two points and leads to the definition of Pareto optimality. One says that x dominates y if $F(x) < F(y)$ componentwise. A point x is then called a Pareto minimizer if it is not dominated by any other one. The so-called Pareto front (or efficient frontier) $F(\mathcal{P})$ is formed by mapping all elements of \mathcal{P} into the decision space \mathbb{R}^2 , $F(\mathcal{P}) = \{F(x) : x \in \mathcal{P}\}$.

Several recent works [23, 24, 25, 27, 37, 38] attempt to capture a complete trade-off between fairness and prediction or clustering quality. In particular, [23, 24, 25] aimed at constructing an entire Pareto front using a posteriori approaches, whereas [27, 37, 38] illustrated the trade-offs with a portion of Pareto front computed by a priori methods where decision-makers' preferences are specified before optimization. The whole Pareto front is not easy to compute especially when the function and gradient evaluations are computationally expensive. Moreover, after obtaining the Pareto front, there is still lack of guideline for the decision-makers to select appropriate solutions. To the best of our knowledge, the existing principles for selecting a certain accuracy-fairness trade-off include the following:

- Given the weights associated with each objective, solving a corresponding weighted-sum problem provides a potential solution candidate.
- Maximizing prediction accuracy [37, 38] subject to the amount of fairness greater than a desired threshold.
- Maximizing the worst performance across sensitive groups [15, 27] leads to lowest performance disparity.

We introduce a new accuracy/fairness trade-off concept inspired from the definition of the Sharpe ratio [32] developed in the sixties for portfolio selection in finance (see, e.g., [14, Section 8.2]). A new predictor, called Sharpe, is defined as the one that provides the highest gain in accuracy return per unit of risk fairness, thus providing a single trade-off with a meaningful property on the Pareto front. The Sharpe predictor depends on the fairness definition used, and each one will lead to a different rate of accuracy or misclassification. The use of the Sharpe predictor will thus allow the comparison of different fairness concepts in fair machine learning. To our knowledge this is the first attempt to identify a prediction data point that could be put in relation to the accuracy/fairness Pareto front, rather than the two front extreme points corresponding to highest accuracy and fairness or any process involving the artificial selection of a weight or a threshold parameter.

2 Fairness in supervised machine learning

In supervised machine learning, training samples consist of pairs of feature vectors (containing a number of features that are descriptive of each instance) and target values/labels. One tries to use the training data to find a predictor, seen as a function mapping feature vectors into target labels, that is accurate on testing data. Such a predictor is typically characterized by a number of parameters, and the process of identifying the optimal parameters is called training or learning. The trained predictor can then be used to predict the labels of test instances. If a

ML predictor inequitably treats samples from different groups defined by *sensitive* or *protected* attributes, such as gender and race, one says that such a predictor is *unfair*. Simply excluding sensitive attributes from the features (called *fairness through unawareness*) does not help as the sensitive attributes (e.g., race) can be inferred from the remaining features (such as the highest educational degree) [30, 39].

Depending on when the fairness criteria are imposed, there are three categories of approaches proposed to handle fairness, namely pre-processing, in-training, and post-processing. Pre-processing approaches [9, 39] transform the input data representation so that the predictions from any model trained on the transformed data are guaranteed to be fair. Post-processing [17, 31] tries to adjust the output of a trained predictor to increase fairness while maintaining the prediction accuracy as much as possible. Even with fairness through unawareness, fairness can be enforced by in-training methods: If the sensitive attributes are accessible in the training samples, in-training methods [3, 6, 18, 38, 36, 37] enforce fairness during the training process either by directly solving fairness-constrained optimization problems, or by adding fairness penalty terms to the learning objective. The approach introduced in this paper falls into the in-training category.

Some of the broadest fairness criteria are *disparate impact* [3] (also called *demographic parity* [6]), *Overall accuracy equality* [5], *equalized odds*, and its special case of *equal opportunity* [17], corresponding to different aspects of fairness. Since many real-world decision-making problems, such as college admission, bank loan application, hiring decisions, etc., can be formulated as binary classification tasks, we will focus in this paper on the binary classification setting. Let $Z \in \mathbb{R}^n$, $A \in \{0, 1\}$, and $Y \in \{-1, +1\}$ respectively denote feature vector, binary-valued sensitive attribute, and target label. Consider a general predictor $\hat{Y} \in \{-1, +1\}$ which depends on Z , A , or both. The predictor is free of disparate impact [3] if the prediction outcome is independent of the sensitive attribute, i.e., for $\hat{y} \in \{-1, +1\}$,

$$\mathbb{P}\{\hat{Y} = \hat{y} | A = 0\} = \mathbb{P}\{\hat{Y} = \hat{y} | A = 1\}. \quad (2.1)$$

Disparate impact may be unrealistic when one group is more likely to be classified as a positive class than the other [19]. In such a case, one can use equalized odds [17] (requiring conditional independence of the sensitive attribute given the true outcome) and equal opportunity [17] (a relaxation of equalized odds imposing such independence only for instances with positive target value). More fairness and bias measures are well-explained in [2, 28, 34].

3 The Sharpe predictor for fairness in machine learning

The Markowitz selection of finance portfolios is done by trading-off return and risk. A portfolio is defined by a vector $x \in X$ denoting the proportions invested in each portfolio asset. Let R denote a vector containing the stochastic returns of all assets. The expected return of the portfolio x is given by $r_1(x) = \mathbb{E}(R)^\top x$. The risk of the portfolio x is given by the variance of its return, $r_2(x) = x^\top Q x$, where Q is the covariance matrix of R . Functions r_1 and r_2 are conflicting. A popular way to determine a portfolio on the Pareto front with desirable properties is by maximizing the so-called Sharpe ratio $(r_1(x) - r_f)/r_2(x)$, where r_f is the expected return of a riskless portfolio ($r_f < \min_{x \in X} r_1(x)$). The Sharpe portfolio gives the highest reward-to-variability ratio or the largest differential of expected return by unit of risk (relatively to a risk-free portfolio). See Figure 1. Sharpe portfolios have other desirable features, in particular being well diversified in their asset composition.




Figure 1: Illustration of Sharpe ratio in portfolio selection.

We are proposing, in the context of fairness in ML, the definition and computation of a *Sharpe predictor*. Prediction accuracy will be called *prediction return*. The lowest the loss or training error $f_1(x)$ is, the highest the return will be. Prediction unfairness will be called *prediction risk*. The highest the fairness measure $f_2(x)$ is, the highest the risk will be. The Sharpe predictor is the one which gives us the highest prediction return per unit of prediction risk (relatively to a risk-free predictor), see Definition 3.1.

Definition 3.1 *The Sharpe predictor x_S is calculated by solving the single-level stochastic optimization problem*

$$\max_x \frac{f_f - f_1(x)}{f_2(x)} = \frac{-f_1(x) - (-f_f)}{f_2(x)}, \quad (3.1)$$

where $-f_1$ is the prediction return (f_1 is the prediction loss), f_2 is the prediction risk, and $f_f > \max_x f_1(x)$ is the prediction loss of a risk-free predictor.

The Sharpe predictor is a Pareto minimizer or non-dominated solution of the bi-objective optimization problem $\min(f_1, f_2)$.

Theorem 3.1 *The Sharpe predictor x_S is a Pareto minimizer of $\min(f_1, f_2)$.*

The proof is elementary and known in the literature Sharpe portfolio selection. If x_S is dominated, there exists a x_d such that $f_1(x_d) < f_1(x_S)$ and $f_2(x_d) < f_2(x_S)$. A contradiction would arise because $(f_f - f_1(x_d))/f_2(x_d) > (f_f - f_1(x_S))/f_2(x_S)$.

Given any predictor x , one has $(f_f - f_1(x_S))/f_2(x_S) > (f_f - f_1(x))/f_2(x)$ (equality can only occur if one of the functions is not strictly convex). Hence, $f_2(x)/f_2(x_S) > (f_f - f_1(x))/(f_f - f_1(x_S))$, which means that if the prediction risk increases from x_S to x , such a change is always larger than the one in prediction return (relative to the return of a risk-free prediction).

4 The Sharpe predictor for accuracy vs disparate impact

Disparate impact (2.1) is one of the most commonly used fairness criteria in supervised machine learning. A general measurement of disparate impact, the so-called CV score [8], is defined by the maximum gap between the probabilities of getting positive outcomes in different sensitive groups, i.e.,

$$\text{CV}(\hat{Y}) = |\mathbb{P}\{\hat{Y} = 1|A = 0\} - \mathbb{P}\{\hat{Y} = 1|A = 1\}|. \quad (4.1)$$

where the binary predictor $\hat{Y} = \hat{Y}(Z; x) \in \{-1, +1\}$ depends on x and maps from Z to Y . As suggested in [24], the trade-offs between prediction accuracy and disparate impact can then be formulated as the following stochastic bi-objective optimization problem:

$$\min (f_1(x), f_2(x)) = \left(\mathbb{E}[\ell(\hat{Y}(Z; x), Y), \text{CV}(\hat{Y}(Z; x))] \right). \quad (4.2)$$

The first objective in (4.2) is a composition function of a loss function $\ell(\cdot, \cdot)$ and the prediction function $\hat{Y}(Z; x)$, and the expectation is taken over the joint distribution of Z and Y .

Let us consider logistic regression model for binary classification, where the parameters are $x = (c, b)^\top$, with $c \in \mathbb{R}^n$ and b a scalar. We assume that N independent samples $\{z_j, a_j, y_j\}_{j=1}^N$ are given. For a given feature vector z_i and the corresponding true label y_i , one searches for a separating hyperplane $\phi(z_j; x) = \phi(z_j; c, b) = c^\top z_j + b$ such that $c^\top z_j + b \geq (<)0$ when $y_j = +1(-1)$, leading to the threshold classifier $\hat{Y}(z_j; c, b) = 2 \times \mathbb{I}(c^\top z_j + b \geq 0) - 1$. The logistic regression loss $\ell(z, y; c, b) = \log(1 + \exp(-y(c^\top z + b)))$ is a smooth and convex surrogate of the 0-1 loss.

The first objective in (4.2) can then be approximated by the empirical logistic regression loss, i.e.,

$$f_1(c, b) = \frac{1}{N} \sum_{j=1}^N \log(1 + \exp(-y_j(c^\top z_j + b))),$$

based on N independent training samples. A regularization term $(\lambda/2)\|c\|^2$ can be added to avoid over-fitting. The second objective is to minimize $\text{CV}(\hat{Y}(Z; x))$. In [38], it was proposed to avoid the non-convexity of the second objective using *decision boundary covariance* as a convex approximate measurement of disparate impact. Specifically, the CV score (4.1) can be approximated by the empirical covariance between the sensitive attributes A and the hyperplane $\phi(Z; c, b)$,

$$\text{Cov}(A, \phi(Z; c, b)) = \mathbb{E}[(A - \bar{A})\phi(Z; c, b)] - \mathbb{E}[A - \bar{A}]\bar{\phi}(Z; c, b) \simeq \frac{1}{N} \sum_{j=1}^N (a_j - \bar{a})\phi(z_j; c, b),$$

where \bar{A} is the expected value of the sensitive attribute, and \bar{a} is an approximated value of \bar{A} using N samples. The intuition behind this approximation is that the disparate impact (2.1) basically requires the predictor to be completely independent of the sensitive attribute. Given that zero covariance is a necessary condition for independence, the second objective can be approximated as:

$$f_2^{\text{DI}}(c, b) = \left[\frac{1}{N} \sum_{j=1}^N (a_j - \bar{a})(c^\top z_j + b) \right]^2,$$

which is monotonically increasing with disparate impact. We construct the finite-sum bi-objective problem

$$\min (f_1(c, b), f_2^{\text{DI}}(c, b)), \quad (4.3)$$

where both objective functions are convex and smooth.

For the illustration of the Sharpe predictor, we will use the experiments reported in [24], where we used the *Adult Income* dataset [21] from the UCI Machine Learning Repository [16]. The cleaned up version of *Adult Income* dataset contains 45,222 samples. Each instance is characterized by 12 nonsensitive attributes (including age, education, marital status, etc.) and a binary sensitive attribute (gender). The goal is to predict whether a person makes over 50K per year. The sensitive attribute chosen was gender. In [24], we have randomly chosen 5,000 training instances, using the remaining instances as the testing set for presenting the results.

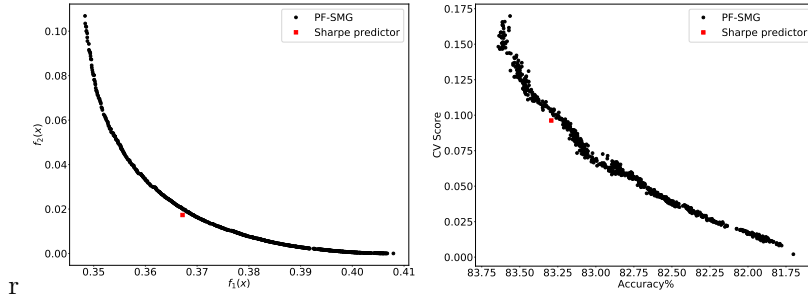


Figure 2: Sharpe predictor and the Pareto front using Adult Income dataset w.r.t. gender. The parameter f_f in (3.1) was set to 0.37. Parameters used in PF-SMG algorithm [23]: step size 0.01 and number of iterations 6,500.

The PF-SMG algorithm [23] was applied to solve (4.3), meaning to compute its Pareto front. The computed Pareto fronts are given in Figure 2, from which we can see the conflicting nature of the two objectives. Figure 2 (left) shows f_1 (prediction error) vs f_2^{DI} (empirical covariance). Figure 2 (right) shows accuracy vs disparate impact (CV score).

The ranges of (f_1, f_2) in ML are not of the same nature of those of (r_1, r_2) in finance, and one difference forces us to consider f_f as the accuracy of a low-risk predictor (instead of a risk-free one). The value of f_f should be selected lower than $\max_x f_1(x)$ but relatively close to it. For example, in order to compute a Sharpe predictor for the bi-objective problem (4.3), we pose the minimization problem $\min (f_1(c, b) - f_f)/f_2^{\text{DI}}(c, b)$, and solve it using the stochastic gradient method.

One can compute a biased stochastic gradient using a batch of samples, as f_1 and f_2^{DI} have finite-sum forms. A gradient estimator g^N is said to be consistent [10] if for any $\epsilon > 0$, $\lim_{N \rightarrow \infty} \mathbb{P}(\|g^N - \nabla f\| > \epsilon) = 0$ holds, where ∇f is the true gradient. Such a consistent gradient estimator can be computed in our implementation by increasing the batch size along the iterations, obtaining increasingly accurate gradients. In our case, where both objectives are convex, an $\mathcal{O}(1/\sqrt{T})$ convergence rate of stochastic gradient descent was proved for consistent but biased gradient estimators (see [10] for more details).

Using the same experiment setting as for the Pareto front calculation, we set $f_f = 0.37$, which is slightly less than the maximum training error (0.394) when using 5,000 training samples. In Figure 2, the testing error, CV score, and testing accuracy using testing samples are evaluated for the obtained Sharpe predictor and marked by a red square. As expected, the Sharpe predictor is on the Pareto front.

5 The Sharpe predictor more broadly

The Sharpe predictor is a broad concept, and it will give researchers and practitioners the possibility to compare and assess a number of different fairness measures proposed for ML. A good candidate for prediction risk is any *subgroup deviation* (see [35, Section 3.3]). In fact, the Sharpe predictor is meaningful form of doing *subgroup aggregation* (see [35, Section 3.4]).

A key advantage of the Sharpe ratio in portfolio selection is that its computation is robust for an observed series of returns without additional information about their distribution. This aspect can be particular interesting in learning from large datasets with high-variance gradients. The

fact that the Sharpe predictor is dimensionless has the potential to introduce novel benchmarking practices when assessing fairness in ML.

The Sharpe ratio is a well-known risk-adjusted performance measure, which enables to evaluate different pairs of return and risk of a portfolio. Another ratio, called information ratio (originally referred to as the appraisal ratio), was introduced in [33]. Instead of taking the risk-free return, the information ratio takes a general benchmark return, in the same vein as we did in the experiment of Section 4. A negative information ratio is an indication of underperformance relative to the selected benchmark. Also, one can choose any measure of risk replacing the standard deviation of the portfolio return in the denominator of these ratios. There are Sharpe type measures that use Value at Risk (VaR) (called *Reward to VaR*) or Conditional Value at Risk (CVaR) (called *Conditional Sharpe Ratio*) [1, 12].

References

- [1] C. Bacon and S. Chairman. How sharp is the sharpe ratio? risk-adjusted performance measures. *Statpro*, nd, 2009.
- [2] S. Barocas, M. Hardt, and A. Narayanan. Fairness in machine learning. *NIPS Tutorial*, 1, 2017.
- [3] S. Barocas and A. D. Selbst. Big data’s disparate impact. *California Law Review*, page 671, 2016.
- [4] 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.
- [5] R. Berk, H. Heidari, S. Jabbari, M. Kearns, and A. Roth. Fairness in criminal justice risk assessments: The state of the art. *Sociological Methods & Research*, 50:3–44, 2021.
- [6] T. Calders, F. Kamiran, and M. Pechenizkiy. Building classifiers with independency constraints. In *2009 IEEE International Conference on Data Mining Workshops*, pages 13–18. IEEE, 2009.
- [7] T. Calders, A. Karim, F. Kamiran, W. Ali, and X. Zhang. Controlling attribute effect in linear regression. In *2013 IEEE 13th international conference on data mining*, pages 71–80. IEEE, 2013.
- [8] T. Calders and S. Verwer. Three naive Bayes approaches for discrimination-free classification. *Data Mining and Knowledge Discovery*, 21:277–292, 2010.
- [9] F. Calmon, D. Wei, B. Vinzamuri, K. N. Ramamurthy, and K. R. Varshney. Optimized pre-processing for discrimination prevention. In *Advances in Neural Information Processing Systems*, pages 3992–4001, 2017.
- [10] J. Chen and R. Luss. Stochastic gradient descent with biased but consistent gradient estimators. *arXiv preprint arXiv:1807.11880*, 2018.
- [11] F. Chierichetti, R. Kuma, S. Lattanzi, and S. Vassilvitskii. Fair clustering through fairlets. In *Advances in Neural Information Processing Systems*, pages 5029–5037, 2017.
- [12] V. Chow and C. W. Lai. Conditional sharpe ratios. *Finance Research Letters*, 12:117–133, 2015.
- [13] S. Corbett-Davies, E. Pierson, A. Feller, S. Goel, and A. Huq. Algorithmic decision making and the cost of fairness. In *Proceedings of the 23rd acm sigkdd international conference on knowledge discovery and data mining*, pages 797–806, 2017.
- [14] G. Cornuejols and R. Tütüncü. *Optimizations Methods in Finance*. Cambridge University Press, Cambridge, 2007.

- [15] E. Diana, W. Gill, M. Kearns, K. Kenthapadi, and A. Roth. Minimax group fairness: Algorithms and experiments. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*, pages 66–76, 2021.
- [16] D. Dua and C. Graff. UCI Machine Learning Repository, 2017.
- [17] M. Hardt, E. Price, and N. Srebro. Equality of opportunity in supervised learning. In *Advances in neural information processing systems*, pages 3315–3323, 2016.
- [18] T. Kamishima, S. Akaho, and J. Sakuma. Fairness-aware learning through regularization approach. In *2011 IEEE 11th International Conference on Data Mining Workshops*, pages 643–650. IEEE, 2011.
- [19] J. Kelly. Women now hold more jobs than men in the u.s. workforce. <https://www.forbes.com/sites/jackkelly/2020/01/13/women-now-hold-more-jobs-than-men>, 2020.
- [20] J. Kleinberg, S. Mullainathan, and M. Raghavan. Inherent trade-offs in the fair determination of risk scores. *arXiv preprint arXiv:1609.05807*, 2016.
- [21] R. Kohavi. Scaling up the accuracy of naive-bayes classifiers: A decision-tree hybrid. In *Proceedings of the Second International Conference on Knowledge Discovery and Data Mining*, KDD96, page 202207. AAAI Press, 1996.
- [22] J. Komiyama, A. Takeda, J. Honda, and H. Shimao. Nonconvex optimization for regression with fairness constraints. In *International conference on machine learning*, pages 2737–2746, 2018.
- [23] S. Liu and L. N. Vicente. The stochastic multi-gradient algorithm for multi-objective optimization and its application to supervised machine learning. Technical report, 2019.
- [24] S. Liu and L. N. Vicente. Accuracy and fairness trade-offs in machine learning: A stochastic multi-objective approach. Technical report, 2020.
- [25] S. Liu and L. N. Vicente. A stochastic alternating balance k -means algorithm for fair clustering. *arXiv preprint arXiv:2105.14172*, 2021.
- [26] N. Martinez, M. Bertran, and G. Sapiro. Fairness with minimal harm: A pareto-optimal approach for healthcare. *arXiv preprint arXiv:1911.06935*, 2019.
- [27] N. Martinez, M. Bertran, and G. Sapiro. Minimax pareto fairness: A multi objective perspective. In *International Conference on Machine Learning*, pages 6755–6764. PMLR, 2020.
- [28] Ninareh Mehrabi, Fred Morstatter, Nripsuta Saxena, Kristina Lerman, and Aram Galstyan. A survey on bias and fairness in machine learning. *ArXiv*, abs/1908.09635, 2019.
- [29] W. R. Monteiro and G. Reynoso-Meza. Proposal of a fair voting classifier using multi-objective optimization.
- [30] D. Pedreshi, S. Ruggieri, and F. Turini. Discrimination-aware data mining. In *Proceedings of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 560–568. ACM, 2008.
- [31] G. Pleiss, M. Raghavan, F. Wu, J. Kleinberg, and K. Q. Weinberger. On fairness and calibration. In *Advances in Neural Information Processing Systems*, pages 5680–5689, 2017.
- [32] W. F. Sharpe. Mutual fund performance. *The Journal of Business*, 39:119–138, 1966.

- [33] J. L. Treynor and F. Black. How to use security analysis to improve portfolio selection. *The Journal of Business*, 46:66–86, 1973.
- [34] S. Verma and J. Rubin. Fairness definitions explained. In *2018 IEEE/ACM International Workshop on Software Fairness (FairWare)*, pages 1–7. IEEE, 2018.
- [35] R. C. Williamson and A. K. Menon. Fairness risk measures. In *International Conference on Machine Learning*, pages 6786–6797, 2019.
- [36] B. Woodworth, S. Gunasekar, M. I. Ohannessian, and N. Srebro. Learning non-discriminatory predictors. In *Conference on Learning Theory*, pages 1920–1953, 2017.
- [37] M. B. Zafar, I. Valera, M. G. Rodriguez, and K. P. Gummadi. Fairness beyond disparate treatment & disparate impact: Learning classification without disparate mistreatment. In *Proceedings of the 26th International Conference on World Wide Web*, pages 1171–1180. International World Wide Web Conferences Steering Committee, 2017.
- [38] M. B. Zafar, I. Valera, M. G. Rodriguez, and K. P. Gummadi. Fairness constraints: Mechanisms for fair classification. In *Artificial Intelligence and Statistics*, pages 962–970, 2017.
- [39] R. Zemel, Y. Wu, K. Swersky, T. Pitassi, and C. Dwork. Learning fair representations. In *International Conference on Machine Learning*, pages 325–333, 2013.