



UNIVERSITY *of* PENNSYLVANIA

Department of Criminology

Working Paper No. 2017-2.0

A Convex Framework for Fair Regression

**Richard Berk
Hoda Heidari
Shahin Jabbari
Matthew Joseph
Michael Kearns
Jamie Morgenstern
Seth Neel
Aaron Roth**

June 9, 2017

This paper can be downloaded from the
Penn Criminology Working Papers Collection:
<http://crim.upenn.edu>

A Convex Framework for Fair Regression

Richard Berk^{1,2}, Hoda Heidari³, Shahin Jabbari³, Matthew Joseph³, Michael Kearns³,
Jamie Morgenstern³, Seth Neel¹, and Aaron Roth³

¹Department of Statistics

²Department of Criminology

³Department of Computer and Information Science
University of Pennsylvania

June 9, 2017

Abstract

We introduce a flexible family of fairness regularizers for (linear and logistic) regression problems. These regularizers all enjoy convexity, permitting fast optimization, and they span the range from notions of group fairness to strong individual fairness. By varying the weight on the fairness regularizer, we can compute the efficient frontier of the accuracy-fairness trade-off on any given dataset, and we measure the severity of this trade-off via a numerical quantity we call the Price of Fairness (PoF). The centerpiece of our results is an extensive comparative study of the PoF across six different datasets in which fairness is a primary consideration.

1 Introduction

The widespread use of machine learning to make consequential decisions about individual citizens (including in domains such as credit, employment, education and criminal sentencing [3, 4, 26, 29]) has been accompanied by increased reports of instances in which the algorithms and models employed can be unfair or discriminatory in a variety of ways [2, 30]. As a result, research on fairness in machine learning and statistics has seen rapid growth in recent years [1, 5–7, 9–11, 13, 14, 18–21, 25, 27], and several mathematical formulations have been proposed as metrics of (un)fairness for a number of different learning frameworks. While much of the attention to date has focused on (binary) classification settings, where standard fairness notions include equal false positive or negative rates across different populations, less attention has been paid to fairness in (linear and logistic) regression settings, where the target and/or predicted values are continuous, and the same value may not occur even twice in the training data.

In this work, we introduce a rich family of fairness metrics for regression models that take the form of a fairness regularizer and apply them to the standard loss functions for linear and logistic regression. Since these loss functions and our fairness regularizer are convex, the combined objective functions obtained from our framework are also convex, and thus permit efficient optimization. Furthermore, our family of fairness metrics covers the spectrum from the type of *group* fairness that is common in classification formulations (where e.g. false arrests in one racial group can be “compensated” for by false arrests in another racial group) to much stronger notions of *individual* fairness (where such cancellations are forbidden, and every injustice is charged to the model). Intermediate fairness notions are also covered. Our framework also permits one to either forbid the use of a “protected” variable (such as race), by demanding that a single model be learned across all groups, or to build different group-dependent models.

Most importantly, by varying the weight on the fairness regularizer, our framework permits us to compute the entire “Pareto curve” or efficient frontier of the trade-off between predictive accuracy and fairness. Such curves are especially important to examine and understand in a domain-specific manner: since demanding fairness of models will always come at a cost of reduced predictive accuracy [8, 11, 15, 34], it behooves practitioners working with fairness-sensitive data sets to understand just how mild or severe this trade-off is in their particular arena, permitting them to make informed modeling and policy decisions.

Our central results take the form of an extensive comparative empirical case study across six distinct datasets in which fairness is a primary concern. For each of these datasets, we compute and examine the corresponding fairness-accuracy efficiency frontier. We introduce an intuitive quantity called the *Price of Fairness (PoF)*, which numerically quantifies the extent to which increased fairness degrades accuracy. We compare the PoF across datasets, fairness notions, and treatments of protected variables.

Our primary contributions are:

- The introduction of a flexible but convex family of fairness regularizers of varying strength that spans the spectrum from group to individual fairness.
- The introduction of a quantitative, data-dependent measure of the severity of the accuracy-fairness tradeoff.
- An extensive empirical comparative study across six fairness-sensitive data sets.

While our empirical study does reveal some reasonably consistent findings across datasets (e.g. efficiency curves show broadly similar shapes; PoF generally higher for individual fairness than group; somewhat surprisingly, PoF not generally improved much when using protected variables), perhaps the most important message is a cautionary one: the detailed trade-off between accuracy and fairness, and the comparison of different fairness notions, appears to be quite domain-dependent and lacking prescriptive “universals”. This is perhaps consistent with the emerging theoretical literature demonstrating the lack of a single “right” definition of fairness [7, 12, 22], and our work adds evidence to the view that fairness is a topic demanding careful domain-specific considerations.

2 The Regression Setting

Consider the standard (linear and logit) regression setting: denote the *explanatory* variables (or *instances*) by $\mathbf{x} \in \mathcal{X} = \mathbb{R}^d$ and the *target* variables (or *labels*) by $y \in \mathcal{Y} = [-1, 1]$. Note that for both linear and logit models, the target values are continuous. Let \mathcal{P} denote the joint distribution over $\mathcal{X} \times \mathcal{Y}$. Suppose every instance \mathbf{x} belongs to exactly one of 2 groups, denoted by 1 and 2.¹ This partition of \mathcal{X} into groups (e.g. into different races or genders) is encoded in a “sensitive” feature \mathcal{X}_{d+1} . Let $S = \{(\mathbf{x}_i, y_i)\}_{i=1}^n$ be a training set of n samples drawn i.i.d. from \mathcal{P} , separated by groups into S_1 and S_2 . Let $n_1 = |S_1|$ and $n_2 = |S_2|$. ($n = n_1 + n_2$.)

This work studies the trade-off between fairness and accuracy for the class of linear and logit regression models. Given a pair of explanatory and target variables (\mathbf{x}, y) , we treat y as the ground truth description of \mathbf{x} ’s merit for the regression task at hand: two pairs $(\mathbf{x}, y), (\mathbf{x}', y')$ with $y \approx y'$ have similar observed outcomes. We aim to design models which treat two such instances with similar observed outcomes similarly, a notion we refer to as *fairness* with respect to the ground truth. For a given accuracy loss ℓ and fairness loss (or *penalty*) f , we define the λ -weighted fairness

¹The generalization to more than 2 groups is straightforward.

loss of a regressor \mathbf{w} on a distribution \mathcal{P} to be $\ell_{\mathcal{P}}(\mathbf{w}) + \lambda f_{\mathcal{P}}(\mathbf{w})$. For our sample S , we analogously define the λ -weighted fairness training loss of \mathbf{w} as $\ell(\mathbf{w}, S) + \lambda f(\mathbf{w}, S)$. For linear regression, we let ℓ be mean-squared error; for logistic regression, we let ℓ be the standard log loss. Finally, we use ℓ_2 regularization for both models, so the overall loss is then $\ell_{\mathcal{P}}(\mathbf{w}) + \lambda f_{\mathcal{P}}(\mathbf{w}) + \gamma \|\mathbf{w}\|_2$.

2.1 A Convex Family of Fairness Regularizers

Our formal definitions of fairness all measure how similarly a model treats two similarly labeled instances, one from group 1 and one from group 2. In particular, all of our definitions have a term for each “cross-group” *pair* of instances/labels, weighted as a function of $|y_i - y_j|$ and also by $|\mathbf{w} \cdot \mathbf{x}_i - \mathbf{w} \cdot \mathbf{x}_j|$. For shorthand, we will refer to pairs of instances (one from each group) as *cross pairs*, and cross pairs with similar labels as similar cross pairs. Each of the below fairness definitions differs in precisely which cross pair disparities can counteract one another. In one extreme (individual fairness in Equation 1), every cross pair disparity increases the fairness penalty of a model. In the other (group fairness in Equation 2), making higher predictions for the group 1 instance of a similar cross pair can be somewhat counterbalanced by making a higher prediction for the group 2 instance of a different similar cross pair. Our notions of fairness for regression align closely to individual and group fairness definitions for classification, both common threads in the fairness literature.

Remark 1. *We assume the sensitive feature \mathcal{X}_{d+1} is available to the learning procedure in one of two ways. In the first setting, which we call the “single model” setting, we assume the algorithm builds a single linear model \mathbf{w} for all of \mathcal{X} (over all but the sensitive features), but can measure the empirical fairness loss of \mathbf{w} using the sensitive feature. In the second setting, which we call the “separate models” setting, we allow the algorithm to build two distinct linear models $\mathbf{w}_1, \mathbf{w}_2$ for the two groups, \mathbf{w}_g based on S_g , thus directly observing the sensitive feature when building these models.*

We specialize the following fairness penalties for the single model setting, but can easily extend them to the separate models setting, by replacing w with w_g when applied to a member of group g .

Individual Fairness The first fairness penalty² we propose is the following:

$$f_1(\mathbf{w}, S) = \frac{1}{n_1 n_2} \sum_{\substack{(\mathbf{x}_i, y_i) \in S_1 \\ (\mathbf{x}_j, y_j) \in S_2}} d(y_i, y_j) (\mathbf{w} \cdot \mathbf{x}_i - \mathbf{w} \cdot \mathbf{x}_j)^2, \quad (1)$$

for some fixed non-negative function d , which we assume is decreasing in $|y_i - y_j|$ (see Section 4 for more details). Since $d(y_i, y_j)$ does not depend upon the decision variables (\mathbf{w}) , one can treat these values as constants in an optimization procedure for selecting \mathbf{w} .

The penalty f_1 corresponds to *individual fairness*; for every cross pair $(\mathbf{x}, y) \in S_1, (\mathbf{x}', y') \in S_2$, a model \mathbf{w} is penalized for how differently it treats \mathbf{x} and \mathbf{x}' (weighted by a function of $|y - y'|$). No cancellation occurs: the penalty for overestimating several of one group’s labels cannot be mitigated by overestimating several of the other group’s labels.

²For simplicity we define our fairness penalties on samples rather than on the underlying distribution \mathcal{P} . The analogous definitions with respect to the distribution can be derived by replacing sums with expectations.

Group Fairness The second fairness penalty we propose is the following:

$$f_2(\mathbf{w}, S) = \left(\frac{1}{n_1 n_2} \sum_{\substack{(\mathbf{x}_i, y_i) \in S_1 \\ (\mathbf{x}_j, y_j) \in S_2}} d(y_i, y_j) (\mathbf{w} \cdot \mathbf{x}_i - \mathbf{w} \cdot \mathbf{x}_j) \right)^2. \quad (2)$$

The penalty f_2 corresponds to *group fairness*: on average, the two groups' instances should have similar labels (weighted by the nearness of the labels of the instances). Unlike f_1 , the penalty f_2 allows for *compensation*: informally, if the model over-values some instances of group 1 relative to group 2 in similar cross pairs, it can compensate on other similar cross-pairs by over-valuing those instances from group 2 relative to group 1.

In both of the above formulations, for any cross pair $(\mathbf{x}_i, y_i) \in S_1$ and $(\mathbf{x}_j, y_j) \in S_2$, any regressor \mathbf{w} will have penalty that increases as $|\mathbf{w} \cdot \mathbf{x}_i - \mathbf{w} \cdot \mathbf{x}_j|$ increases, weighted by $d(y_i, y_j)$. If the cross pair is similar (y_i is close to y_j and $d(y_i, y_j)$ is large), a regressor which makes very different predictions for \mathbf{x}_i and \mathbf{x}_j will incur large loss. If the cross pair is less similar (y_i is far from y_j and $d(y_i, y_j)$ is smaller), there is less penalty for having a regressor for which $|\mathbf{w} \cdot \mathbf{x}_i - \mathbf{w} \cdot \mathbf{x}_j|$ is large.

Hybrid notions of fairness Note that group and individual fairness correspond to two extremes: in one extreme the fairness penalty considers each cross pair separately and in the other one the fairness penalty considers all the cross pairs together. Mathematically one could define different notions of fairness by grouping the cross pairs in different manners or even restrict the fairness penalty only on a subset of cross pairs (called *bucketing*). In particular, for binary labeled data (where $\mathcal{Y} = \{-1, 1\}$) one natural choice is to group the cross pairs based on their labels. This would result in the following definition of fairness which we call hybrid fairness:

$$f_3(\mathbf{w}, S) = \left(\sum_{\substack{(\mathbf{x}_i, y_i) \in S_1 \\ (\mathbf{x}_j, y_j) \in S_2 \\ y_i = y_j = 1}} \frac{d(y_i, y_j) (\mathbf{w} \cdot \mathbf{x}_i - \mathbf{w} \cdot \mathbf{x}_j)}{n_{1,1} n_{2,1}} \right)^2 + \left(\sum_{\substack{(\mathbf{x}_i, y_i) \in S_1 \\ (\mathbf{x}_j, y_j) \in S_2 \\ y_i = y_j = -1}} \frac{d(y_i, y_j) (\mathbf{w} \cdot \mathbf{x}_i - \mathbf{w} \cdot \mathbf{x}_j)}{n_{1,-1} n_{2,-1}} \right)^2, \quad (3)$$

where $n_{g,t}$ denotes the size of group $g \in \{1, 2\}$ with label $t \in \{-1, 1\}$ in the sample. Intuitively, hybrid fairness requires both positive and both negatively labeled cross pairs to be treated similarly in average over the two groups. Some compensation might occur, but only amongst instances with the same label: over-valuing positive instances from group 1 can mitigate over-valuing positive instances from group 2, but does not mitigate under-valuing negative instances from group 1. These two terms can be weighted differently for applications where the treatment of positive (or negative) instances are more important. See Sections 3 and 4 for more details.

2.2 Discussion of Our Notions of Fairness

We now discuss several salient features of our fairness notions.

Why are these fairness notion different from accuracy? All of our fairness penalties are small for any perfect regressor (any \mathbf{w} such that $\mathbf{w} \cdot \mathbf{x} = y$ for all $(\mathbf{x}, y) \sim \mathcal{P}$): for a similar cross pair, $y_i \approx y_j$ and also $\mathbf{w} \cdot \mathbf{x}_i \approx \mathbf{w} \cdot \mathbf{x}_j$ for a perfect regressor \mathbf{w} . Our fairness regularizers might then be interpreted as an unusual proxy for standard accuracy rather than as fairness notions. However,

perfect (linear or otherwise) regressors almost never exist in practice; and between two models with similar accuracy, these definitions bias a learning procedure towards those which have similar treatment of similarly labeled instances from different groups.

What minimizes these penalties? We note that any constant regressor (any \mathbf{w} such that $\mathbf{w} \cdot \mathbf{x} = c$ for all $\mathbf{x} \in \mathcal{X}$ and some $c \in \mathbb{R}$) exactly minimizes all of our fairness regularizers. As we seen empirically, this implies that as the fairness regularization factor λ increases, we transition from an unfair model with minimum accuracy loss to a constant (and therefore perfectly fair but trivial) model, whose accuracy is the best any constant model can achieve.

3 Related Work

Recent work has shown that different fairness notions are often mutually exclusive [7, 12, 22]. Unsurprisingly then, different fairness notions have corresponded to different algorithms and optimization frameworks. Previously introduced fairness notions have generally split along several axis: classification vs. regression and individual vs. group fairness and disparate treatment. Most of previous work has focused on classification, despite the ubiquity of regression in real world applications with fairness concerns.

In classification, one line of work aims to achieve the group fairness notion known as *statistical parity*, i.e. to avoid disparate impact (see e.g. [1, 5, 10, 11, 13, 18–21, 25, 27]). Statistical parity requires a predictor to predict each label at similar rates across different groups. This definition can be at odds with accuracy especially when the two groups are inherently different. Hardt et al. [14] introduced a new notion of group fairness called *equality of odds*, partially to alleviate this friction, and partially arguing that equality of odds more accurately captured what it would mean for a classifier to be equally “good” for two groups. Equality of odds has a very intuitive interpretation for classification: it requires similarity of misclassification rates across groups (rather than forcing the marginal classification rates in the two groups to coincide). Optimizing for accuracy subject to an equality of odds constraint was recently shown to be NP-hard [31]; work following this result presented efficient heuristics for the problem [31, 34]. We also study fairness definitions which we can implement efficiently but our interest is in studying the trade-off between fairness and accuracy and the relationship between different notions of fairness.

Although equality of odds is also defined for regression, it is very difficult to determine empirically whether a regressor’s output is conditionally independent of the protected attribute (conditioned on the true label), as each true label may be seen only once.

Calders et al. [6] introduced the study of statistical parity’s analog in regression settings, (called *equal means* and *balanced residuals*). More recently Johnson et al. [16] also studied fairness for regression problems and formalized several notions for *impartial estimates* based on the causal relationship between *sensitive attributes*, *legitimate attributes*, *suspect attributes* and the label. Both groups consider group fairness; our group fairness notion differs from these as we incorporate the similarity of pairs (through the function d) in our definition though the specific choice of $d(y, y') = c$ for some $c \in \mathbb{R}$ and all y and y' would recover equal means.

To achieve any of these fairness notions, one needs to decide whether or not to allow for *disparate treatment* (allowing for different treatment, or different models, for different groups)³, and where in the learning process to enforce fairness: preprocessing of data (e.g. [18]); inprocessing, during the

³Any classifier that uses sensitive attributes in its decision making is implicitly fitting separate models to the two populations, and while this might seem unfair, it has been argued that it is actually necessary for fairness (most notably in Dwork et al. [9]).

training of a model as either a constraint or incorporated into the objective function (e.g. [11, 21, 33]); or postprocessing, where data is labeled by some black-box model and then relabeled as a function only of the original labels (e.g. [14]). Our approach in this paper falls into the in-processing category, by encoding fairness as a regularizer (an approach previously studied in e.g. [21, 33, 34]). Our work differs from previous work in several aspects by primarily focusing on regression, and that our family of fairness measures draws inspiration from the idea that *similar* instances should be treated *similarly* [9, 35].

4 A Comparative Empirical Case Study

In this section we describe an empirical case study in which we apply our regularization framework to six different datasets in which fairness is a central concern. These datasets include cases in which the observed labels are real-valued, and cases in which they are binary-valued. For the real-valued datasets, we apply standard linear regression with our various fairness regularizers. For the binary-valued datasets, we apply standard logistic regression, again along with fairness regularizers. For datasets with real-valued targets we normalized the inputs and outputs to be zero mean and unit variance, and we set the cross-group fairness weights as $d(y_i, y_j) = e^{-(y_i - y_j)^2}$; for datasets with binary targets we set $d(y_i, y_j) = \mathbb{1}[y_i = y_j]$.

For each dataset S , our framework requires that we solve optimization problems of the form $\min_{\mathbf{w}} \ell(\mathbf{w}, S) + \lambda f(\mathbf{w}, S) + \gamma \|\mathbf{w}\|_2$ for variable values of λ , where $\ell(\mathbf{w}, S)$ is either MSE (linear regression) or the logistic regression loss. For each λ we picked γ as a function of this λ by cross validation (see Appendix A.1 for more details). All optimization problems are solved using the CVX solver in Matlab (for real-valued datasets) or python (for binary-valued datasets).⁴ Furthermore, all the results are reported using 10-fold cross validation (see more details in Appendix A.2).

The datasets themselves are summarized in Table 1, where we specify the size and dimensionality of each, along with the “protected” feature (race or gender) that thus defines the subgroups across which we apply our fairness criteria (see Appendix A.3 for more details). The datasets vary considerably in the number of observations, their dimensionality, and the relative size of the minority subgroup.

The *Adult* dataset [23, 24] from the UC Irvine Repository contains 1994 Census data, and the goal is to predict whether the income of an individual in the dataset is more than 50K per year or not. The sensitive or protected attribute is gender.⁵ The *Communities and Crime* dataset [24] includes features relevant to per capita violent crime rates in different communities in the United States, and the goal is to predict this crime rate; race is the protected variable. The *COMPAS* dataset contains data from Broward County, Florida, originally compiled by ProPublica [2], in which the goal is to predict whether a convicted individual would commit a violent crime in the following two years or not. The protected attribute is race, and the data was filtered in a fashion similar to that of Corbett-Davies et al. [8]. The *Default* dataset [24, 32] contains data from Taiwanese credit card users, and the goal is to predict whether an individual will default on payments. The protected attribute is gender. The *Law School* dataset⁶ consists of the records of law students who went on to take the bar exam. The goal is to predict whether a student will pass the exam based on features such as LSAT score and undergraduate GPA. The protected attribute is gender. The *Sentencing* dataset contains information from a state department of corrections regarding inmates in 2010. The

⁴See <http://www.cvxr.com> and <http://www.cvxpy.org> for more details. We set the number of iterations to be 1000 in our optimization solvers.

⁵We only used the data in Adult.data in our experiments.

⁶<http://www2.law.ucla.edu/sander/Systemic/Data.htm>

goal is to predict the sentence length given by the judge based on factors such as previous criminal records and the crimes for which the conviction was obtained. The protected attribute is gender.

Data Set	Type	n	d	Minority n	Protected
Adult	logit	32561	14	10771	gender
Communities and Crime	linear	1994	128	227	race
COMPAS	logit	3373	19	1455	race
Default	logit	30000	24	11888	gender
Law School	logit	27478	36	12079	gender
Sentencing	linear	5969	17	385	gender

Table 1: Summary of datasets. Type indicates whether regression is logistic or linear; n is total number of data points; d is dimensionality; Minority n is the number of data points in the smaller population; Protected indicates which feature is protected or fairness-sensitive.

4.1 Accuracy-Fairness Efficient Frontiers

We begin by examining the efficient frontier of accuracy vs. fairness for the six datasets. These curves are shown in Figure 1, and are obtained by varying the weight λ on the fairness regularizer, and for each value of λ finding the model which minimizes the associated regularized loss function. For the logistic regression cases, we extract probabilities from the learned model \mathbf{w} as $\Pr[y_i = 1] = \exp(\mathbf{w} \cdot x_i) / (1 + \exp(\mathbf{w} \cdot x_i))$ and evaluate these probabilities as predictions for the binary labels using MSE.⁷ In all of the datasets, as λ increases, the models converge to the best constant predictor, which minimizes the fairness penalties.

Perhaps the most striking aspect of Figure 1 is the great diversity of tradeoffs across different datasets and different fairness regularizers. For instance, if we examine the individual fairness regularizer, on four of the datasets (Adult, Communities and Crime, Law School and Sentencing), the curvature is relatively mild and constant — there is an approximately fixed rate at which fairness can be traded for accuracy. In contrast, on COMPAS and Default, fairness loss can be reduced almost for “free” until some small threshold value, at which point the accuracy cost increases dramatically. Similar comments can be made regarding hybrid fairness in the logistic regression cases.

Individual fairness appears to be strictly more costly than group fairness for the entire regime between the extremes of $\lambda = 0$ and $\lambda \rightarrow \infty$ for a majority of these datasets (with the exception of COMPAS and Default datasets). In COMPAS and Default, for small amounts of unfairness, group unfairness may be more costly than the same level of individual unfairness.

Perhaps surprisingly, building separate models for each population barely improves the tradeoff (and in some cases, hurts the tradeoff for some values of λ) across almost all datasets and fairness regularizers. This suggests that the academic discussion about whether to allow disparate treatment (as explicitly allowed in e.g. [9, 17]) — i.e. whether sensitive attributes such as race and gender should be “forbidden” vs. used in building more accurate models for each subpopulation, is perhaps less consequential than expected (at least on these datasets and using linear or logistic regression models).

Note that in theory group fairness is strictly less costly than individual fairness for any particular model \mathbf{w} (by Jensen’s inequality), and using separate models (one for each group) should strictly

⁷Note that assessing the MSE of these probabilities, interpreted as predictions, is a sensible choice. Since squared error is a proper scoring rule, if the labels are indeed generated according to a logistic regression model, minimizing the squared error of a predictor using mean squared error will elicit the true model as its minimizer.

improve the fairness/accuracy trade-off for any of these notions of fairness. However, both group fairness and separate models are more prone to overfitting (than individual fairness and single model), and hence a larger ℓ_2 -regularization parameter γ tends to be selected in cross validation in these settings (see Appendix A.1 for more details). This is a surprising interaction between the strength of the fairness penalty and the generalization ability of the model, and results in group fairness sometimes having a more severe tradeoff with accuracy when compared to individual fairness, and separate models having little benefit out of sample, although they can appear to have a large effect in-sample (because of the effects of over-fitting).

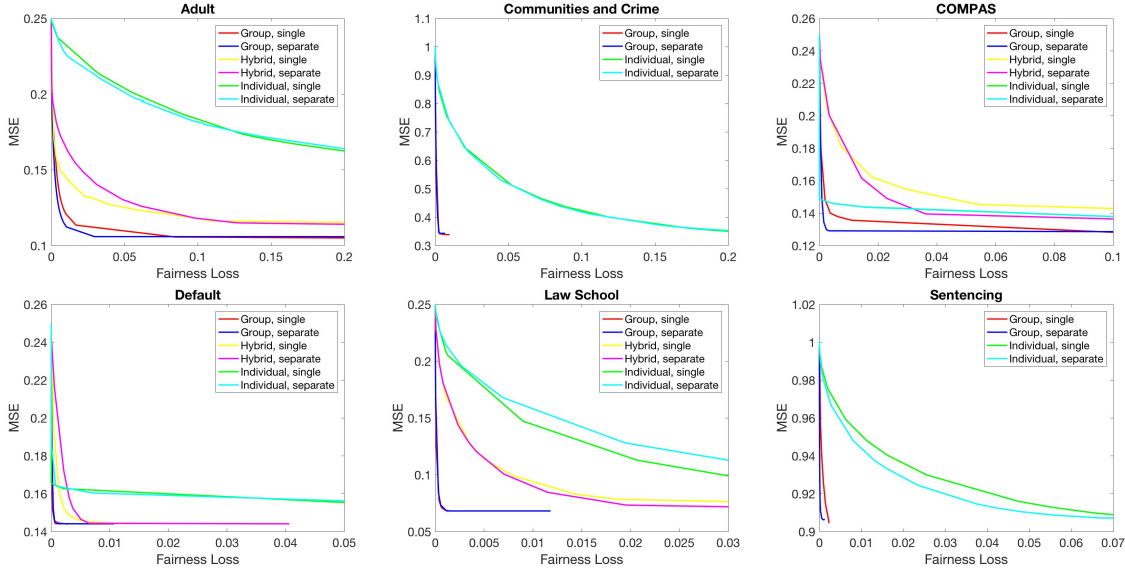


Figure 1: Efficient frontiers of accuracy vs. fairness for each dataset. For datasets with binary-valued targets (logistic regression), we consider three fairness notions (group, individual and hybrid), and for each examine building a single model or separate models for each group, yielding a total of six curves. For real-valued targets (linear regression), we consider two fairness notions (group and individual), and again single or separate models, yielding a total of four curves.

4.2 Price of Fairness

The efficient fairness/accuracy frontiers pictured in Figure 1 can be compared across data sets in a qualitative sense — e.g. to see that in some datasets, the fairness penalty can be substantially decreased with little cost to accuracy. However, they are difficult to compare quantitatively, because the scale of the fairness loss differs substantially from data set to data set. In this section, we give a cross-dataset comparison using a measure (which we call *Price of Fairness*) which has the effect of normalizing the fairness loss across data sets to lie on the same scale.

For a given data set and regression type (linear or logistic), let \mathbf{w}^* be the optimal model absent any fairness penalty (i.e. the empirical risk minimizer when the fairness “regularization” weight $\lambda = 0$). This model will suffer some fairness penalty: it represents the “maximally unfair” point on the fairness/accuracy frontiers from Figure 1. For each dataset, we will fix a normalization such that this fairness penalty is rescaled to be 1, and ask for the cost (in terms of the relative increase in mean squared error) of constraining our predictor to have fairness penalty $\alpha \leq 1$. Equivalently, this is measuring the relative increase in MSE that results from constraining a predictor to have fairness

penalty that is no more than *an α fraction of the fairness penalty of the unconstrained optimal predictor*.

More formally, let $\mathbf{w}^* = \arg \min_{\mathbf{w}} \ell_{\mathcal{P}}(\mathbf{w})$. For any value of $\alpha \in [0, 1]$ we define the *price of fairness* (PoF) as follows:

$$\text{PoF}(\alpha) = \frac{\min_{\mathbf{w}} \ell_{\mathcal{P}}(\mathbf{w}) \text{ subject to } f_{\mathcal{P}}(\mathbf{w}) \leq \alpha f_{\mathcal{P}}(\mathbf{w}^*)}{\ell_{\mathcal{P}}(\mathbf{w}^*)}.$$

Note that by definition, $\text{PoF}(\alpha) \geq 1$, $\text{PoF}(1) = 1$, and that $\text{PoF}(\alpha)$ increases monotonically as α decreases. Larger values represent more severe costs for imposing fairness constraints that ask that the measure of unfairness be small relative to the unconstrained optimum. It is important to note that because this measure asks for the cost of *relative* improvements over the unconstrained optimum, it can be, for example, that the PoF for one fairness penalty case is larger than for another, even if the *absolute* fairness loss for both the numerator and the denominator is smaller in the second case. With this observation in mind, we can move to the empirical findings.

Figure 2 displays the PoF on each of the 6 datasets we study, for each fairness regularizer (individual, hybrid, and group), and for the single and separate model case. We first note that even when normalized on a common scale, we continue to see the diversity across datasets that was apparent in Figure 1. For some datasets (e.g. COMPAS and Sentencing), increasing the fairness constraint by decreasing α has only a mild cost in terms of error. For others (e.g. Communities and Crime, and Law School), the cost increases steadily as we decrease α .

Next, we observe that with this normalization, although the difference between separate and single models remains small on most datasets, on two datasets, differences emerge. In the Law School dataset, restricting to a single model leads to a significantly higher PoF when considering the group fairness metric, compared to allowing separate models. In contrast, on the Adult dataset, restricting to a single model substantially *reduces* the PoF when considering the individual fairness metric.

Finally, this normalization allows us to observe variation across fairness penalties in the *rate of change* in the PoF as α is decreased. In some datasets (e.g. Communities and Crime, and Sentencing), the PoF changes in lock-step across all measures of unfairness. However, for others (e.g. Default), the PoF increases substantially with α when we consider group or hybrid fairness measures, but is much more stable for individual fairness.

5 Conclusions

The use of a complexity regularizer to control overfitting is both standard and well-understood in machine learning. While the use of such a regularizer introduces a trade-off — goodness of fit vs. model complexity — it does not introduce a *tension*, because complexity regularization is always in service of improving generalization, and is usually not a goal in its own right.

In contrast, in this work we have studied a variety of *fairness* regularizers for regression problems, and applied them to data sets in which fairness is not subservient to generalization, but is instead a first-order consideration. Our empirical study has demonstrated that the choice of fairness regularizer (group, individual, hybrid, or other) and the particular data set can have qualitative effects on the trade-off between accuracy and fairness. Combined with recent theoretical results [7, 12, 22] that also highlight the incompatibility of various fairness measures, our results highlight the care that must be taken by practitioners in defining the type of fairness they care about for their particular application, and in determining the appropriate balance between predictive accuracy and fairness.

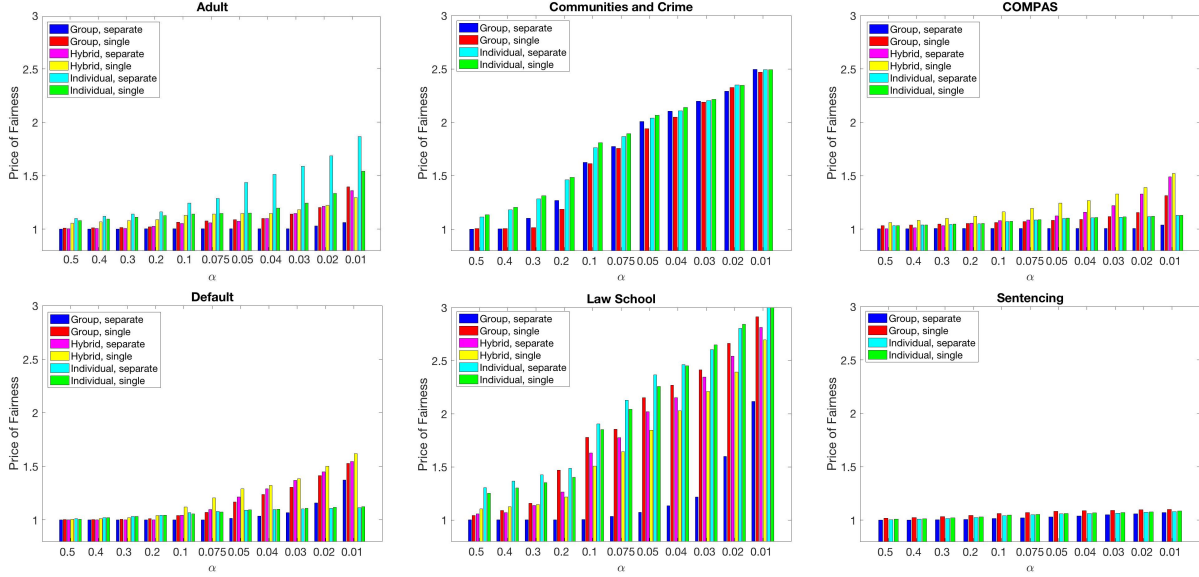


Figure 2: The “Price of Fairness” across data sets, for each type of fairness regularizer, in both the single and separate model case.

References

- [1] Philip Adler, Casey Falk, Sorelle Friedler, Gabriel Rybeck, Carlos Scheidegger, Brandon Smith, and Suresh Venkatasubramanian. Auditing black-box models for indirect influence. In *Proceedings of the 16th International Conference on Data Mining*, pages 1–10, 2016.
- [2] Julia Angwin, Jeff Larson, Surya Mattu, and Lauren Kirchner. Machine bias. *Propublica*, 2016.
- [3] Anna Barry-Jester, Ben Casselman, and Dana Goldstein. The new science of sentencing. *The Marshall Project*, August 8 2015. URL <https://www.themarshallproject.org/2015/08/04/the-new-science-of-sentencing>. Retrieved 4/28/2016.
- [4] Nanette Byrnes. Artificial intolerance. *MIT Technology Review*, March 28 2016. URL <https://www.technologyreview.com/s/600996/artificial-intolerance/>. Retrieved 4/28/2016.
- [5] Toon Calders and Sicco Verwer. Three naive Bayes approaches for discrimination-free classification. *Data Mining and Knowledge Discovery*, 21(2):277–292, 2010.
- [6] Toon Calders, Asim Karim, Faisal Kamiran, Wasif Ali, and Xiangliang Zhang. Controlling attribute effect in linear regression. In *Proceedings of the 13th International Conference on Data Mining*, pages 71–80, 2013.
- [7] Alexandra Chouldechova. Fair prediction with disparate impact: A study of bias in recidivism prediction instruments. *coRR*, abs/1703.00056, 2017.
- [8] Sam Corbett-Davies, Emma Pierson, Avi Feller, Sharad Goel, and Aziz Huq. Algorithmic decision making and the cost of fairness. *CoRR*, abs/1701.08230, 2017.
- [9] Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard Zemel. Fairness through awareness. In *Proceedings of the 3rd Conference on Innovations in Theoretical Computer Science*, pages 214–226, 2012.

- [10] Michael Feldman, Sorelle Friedler, John Moeller, Carlos Scheidegger, and Suresh Venkatasubramanian. Certifying and removing disparate impact. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 259–268, 2015.
- [11] Benjamin Fish, Jeremy Kun, and Ádám Dániel Lelkes. A confidence-based approach for balancing fairness and accuracy. In *Proceedings of the 16th SIAM International Conference on Data Mining*, pages 144–152, 2016.
- [12] Sorelle Friedler, Carlos Scheidegger, and Suresh Venkatasubramanian. On the (im)possibility of fairness. *CoRR*, abs/1609.07236, 2016.
- [13] Sara Hajian and Josep Domingo-Ferrer. A methodology for direct and indirect discrimination prevention in data mining. *IEEE Transactions on Knowledge and Data Engineering*, 25(7):1445–1459, 2013.
- [14] Moritz Hardt, Eric Price, and Nathan Srebro. Equality of opportunity in supervised learning. In *Proceedings of the 30th Annual Conference on Neural Information Processing Systems*, pages 3315–3323, 2016.
- [15] Zubin Jelveh and Michael Luca. Towards diagnosing accuracy loss in discrimination-aware classification: An application to predictive policing. In *2nd Workshop on Fairness, Accountability, and Transparency in Machine Learning*, 2015.
- [16] Kory Johnson, Dean Foster, and Robert Stine. Impartial predictive modeling: Ensuring fairness in arbitrary models. *CoRR*, abs/1608.00528, 2016.
- [17] Matthew Joseph, Michael Kearns, Jamie Morgenstern, and Aaron Roth. Fairness in learning: classic and contextual bandits. In *Proceedings of The 30th Annual Conference on Neural Information Processing Systems*, pages 325–333, 2016.
- [18] Faisal Kamiran and Toon Calders. Data preprocessing techniques for classification without discrimination. *Knowledge and Information Systems*, 33(1):1–33, 2011.
- [19] Faisal Kamiran, Toon Calders, and Mykola Pechenizkiy. Discrimination aware decision tree learning. In *Proceedings of the 10th IEEE International Conference on Data Mining*, pages 869–874, 2010.
- [20] Faisal Kamiran, Asim Karim, and Xiangliang Zhang. Decision theory for discrimination-aware classification. In *Proceedings of the 12th IEEE International Conference on Data Mining*, pages 924–929, 2012.
- [21] Toshihiro Kamishima, Shotaro Akaho, Hideki Asoh, and Jun Sakuma. Fairness-aware classifier with prejudice remover regularizer. In *Proceedings of the European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 35–50, 2012.
- [22] Jon Kleinberg, Sendhil Mullainathan, and Manish Raghavan. Inherent trade-offs in the fair determination of risk scores. In *Proceedings of the 8th Conference on Innovations in Theoretical Computer Science*, 2017.
- [23] Ron Kohavi. Scaling up the accuracy of naive Bayes classifiers: A decision-tree hybrid. In *Proceedings of the 2nd International Conference on Knowledge Discovery and Data Mining*, pages 202–207, 1996.

- [24] Moshe Lichman. UCI machine learning repository, 2013. URL <http://archive.ics.uci.edu/ml>.
- [25] Binh Thanh Luong, Salvatore Ruggieri, and Franco Turini. k-NN as an implementation of situation testing for discrimination discovery and prevention. In *Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 502–510. ACM, 2011.
- [26] Clair Miller. Can an algorithm hire better than a human? *The New York Times*, June 25 2015. URL <http://www.nytimes.com/2015/06/26/upshot/can-an-algorithm-hire-better-than-a-human.html/>. Retrieved 4/28/2016.
- [27] Dino Pedreshi, Salvatore Ruggieri, and Franco 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.
- [28] Michael Redmond and Alok Baveja. A data-driven software tool for enabling cooperative information sharing among police departments. *European Journal of Operational Research*, 141(3):660–678, 2002.
- [29] Cynthia Rudin. Predictive policing using machine learning to detect patterns of crime. *Wired Magazine*, August 2013. URL <http://www.wired.com/insights/2013/08/predictive-policing-using-machine-learning-to-detect-patterns-of-crime/>. Retrieved 4/28/2016.
- [30] Latanya Sweeney. Discrimination in online ad delivery. *Communications of the ACM*, 56(5):44–54, 2013.
- [31] Blake Woodworth, Suriya Gunasekar, Mesrob Ohannessian, and Nathan Srebro. Learning non-discriminatory predictors. *CoRR*, abs/1702.06081, 2017.
- [32] I-Cheng Yeh and Che-hui Lien. The comparisons of data mining techniques for the predictive accuracy of probability of default of credit card clients. *Expert Systems with Applications*, 36(2):2473–2480, 2009.
- [33] Muhammad Zafar, Isabel Valera, Manuel Gomez-Rodriguez, and Krishna Gummadi. Fairness constraints: A mechanism for fair classification. *CoRR*, abs/1507.05259, 2015.
- [34] Muhammad Bilal Zafar, Isabel Valera, Manuel Gomez-Rodriguez, and Krishna P. Gummadi. Fairness beyond disparate treatment & disparate impact: Learning classification without disparate mistreatment. In *Proceedings of the 26th International Conference on World Wide Web*, pages 1171–1180, 2017.
- [35] Richard Zemel, Yu Wu, Kevin Swersky, Toni Pitassi, and Cynthia Dwork. Learning fair representations. In *Proceedings of the 30th International Conference on Machine Learning*, pages 325–333, 2013.

A Missing Details from the Experiments

A.1 Cross Validation for Picking γ

In this section we show how we used cross validation in our experiments to find γ . For each dataset S , our framework requires that we solve optimization problems of the form $\min_{\mathbf{w}} \ell(\mathbf{w}, S) + \lambda f(\mathbf{w}, S) + \gamma \|\mathbf{w}\|_2$ for variable values of λ , where $\ell(\mathbf{w}, S)$ is either MSE (linear regression) or the logistic regression loss. For each λ we picked γ as a function of this λ as follows:

We divided the data into 10 folds (partitions). Let S_i denote the data of fold i and S_{-i} denote all the data except fold i . Let $\Gamma = \{\gamma_1, \dots, \gamma_k\}$ be a set of k potential values for γ we are selecting from and L be a vector of size k initialized to all zero.

Algorithm 1 Procedure for Picking γ

```

for  $j = 1, \dots, k$  (repeat over values of  $\gamma$  in  $\Gamma$ ) do

    for  $i = 1, \dots, 10$  (10-fold cross-validation loop) do
        Compute  $w_i(\gamma_j) \in \arg \min_w \ell(w, S_{-i}) + \lambda f(w, S_{-i}) + \gamma_j \|w\|_2$ 
                                                     $\triangleright$  train a model  $w_i(\gamma_j)$  on all the data except fold  $i$ 
        Set  $\text{Loss}(i, j) = \ell(w_i(\gamma_j), S_i) + \lambda f(w_i(\gamma_j), S_i)$ 
                                                     $\triangleright$  record the test loss of the computed model  $w_i(\gamma_j)$  on fold  $i$ 
         $L(j) = L(j) + \text{Loss}(i, j)$ .
                                                     $\triangleright$  add the test loss to the total loss of  $\gamma_j$  so far.

```

After the loops are over, let k^* denote the entry of L with the smallest value. Then we pick $\gamma_{k^*} \equiv \gamma(\lambda)$ as the regularizer. So our regularized objective function for this particular λ becomes

$$\min_w \ell(w, S) + \lambda f(w, S) + \gamma_\lambda \|w\|$$

We observed in our experiments that γ increases as λ increases. Moreover, for any fixed λ , the value of γ chosen for separate models is usually higher than the value of γ chosen for single model.

A.2 Additional Details

We used 10-fold cross validation to evaluate the performance (fairness and accuracy losses) of a trained model on test data. For smaller datasets (Communities and Crime, COMPAS and Sentencing) we used one run of 10-fold cross validation. For larger datasets (Adult, Default and Law School) we first randomly sampled 30-50% of the dataset (depending on the dataset), ran the 10-fold cross validation on the sampled data and then repeated the experiment 3 times each time with a new random sample. This is because both the running time and the instability of the CVXPY solver would increase when we used the whole datasets for the larger datasets.

While the fairness losses 1, 2, and 3 are defined using all the $n_1 \times n_2$ cross pairs in the dataset, in our experiments we only used $2 \times \text{Minority } n$ random cross pairs where $\text{Minority } n = \min\{n_1, n_2\}$ (see Table 1). This is because: (1) using more cross pairs did not substantially improve the efficiency curves in Figure 1, (2) the CVXPY solver for binary-valued problems would become unstable when using individual fairness if we increase the number of cross pairs significantly. Furthermore, we used the same cross validation folds and cross pairs when experimenting with our different notions of fairness.

Finally, linear models perform poorly on some of our datasets (see e.g. [32]). However, our goal in this work was not to obtain the most accurate models for particular datasets but to study the trade-off between fairness and accuracy and also different notions of fairness for regression problems.

A.3 Datasets

The following is a more through description of the datasets used in this work. We did not use any feature selection methods for any of the datasets. Although, for each dataset, we removed the uninformative features (e.g. various kinds of IDs). Categorical variables were converted to dummy/indicator variables, and an indicator for missing values were included per column with missing values.

The Adult dataset The Adult dataset [23, 24] from the UCI repository contains Census data (hours worked per week, education, sex, age, marital status, and so forth), and has been used to train predictors as to whether an individual earns more or less than \$50,000 per year. Predicting income based upon less sensitive attributes, such as education and occupation, can be used to aid in other economic decisions (related to lending or hiring). It is therefore important to design models which are equally predictive of income regardless of gender or race. We created groups for this dataset based on gender. Furthermore, we extracted the data only from the Adult.data file.

The Communities and Crime dataset The Communities and Crime Dataset [24, 28], from the UCI repository is a dataset which includes many features deemed relevant to violent crime rates (such as the percentage of the community’s population in an urban area, the community’s racial makeup, law enforcement involvement and racial makeup of that law enforcement in a community, amount a community’s law enforcement allocated to drug units) for different communities. This data is provided to train regression models based on this data to predict the amount of violent crime (murder, rape, robbery, and assault) in a given community. If police departments allocate their units to communities based upon a prediction of this form, this should draw concerns of fairness if the predictive behavior is much better for violent predominantly white community than for a violent predominantly nonwhite communities. We consider communities racial makeup as sensitive variable, however, the racial makeup consists of more than two races. We created two groups for this dataset based on whether the percentage of black people in a community (blackPerCap) is higher than the percentage of whites (whitePerCap), indians (indianPerCap), asians (AsianPerCap) and hispanics (HispPerCap) in that community.

The COMPAS dataset The COMPAS dataset contains data from Broward County, Florida originally compiled by ProPublica [2] in which the goal is to predict whether a convicted individual would commit a violent crime in the following two years or not. Following the analysis of Propublica and Corbett-Davies et al. [8] we considered black and white defendants who were assigned COMPAS risk scores within 30 days of their arrest. Furthermore, we restricted ourselves to defendants who “spent at least two years outside a correctional facility without being arrested for a violent crime, or were arrested for a violent crime within this two-year period [8]”. Our goal is to predict the two-year violent recidivism. The dataset includes race, age, sex, number of prior convictions, and COMPAS violent crime risk score (a score between 1 and 10 which we categorized to low, medium and high levels of risks). As mentioned earlier, we created groups for this dataset based on race.

The Default dataset The Default of Credit Card Clients dataset [24, 32] from the UCI repository, contains data from Taiwanese credit card users, such as their credit limit, gender, education, marital

status, history of payment, bill and payment amounts. These features are designed to be used in the prediction of the probability of default payments, although the data has binary labels. It has been noted [32] that linear models tend to perform quite poorly on this task. Given the sensitive nature of lending (and historical inequity of the allocation of credit across races and genders), it is important to understand the tradeoff between fairness and accuracy in settings related to credit. We created groups for this dataset based on gender.

The Law School dataset The bar passage study was initiated in 1991 by Law School Admission Council national longitudinal. The dataset contains records for law students who took the bar exam. The binary outcome indicates whether the student passed the bar exam or not. The features include variables such as cluster, lsat score, undergraduate GPA, zfyGPA, zGPA, full-time status, family income, age and also sensitive variables such as race and gender. The variable cluster is the result of a clustering of similar law schools (which is done apriori), and is used to adjust for the effect of type of law school. zGPA is the z-scores of the students overall GPA and zfyGPA is the first year GPA relative to students at the same law school. We created groups for this dataset based on gender. For more information about this dataset refer to <http://www2.law.ucla.edu/sander/Systemic/Data.htm>.

The Sentencing dataset The Sentencing dataset contains 5969 random sample of all prison inmates primarily from the department of corrections of a state. Moreover data from state agencies that have information of arrests are also added to the dataset. To be sentenced to a state prison, an offender must have been convicted of a felony or must have pled guilty to a felony. The outcome variable, sentence length, is the sentence given by a judge, but is not necessarily the amount of time served in prison. There are usually provisions for review by the state parole board, which can lead to release well before the full nominal sentence is served. The predictors include variables that are legitimate factors in sentencing (e.g., the crime for which a conviction was obtained) as well as sensitive variables like gender. We created groups for this dataset based on gender.