

---

# Selective Regression Under Fairness Criteria

---

Abhin Shah<sup>\*1</sup> Yuheng Bu<sup>\*1</sup> Joshua Ka-Wing Lee<sup>2</sup> Subhro Das<sup>3</sup> Rameswar Panda<sup>3</sup> Prasanna Sattigeri<sup>3</sup>  
Gregory W. Wornell<sup>1</sup>

## Abstract

Selective regression allows abstention from prediction if the confidence to make an accurate prediction is not sufficient. In general, by allowing a reject option, one expects the performance of a regression model to increase at the cost of reducing coverage (i.e., by predicting on fewer samples). However, as we show, in some cases, the performance of a minority subgroup can decrease while we reduce the coverage, and thus selective regression can magnify disparities between different sensitive subgroups. Motivated by these disparities, we propose new fairness criteria for selective regression requiring the performance of every subgroup to improve with a decrease in coverage. We prove that if a feature representation satisfies the *sufficiency* criterion or is *calibrated for mean and variance*, then the proposed fairness criteria is met. Further, we introduce two approaches to mitigate the performance disparity across subgroups: (a) by regularizing an upper bound of conditional mutual information under a Gaussian assumption and (b) by regularizing a contrastive loss for conditional mean and conditional variance prediction. The effectiveness of these approaches is demonstrated on synthetic and real-world datasets.

## 1. Introduction

As the adoption of machine learning (ML) based systems accelerates in a wide range of applications, including critical workflows such as healthcare management (Bellamy et al., 2018), employment screening (Selbst et al., 2019), automated loan processing, there is a renewed focus on the

trustworthiness of such systems. An important attribute of a trustworthy ML system is to reliably estimate the uncertainty in its predictions. For example, consider a loan approval ML system where the prediction task is to suggest appropriate loan terms (e.g., loan approval, interest rate). If the model’s uncertainty in its prediction is high for an applicant, the prediction can be rejected to avoid potentially costly errors. The users of the system, i.e., the decision-maker, can intervene, and take remedial actions such as gathering more information for applicants with rejected model predictions or involving a special human credit committee before arriving at a decision. This paradigm is known as *prediction with reject-option* or *selective prediction*.

By making the tolerance for uncertainty more stringent, the user expects the error rate of the predictions made by the system to decrease as the system makes predictions for fewer samples, i.e., as coverage is reduced. Although the error may lessen over the entire population, (Jones et al., 2020) demonstrated that for *selective classification*, this may not be true for different sub-populations. In other words, selective classification could magnify disparities across different sensitive groups (e.g., race, gender). For example, in the loan approval ML system, the error rate for a sensitive group could increase with a decrease in coverage. To mitigate such disparities, (Lee et al., 2021; Schreuder & Chzhen, 2021) proposed methods for performing fair selective classification.

In this work, we demonstrate and investigate the performance disparities across different subgroups for *selective regression* as well as develop novel methods to mitigate such disparities. Similar to (Lee et al., 2021), we do not assume access to the identity of sensitive groups at test time. Compared to selective classification, one major challenge to tackling the aforementioned disparities in selective regression is as follows: in selective classification, generating an *uncertainty measure* (i.e., the model’s uncertainty for its prediction) from an existing classifier is straightforward. For example, one could take the softmax output of an existing classifier as an uncertainty measure. In contrast, there is no direct method to extract an uncertainty measure from an existing regressor designed only to predict the conditional mean.

<sup>\*</sup>Equal contribution <sup>1</sup>Massachusetts Institute of Technology

<sup>2</sup>This work was done while J. Lee was at Massachusetts Institute of Technology; the author is now with Snap <sup>3</sup>MIT-IBM Watson AI Lab, IBM Research. Correspondence to: Abhin Shah <abhin@mit.edu>, Yuheng Bu <buyuheng@mit.edu>.

**Contributions.** First, we show via the Insurance dataset and a toy example that selective regression, like selective classification, can decrease the performance of some subgroups when *coverage* (the fraction of samples for which a decision is made) is reduced (see Section 3.1).

1. Motivated by this, we provide a novel fairness criteria (Definition 1) for selective regression, namely, *monotonic selective risk*, which requires the risk of each subgroup to monotonically decrease with a decrease in coverage (see Section 3.2).
2. We prove that if a feature representation satisfies the standard *sufficiency* criterion or is *calibrated for mean and variance* (Definition 2), then the *monotonic selective risk* criteria is met (see Theorem 4.1 and 4.2).
3. We provide two neural network-based algorithms: one to impose the sufficiency criterion by regularizing an upper bound of conditional mutual information under a Gaussian assumption (see Section 5.1) and the other to impose the calibration for mean and variance by regularizing a contrastive loss (see Section 5.2).
4. Finally, we empirically<sup>1</sup> demonstrate the effectiveness of these algorithms on real-world datasets (see Section 6).

## 2. Background

### 2.1. Fair Regression

In standard (supervised) regression, given pairs of input variables  $X \in \mathcal{X}$  (e.g., demographic information) and target variable  $Y \in \mathbb{R}$  (e.g., annual medical expenses), we want to find a predictor  $f : \mathcal{X} \rightarrow \mathbb{R}$  that best estimates the target variable for new input variables. Formally, given a set of predictors  $\mathcal{F}$  and a set of training samples of  $X$  and  $Y$ , i.e.,  $\{(x_1, y_1), \dots, (x_n, y_n)\}$ , the goal is to construct  $f^* \in \mathcal{F}$  which minimizes the mean-squared error (MSE):

$$f^* = \arg \min_{f \in \mathcal{F}} \mathbb{E}[(Y - f(X))^2]. \quad (1)$$

In fair regression, we augment the goal in (1) by requiring our predictor to retain “fairness” w.r.t. some sensitive attributes  $D \in \mathcal{D}$  (e.g., race, gender). For example, we may want our predictions of annual medical expenses using the demographic information not to discriminate w.r.t. race.

In this work, we assume  $D$  to be discrete and consider members with the same value of  $D$  as being in the same *subgroup*. While numerous criteria have been proposed to enforce fairness in machine learning, we focus on the notion of *subgroup risks* (Williamson & Menon, 2019), which ensures that the predictor  $f$  behaves similarly (in terms of risks) across all subgroups. This notion, also known as accuracy disparity, has been used frequently

<sup>1</sup>The source code is available at [github.com/Abhin02/fair-selective-regression](https://github.com/Abhin02/fair-selective-regression).

in fair regression, e.g., (Chzhen & Schreuder, 2020; Chi et al., 2021), and has also received attention in the field of domain generalization, e.g., (Krueger et al., 2021). Formally, given a set of training samples of  $X$ ,  $Y$ , and  $D$ , i.e.,  $\{(x_1, y_1, d_1), \dots, (x_n, y_n, d_n)\}$ , the goal is to construct  $f^* \in \mathcal{F}$  which minimizes the overall MSE subject to the subgroup MSE being equal for all subgroups:

$$f^* = \arg \min_{f \in \mathcal{F}} \mathbb{E}[(Y - f(X))^2] \text{ s.t. } \forall d, d' \in \mathcal{D}, \\ \mathbb{E}[(Y - f(X))^2 | D = d] = \mathbb{E}[(Y - f(X))^2 | D = d']. \quad (2)$$

In this work, we consider the scenario where the sensitive attribute is available only during training i.e., we do not assume access to the sensitive attribute at test time.

### 2.2. Selective Regression

In selective regression, given pairs of input variables  $X \in \mathcal{X}$  and target variable  $Y \in \mathbb{R}$ , for new input variables, the system has a choice to: (a) make a prediction of the target variable or (b) abstain from a prediction (if it is not sufficiently confident). In the example of predicting annual medical expenses, we may prefer abstention in certain scenarios to avoid harms arising from wrong predictions. By only making predictions for those input variables with low prediction uncertainty, the performance (in terms of MSE) is expected to improve. Formally, in addition to a predictor  $f : \mathcal{X} \rightarrow \mathbb{R}$  that best estimates the target variable for new input variables, we need to construct a rejection rule  $\Gamma : \mathcal{X} \rightarrow \{0, 1\}$  that decides whether or not to make a prediction for new input variables. Thereafter, for  $X = x$ , the system outputs  $f(x)$  as the prediction when  $\Gamma(x) = 1$ , and makes no prediction if  $\Gamma(x) = 0$ .

There are two quantities that characterize the performance of selective regression: (i) *coverage*, i.e., the fraction of samples that the system makes predictions on, which is denoted by  $c(\Gamma) = \mathbb{P}(|\Gamma(X)| = 1)$  and (ii) the MSE when prediction is performed

$$\mathbb{E}[(Y - f(X))^2 | \Gamma(X) = 1].$$

In order to construct a rejection rule  $\Gamma$ , we need some measure of the uncertainty  $g(\cdot)$  associated with a prediction  $f(\cdot)$ . Then, the rejection rule  $\Gamma$  can be defined as:

$$\Gamma(x) := \begin{cases} 1, & \text{if } g(x) \leq \tau \\ 0, & \text{otherwise.} \end{cases}$$

where  $\tau$  is the parameter that balances the MSE vs. coverage tradeoff: larger  $\tau$  results in a larger coverage but also yields a larger MSE. Therefore,  $\tau$  can be interpreted as the cost for not making predictions. As discussed in (Zaoui et al., 2020), a natural choice for the uncertainty measure  $g(\cdot)$  could be the conditional variance of  $Y$  given  $X$ .

The goal of selective regression is to build a model with (a) high coverage and (b) low MSE. However, there may not exist any  $\tau$  for which both (a) and (b) are satisfied simultaneously. Therefore, in practice, the entire MSE vs. coverage tradeoff curve is generated by sweeping over all possible values of  $\tau$  allowing the system designer to choose any convenient operating point.

### 2.3. Related Work

**Selective Regression.** While selective classification has received a lot of attention in the machine learning community (Chow, 1957; 1970; Hellman, 1970; Herbei & Wegkamp, 2006; Bartlett & Wegkamp, 2008; Nadeem et al., 2009; Lei, 2014; Geifman & El-Yaniv, 2017), there is very limited work on selective regression. It is also known that existing methods for selective classification cannot be used directly for selective regression (Jiang et al., 2020). (Wiener & El-Yaniv, 2012) studied regression with reject option to make predictions inside a ball of small radius around the regression function with high probability. (Geifman & El-Yaniv, 2019) proposed SelectiveNet, a neural network with an integrated reject option, to optimize both classification (or regression) performance and rejection rate simultaneously. (Zaoui et al., 2020) considered selective regression with a fixed rejection rate and derived the optimal rule which relies on thresholding the conditional variance. (Jiang et al., 2020) analyzed selective regression with the goal to minimize the rejection rate given a regression risk bound. We emphasize that none of these works study the question of fairness in selective regression.

**Fair Regression.** (Calders et al., 2013), one of the first works on fair regression, studied linear regression with constraints on the mean outcome or residuals of the models. More recently, several works including (Berk et al., 2017; Pérez-Suay et al., 2017; Komiyama & Shimao, 2017; Komiyama et al., 2018; Fitzsimons et al., 2018; Raff et al., 2018; Agarwal et al., 2019; Nabi et al., 2019; Oneto et al., 2020) considered fair regression employing various fairness criteria. (Mary et al., 2019) and (Lee et al., 2020) enforced independence between prediction and sensitive attribute by ensuring that the maximal correlation is below a fixed threshold. (Chzhen et al., 2020) considered learning an optimal regressor requiring the distribution of the output to remain the same conditioned on the sensitive attribute. We emphasize that none of these works could be used for selective regression as they are designed to only predict the conditional mean and not the conditional variance (i.e., the uncertainty).

**Fairness Criteria.** Numerous metrics and criteria have been proposed to enforce fairness in machine learning (Verma & Rubin, 2018). Many of these criteria are mutually

exclusive outside of trivial cases (Gölz et al., 2019). Further, the existing approaches also differ in the way they enforce these criteria: (a) pre-processing methods (Zemel et al., 2013; Louizos et al., 2015; Calmon et al., 2017) modify the training set to ensure fairness of any learned model, (b) post-processing methods (Hardt et al., 2016; Pleiss et al., 2017; Corbett-Davies et al., 2017) transform the predictions of a trained model to satisfy a measure of fairness, and (c) in-processing methods (Kamishima et al., 2011; Zafar et al., 2017; Agarwal et al., 2018) modify the training process to directly learn fair predictors e.g., minimizing a loss function that accounts for both accuracy and fairness as in (2). Additionally, these criteria also differ in the kind of fairness they consider (see (Mehrabi et al., 2021; Castelnovo et al., 2022) for details): (a) group fairness ensures that subgroups that differ by sensitive attributes are treated similarly and (b) individual fairness ensures that individuals who are similar (with respect to some metric) are treated similarly. In this work, we consider group fairness and propose a novel fairness criteria specific to selective regression. Our approach falls under the umbrella of in-processing methods as will be evident in Section 5 (see (5), (8), and (9)).

## 3. Fair Selective Regression


While fair regression and selective regression have been independently studied before, consideration of fair selective regression (i.e., selective regression while ensuring fairness) is missing in the literature. In this section, we explore the disparities between different subgroups that may arise in selective regression. Building on top of this, we formulate a notion of fairness for selective regression.

### 3.1. Disparities in Selective Regression

(Jones et al., 2020) argued that, in selective classification, increasing abstentions (i.e., decreasing coverage) could *decrease* accuracies on some subgroups and observed this behavior for CelebA dataset. In this section, we show that a similar phenomenon can be observed in selective regression.

**Insurance dataset.** Consider the task of predicting the annual medical expenses charged to patients from input variables such as age, BMI, number of children, etc., as in Insurance dataset. Suppose we construct our predictor as the conditional expectation and our uncertainty measure as the conditional variance following (Zaoui et al., 2020). Then, generating the subgroup MSE vs. coverage tradeoff curve<sup>2</sup> across the subgroups induced by gender, as shown in Figure 1a, we see that while decreasing the coverage improves the performance for the majority subgroup (i.e., females), the performance for the minority subgroup (i.e., males)

<sup>2</sup>We have used `Baseline 1` to generate this (see Section 6).



(a) The disparity for the Insurance dataset when the predictor is the conditional expectation and the uncertainty measure is the conditional variance.

(b) The disparity for toy example when the predictor is the conditional expectation and the uncertainty measure is the conditional variance of  $Y$  given  $X_1$  and  $X_2$ .

(c) The disparity could be mitigated for toy example when the predictor is the conditional expectation and the uncertainty measure is the conditional variance of  $Y$  given  $X_1$ .

Figure 1: Subgroup MSE vs. coverage to illustrate disparities in selective regression via Insurance dataset and toy example.

degrades. Thus, unknowingly, selective regression can magnify disparities between different sensitive subgroups.

Motivated to further understand this phenomenon, we explicitly recreate it via the following toy example.

**Toy example.** Consider predicting  $Y$  from two normalized input variables  $X_1$  and  $X_2$  that are generated i.i.d. from the uniform distribution over  $[0, 1]$ . Suppose we have a binary sensitive attribute  $D$  with  $\mathbb{P}(D=0) = 0.9$ , where  $D = 0$  represents majority and  $D = 1$  represents minority. To illustrate the disparities that may arise in selective regression, we let the distribution of  $Y$  differ with  $D$ . More specifically, for majority subgroup, we let the target be

$$Y|_{D=0} = X_1 + X_2 + \mathcal{N}(0, 0.1X_1 + 0.15X_2)$$

and, for minority subgroup, we let the target be

$$Y|_{D=1} = X_1 + X_2 + \mathcal{N}(0, 0.1X_1 + 0.15(1 - X_2)).$$

To summarize, the only difference is that for the majority, the variance of  $Y$  increases in  $X_2$ , and for the minority, the variance of  $Y$  decreases in  $X_2$ .

In this case, the conditional variance  $\text{Var}(Y|X)$ , i.e., our uncertainty measure would mainly capture the behavior of the majority subgroup  $D = 0$ , i.e., the subgroup with more samples. Since  $\text{Var}(Y|X, D=0)$  differs significantly from  $\text{Var}(Y|X, D=1)$ ,  $\text{Var}(Y|X)$  may not be a good measure of uncertainty for the minority subgroup  $D = 1$ . As a result, for minority subgroup, when we decrease the coverage, we may make predictions on samples that erroneously achieve low uncertainty based on  $\text{Var}(Y|X)$ . This results in an increase in the MSE for the subgroup  $D = 1$  (Figure 1b).

An alternative could be to use the conditional variance of  $Y$  given  $X_1$  (instead of both  $X_1$  and  $X_2$ ) as our uncertainty measure. While this may be a slightly worse measure of uncertainty for  $D = 0$  than  $\text{Var}(Y|X)$ , it is a much better measure of uncertainty for  $D = 1$ . All in all, using this

uncertainty measure, when we decrease the coverage, we make predictions on samples with low uncertainty for all subgroups resulting in a decrease in the MSE for every subgroup as shown in Figure 1c (albeit at the cost of a slight increase in the overall MSE).

It is important to note that the toy example is designed to highlight the disparities when the uncertainty measure, a component of selective regression, is designed unfairly. The disparities could generally occur due to the predictor or the uncertainty measure (or both).

### 3.2. When is Selective Regression Fair?

Motivated by the disparities (that may arise) in selective regression, as shown above, we define the first notion of fairness for selective regression, which is called *monotonic selective risk*. This notion requires our predictor and uncertainty measure to ensure the primary goal of selective regression that the subgroup MSE decreases monotonically with a decrease in coverage for every subgroup. The subgroup MSE for  $d \in \mathcal{D}$ , as a function of the predictor  $f$  and the uncertainty measure  $g$ , for a fixed coverage (parameterized by  $\tau$ ) is given by

$$\text{MSE}(f, g, \tau, d) = \mathbb{E}[(Y - f(X))^2 | g(X) \leq \tau, D = d].$$

Now, we formalize the criteria of *monotonic selective risk* which ensures that no subgroup is discriminated against when the coverage is reduced in selective regression.

**Definition 1.** We say that a predictor  $f$  and an uncertainty measure  $g$  satisfy *monotonic selective risk* if for any  $\tau < \tau'$

$$\text{MSE}(f, g, \tau, d) \leq \text{MSE}(f, g, \tau', d) \quad \forall d \in \mathcal{D}.$$

Inspired by the success of representation based learning in machine learning (Bengio et al., 2013), we seek to find a representation  $\Phi : \mathcal{X} \rightarrow \mathcal{H}$  that maps the input variable  $X \in \mathcal{X}$  to an intermediate representation  $\Phi(X) \in \mathcal{H}$  and

use  $\Phi(X)$  to construct our predictor  $f : \mathcal{H} \rightarrow \mathbb{R}$  and our uncertainty measure  $g : \mathcal{H} \rightarrow \mathbb{R}_{\geq 0}$ . Then, for  $X = x$ , our prediction and uncertainty measure are (with a slight abuse of notation)  $f(\Phi(x))$  and  $g(\Phi(x))$  respectively.

## 4. Theoretical Results

In this section, we show that under certain conditions on the feature representation  $\Phi$ , the conditional mean as the predictor and the conditional variance as the uncertainty measure satisfy *monotonic selective risk* (Definition 1).

### 4.1. Sufficiency

The *sufficiency* criterion requires<sup>3</sup>  $Y \perp D | \Phi(X)$ , i.e., the learned representation  $\Phi(X)$  completely subsumes all information about the sensitive attribute that is relevant to the target variable (Cleary, 1966). Sufficiency is closely tied with learning domain-invariant feature representation (Arjovsky et al., 2019; Creager et al., 2021) and has been used in fair selective classification (Lee et al., 2021).

The theorem below shows that if the feature representation is sufficient, then the choice of conditional mean as the predictor and the conditional variance as the uncertainty measure ensures the fairness criteria of *monotonic selective risk*, i.e., the subgroup MSE decreases monotonically with coverage for all subgroups. See Appendix A for a proof.

**Theorem 4.1.** *Suppose the representation  $\Phi(X)$  is sufficient i.e.,  $Y \perp D | \Phi(X)$ . Let  $f(\Phi(X)) = \mathbb{E}[Y | \Phi(X)]$  and  $g(\Phi(X)) = \text{Var}[Y | \Phi(X)]$ . Then, for any  $d \in \mathcal{D}$  and any  $\tau < \tau'$ , we have  $\text{MSE}(f, g, \tau, d) < \text{MSE}(f, g, \tau', d)$ .*

### 4.2. Calibration for mean and variance

In practice, the conditional independence  $Y \perp D | \Phi(X)$  required by sufficiency may be too difficult to satisfy. Since we only care about MSE, which depends on the first and second-order moments, one could relax the sufficiency condition by requiring the representation  $\Phi$  to be such that these moments are the same across all subgroups. This inspires our notion of  $\Phi$  *calibrated for mean and variance*.

**Definition 2.** We say a representation  $\Phi(X)$  is *calibrated for mean and variance* if

$$\begin{aligned} \mathbb{E}[Y | \Phi(X), d] &= \mathbb{E}[Y | \Phi(X)] \quad \forall d \in \mathcal{D}, \\ \text{Var}[Y | \Phi(X), d] &= \text{Var}[Y | \Phi(X)] \quad \forall d \in \mathcal{D}. \end{aligned}$$

The theorem below shows that if the feature representation is *calibrated for mean and variance*, then the choice of

<sup>3</sup>While conventionally sufficiency requires  $Y \perp D | \mathbb{E}[\Phi(X)]$ , this notion has been adapted to incorporate feature representation/score function (e.g., Sec 1.1 in (Liu et al., 2019)), Sec 3.4 in (Castelnuovo et al., 2022).

conditional mean as the predictor and the conditional variance as the uncertainty measure ensures the fairness criteria of *monotonic selective risk*, i.e., the subgroup MSE decreases monotonically with coverage for all subgroups. The proof is similar to the proof of Theorem 4.1 and is omitted.

**Theorem 4.2.** *Suppose the representation  $\Phi(X)$  is calibrated for mean and variance. Let  $f(\Phi(X)) = \mathbb{E}[Y | \Phi(X)]$  and  $g(\Phi(X)) = \text{Var}[Y | \Phi(X)]$ . Then, for any  $d \in \mathcal{D}$  and any  $\tau < \tau'$ , we have  $\text{MSE}(f, g, \tau, d) < \text{MSE}(f, g, \tau', d)$ .*

## 5. Algorithm Design

In this section, we provide two neural network-based algorithms: one to impose sufficiency and the other to impose the calibration for mean and variance.

### 5.1. Imposing sufficiency

To simplify our algorithm when directly enforcing sufficiency, we utilize the framework of heteroskedastic neural network (Gal, 2016). A heteroskedastic neural network, which requires training only a single neural network, assumes that the distribution of  $Y$  conditioned on  $X$  is Gaussian. Then, it is trained by minimizing the negative log likelihood:

$$L_G(\Phi, \theta) \triangleq - \sum_{i=1}^n \log \mathbb{P}_G(y_i | \Phi(x_i); \theta)$$

where  $\mathbb{P}_G(y | \Phi(x); \theta)$  represents a Gaussian distribution with  $f(\Phi(x); \theta_f)$  and  $g(\Phi(x); \theta_g)$  as the conditional mean and the conditional variance (of  $Y$  given  $\Phi(X)$ ) respectively. The feature representation  $\Phi$  is parameterized by  $\theta_\Phi$  and the neural network is supposed to learn the parameters  $\theta_\Phi$  and  $\theta = (\theta_f, \theta_g)$ .

To impose sufficiency, we augment minimizing the negative log likelihood as follows:

$$\min_{\theta, \Phi} L_G(\Phi, \theta), \quad \text{s.t. } Y \perp D | \Phi(X).$$

To relax the hard constraint of  $Y \perp D | \Phi(X)$  into a soft constraint, we use the conditional mutual information, since  $Y \perp D | \Phi(X)$  is equivalent to  $I(Y; D | \Phi(X)) = 0$ . For  $\lambda \geq 0$ ,

$$\min_{\theta, \Phi} L_G(\Phi, \theta) + \lambda I(Y; D | \Phi(X)).$$

As discussed in (Lee et al., 2021), existing methods using mutual information for fairness are ill-equipped to handle conditioning on the feature representation  $\Phi(\cdot)$ . Therefore, we further relax the soft constraint by using the following

**Algorithm 1** Heteroskedastic neural network with sufficiency-based regularizer

---

**Input:** training samples  $\{(x_i, y_i, d_i)\}_{i=1}^n$ , regularizer  $\lambda$   
**Draw:**  $\{\tilde{d}_1, \dots, \tilde{d}_n\}$  drawn i.i.d. from  $\hat{\mathbb{P}}_D$   
**Initialize:**  $\theta, \theta_\Phi$ , and  $w^{(d)}$  with pre-trained models  
**Initialize:**  $n_d =$  number of samples in group  $d, \forall d \in \mathcal{D}$   
**for** each training iteration **do**  
   **for** each batch **do**  
     **for**  $d = 1, \dots, |\mathcal{D}|$  **do**  
       # update subgroup-specific mean/variance predictor  
        $w^{(d)} \leftarrow w^{(d)} - \frac{1}{n_d} \eta_w \nabla_w L_d(w)$   
     **end for**  
   **end for**  
   **for** each batch **do**  
     # update feature extractor  
      $\theta_\Phi \leftarrow \theta_\Phi - \frac{1}{n} \eta \nabla_{\theta_\Phi} (L_G(\Phi, \theta) + \lambda L_R(\Phi))$   
     # update mean/variance predictor  
      $\theta \leftarrow \theta - \frac{1}{n} \eta \nabla_\theta L_G(\Phi, \theta)$   
   **end for**  
**end for**

---

upper bound for  $I(Y; D | \Phi(X))$  from (Lee et al., 2021):

$$I(Y; D | \Phi(X)) \leq \mathbb{E}_{\Phi(X), Y, D} [\log \mathbb{P}(Y | \Phi(X), D)] \quad (3)$$

$$- \mathbb{E}_D [\mathbb{E}_{\Phi(X), Y} [\log \mathbb{P}(Y | \Phi(X), D)]] .$$

where equality is achieved if and only if  $Y \perp D | \Phi(X)$ .

In order to compute the upper bound in (3), we need to learn the unknown distribution  $\mathbb{P}(y | \Phi(x), d)$ . We approximate this by  $\mathbb{P}_G(y | \Phi(x), d; w)$  which is a Gaussian distribution with  $f(\Phi(x), d; w_f)$  and  $g(\Phi(x), d; w_g)$  as the conditional mean and the conditional variance of  $Y$  given  $\Phi(X)$  and  $D$ , respectively. The neural network is supposed to learn the parameters  $w = (w_f, w_g)$ .

In scenarios where  $\Phi(X)$  is high-dimensional compared to  $D$ , it would be preferred to approximate  $\mathbb{P}_G(y | \Phi(x), d; w)$  by  $\mathbb{P}_G(y | \Phi(x); w^{(d)})$ , i.e., train a *subgroup-specific* Gaussian model with parameters  $w^{(d)}$  for each  $d \in \mathcal{D}$  instead of using  $D$  as a separate input to ensure that  $D$  has an effect in  $\mathbb{P}_G(y | \Phi(x), d; w)$ . Then, for  $d \in \mathcal{D}$ ,

$$w^{(d)} = \arg \min_w L_d(w), \quad \text{where} \quad (4)$$

$$L_d(w) \triangleq - \sum_{i: d_i=d} \log \mathbb{P}_G(y_i | \Phi(x_i); w).$$

To summarize, the first term of the upper bound in (3) is approximated by the log-likelihood of the training samples using  $\mathbb{P}_G(y | \Phi(x); w^{(d_i)})$  for each subgroup  $\tilde{d}_i \in \mathcal{D}$  (i.e., subgroup-specific loss). Then, drawing  $\tilde{d}_i$  i.i.d from  $\mathbb{P}_D$  i.e., the marginal distribution of  $D$  (which could be approximated by  $\hat{\mathbb{P}}_D$ ), the second term of the upper bound in (3) is approximated by the negative log-likelihood

of the samples using the randomly-selected Gaussian model  $\mathbb{P}_G(y | \Phi(x); w^{(\tilde{d}_i)})$  for each subgroup  $\tilde{d}_i \in \mathcal{D}$  (i.e., subgroup-agnostic loss). Combining everything and replacing all the expectations in (3) with empirical averages, the regularizer is given by

$$L_R(\Phi) \triangleq \sum_{i=1}^n \log \left( \frac{\mathbb{P}_G(y_i | \Phi(x_i); w^{(d_i)})}{\mathbb{P}_G(y_i | \Phi(x_i); w^{(\tilde{d}_i)})} \right),$$

where  $\tilde{d}_i$  are drawn i.i.d. from the marginal distribution  $\hat{\mathbb{P}}_D$ . Summarizing, the overall objective is

$$\min_{\theta, \Phi} L_G(\Phi, \theta) + \lambda L_R(\Phi). \quad (5)$$

As shown in Algorithm 1, we train our model by alternating between the fitting subgroup-specific models in (4) and feature updating in (5).

## 5.2. Imposing calibration for mean and variance

To achieve the calibration for mean and variance, we let the representation  $\Phi = (\Phi_1, \Phi_2)$ . Then, to enable the use of the residual-based neural network (Hall & Carroll, 1989), we let the conditional expectation depend only on  $\Phi_1$  i.e.,  $f(\Phi(X)) = f(\Phi_1(X))$  and let the conditional variance depend only on  $\Phi_2$  i.e.,  $g(\Phi(X)) = g(\Phi_2(X))$ . This method is useful in scenarios where the conditional Gaussian assumption in the Section 5.1 does not hold.

In a residual-based neural network, the conditional mean-prediction network is trained by minimizing:

$$L_{S1}(\Phi_1, \theta_f) \triangleq \sum_{i=1}^n (y_i - f(\Phi_1(x_i); \theta_f))^2.$$

The feature representation  $\Phi_1$  is parameterized by  $\theta_{\Phi_1}$  and the mean-prediction network is supposed to learn the parameters  $\theta_{\Phi_1}$  and  $\theta_f$ .

Then, the conditional variance-prediction network is trained by fitting the residuals obtained from the mean-prediction network, i.e.,  $r_i \triangleq (y_i - f(\Phi_1(x_i); \theta_f))^2$  by minimizing:

$$L_{S2}(\Phi_2, \theta_g) \triangleq \sum_{i=1}^n (r_i - g(\Phi_2(x_i); \theta_g))^2.$$

The feature representation  $\Phi_2$  is parameterized by  $\theta_{\Phi_2}$  and the variance-prediction network is supposed to learn the parameters  $\theta_{\Phi_2}$  and  $\theta_g$ .

To impose calibration under mean, we need to convert the following hard constraint

$$\mathbb{E}[Y | \Phi_1(X), D] = \mathbb{E}[Y | \Phi_1(X)] \quad (6)$$

into a soft constraint. We do this by using the following contrastive loss:

$$\mathbb{E}_D \left[ \mathbb{E}_{\Phi(X), Y} \left[ (Y - \mathbb{E}[Y|\Phi_1(X), D])^2 \right] \right. \\ \left. - \mathbb{E}_{\Phi(X), Y, D} \left[ (Y - \mathbb{E}[Y|\Phi_1(X), D])^2 \right] \right], \quad (7)$$

which is inspired from (3) and obtained by replacing the negative log-likelihood  $-\log \mathbb{P}(Y|\Phi(X), D)$  in (3) by the MSE achieved using the representation  $\Phi_1(X)$  and sensitive attribute  $D$ . We emphasize that (7) is zero when (6) holds and therefore (7) is a relaxation of (6).

To compute (7), we need to learn the unknown conditional expectation  $\mathbb{E}[Y|\Phi_1(X), D]$ . Similar to Section 5.1, we approximate this by  $f(y|\Phi_1(x); w_f^{(d)})$  (i.e., train a *subgroup-specific* mean-prediction model with parameters  $w_f^{(d)}$  for each  $d \in \mathcal{D}$ , instead of using  $D$  as a separate input). Similar to Section 5.1, combining everything and replacing all the expectations in (7) with empirical averages, the regularizer for the mean-prediction network is given by

$$L_{R1}(\Phi_1) \triangleq \sum_{i=1}^n \left( (y_i - f(\Phi_1(x_i); w_f^{(\tilde{d}_i)}))^2 \right. \\ \left. - (y_i - f(\Phi_1(x_i); w_f^{(d_i)}))^2 \right),$$

where  $\tilde{d}_i$  are drawn i.i.d. from  $\mathbb{P}_D$  i.e., the marginal distribution of  $D$  (approximated by  $\hat{\mathbb{P}}_D$ ) and for  $d \in \mathcal{D}$ ,

$$w_f^{(d)} = \arg \min_w \sum_{i: d_i=d} (y_i - f(\Phi_1(x_i); w))^2.$$

In summary, the overall objective for mean-prediction is

$$\min_{\theta_f, \Phi_1} L_{S1}(\Phi_1, \theta_f) + \lambda_1 L_{R1}(\Phi_1). \quad (8)$$

Once the mean-prediction network is trained, we obtain the residuals and train the variance-prediction network  $g$  using a similar regularizer:

$$\min_{\theta_g, \Phi_2} L_{S2}(\Phi_2, \theta_g) + \lambda_2 L_{R2}(\Phi_2), \quad (9)$$

where  $L_{S2}$  and  $L_{R2}$  are defined in a similar manner as  $L_{S1}$  and  $L_{R1}$ . More details about these and the pseudo-code (i.e., Algorithm 2) are provided in the Appendix B.

## 6. Experimental Results

**Datasets.** We test our algorithms on Insurance and Crime datasets, and provide an application of our method in Causal Inference via IHDP dataset. These datasets (summarized in Table 1) are selected due to their potential fairness concerns, e.g., (a) presence of features often associated with possible discrimination, such as race and sex, and (b) potential sensitivity regarding the predictions being made such as medical expenses, violent crimes, and cognitive test score.

Table 1: Summary of datasets.

Dataset	Target	Attribute
Insurance	Medical Expenses	Sex
Crime	Crimes per Population	Race
IHDP	Cognitive Test Score	Sex

**Insurance.** The Insurance dataset (Lantz, 2019) considers predicting total annual medical expenses charged to patients using demographic statistics. Following (Chi et al., 2021), we use sex as the sensitive attribute:  $D = 1$  (i.e., minority) if male otherwise 0. After preprocessing (see Appendix C), the dataset contains 1000 samples (338 with  $D = 1$  and 662 with  $D = 0$ ) and 5 features.

**Communities and Crime.** The Crime dataset (Redmond & Baveja, 2002) considers predicting the number of violent crimes per 100K population using socio-economic information of communities in U.S. Following (Chi et al., 2021), we use race (binary)<sup>4</sup> as the sensitive attribute:  $D = 1$  (i.e., minority) if the population percentage of the black is more or equal to 20 otherwise 0. After preprocessing (see Appendix C), the dataset contains 1994 samples (532 with  $D = 1$  and 1462 with  $D = 0$ ) and 99 features.

**IHDP.** The IHDP dataset (Hill, 2011) is generated based on a randomized control trial targeting low-birth-weight, premature infants. In the treated group, the infants were provided with both intensive and high-quality childcare and specialist home visits. The task is to predict the infants' cognitive test scores. We let sex be the sensitive attribute and observe that *male* is the minority group (i.e.,  $D = 1$ ) in both the control group and the treatment group. After preprocessing (see Appendix C), the control group has 608 samples (296 with  $D = 1$  and 312 with  $D = 0$ ) and the treatment group has 139 samples (67 with  $D = 1$  and 72 with  $D = 0$ ). The dataset contains 25 features.

**Choice of  $\lambda$ .** We observe our algorithms to be agnostic to the choice of  $\lambda$  as long as it is in a reasonable range, i.e.,  $\lambda \in [0.5, 3]$ . To be consistent, we set  $\lambda = 1$  throughout.

**Baselines.** We compare against the following baselines:

- Baseline 1: Heteroskedastic neural network without any regularizer i.e., Algorithm 1 with  $\lambda = 0$ .
- Baseline 2: Residual-based neural network without any regularizer i.e., Algorithm 2 with  $\lambda = 0$ .

**Experimental Details.** In all of our experiments, we use two-layer neural networks, and train our model only once on a fixed training set. We evaluate and report the empirical findings on a held-out test set with a train-test split ratio

<sup>4</sup>We provide results for the scenario where race can take more than two values in in Appendix C.8.

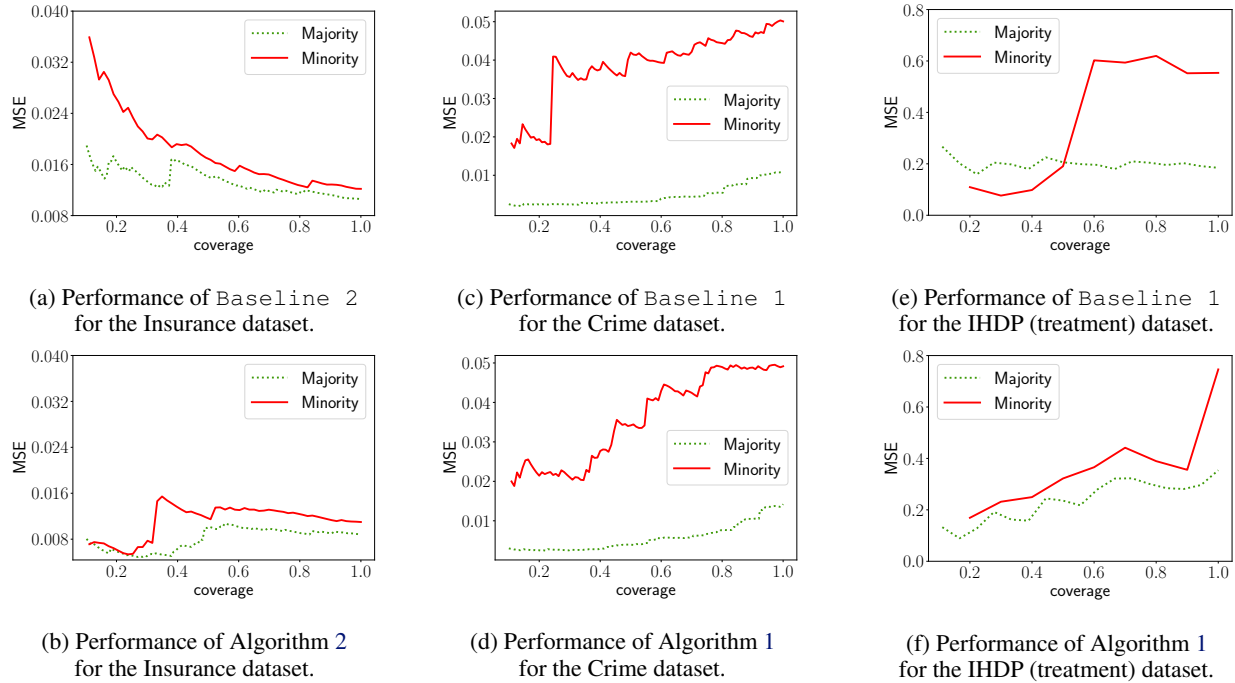


Figure 2: Subgroup MSE vs. coverage plots. Compared to baselines (top), our algorithms (bottom) (a) show a consistent trend of decreasing MSE with decrease in coverage for both subgroups, (b) achieve better minority MSE for fixed coverage, i.e., a smaller AUC for minority subgroup (red), (c) achieve comparable majority MSE for fixed coverage, i.e., a comparable AUC for majority subgroup (green), (d) reduce gap between the subgroup MSE curves, i.e., a smaller AUADC.

of 0.8/0.2. More experimental details can be found in Appendix C.

**Comparison in terms of selective regression.** To compare different algorithms in terms of how well they perform selective regression (i.e., without fairness), we look at area under MSE vs. coverage curve (AUC), which encapsulates performance across different coverage (Franc & Prusa, 2019; Lee et al., 2021). We provide the results in Table 2 (smaller AUC indicates better performance) and observe that our algorithms are competitive (if not better) than baselines. We provide MSE vs. coverage curves in Appendix C.

**Comparison in terms of fairness.** To compare different algorithms in terms of *fair* selective regression, we look at subgroup MSE vs coverage curves. for the Insurance dataset, we show these curves for Baseline 2 in Figure 2a and Algorithm 2 in Figure 2b. for the Crime dataset, we show these curves for Baseline 1 in Figure 2c and Algorithm 1 in Figure 2d. See Appendix C for remaining set of curves.

For the Insurance dataset, we see that subgroup MSE for Baseline 2 increases with decrease in coverage for both majority and minority subgroups (Figure 2a). In contrast, the subgroup MSE for Algorithm 2 tends to decrease with a decrease in coverage for both subgroups (Figure 2b). For the Crime dataset, we see that the subgroup

MSE for Baseline 1 as well as Algorithm 1 tends to decrease with a decrease in coverage for both subgroups (Figure 2c and Figure 2d). However, for a particular coverage, Algorithm 1 achieves a better MSE for the minority subgroup, a comparable MSE for the majority subgroup, and reduces the gap between the subgroup curves than Baseline 1.

As hinted above, in addition to ensuring *monotonic selective risk*, one may wish to (a) achieve a better performance focusing solely on the majority subgroup, (b) achieve a better performance focusing solely on the minority subgroup, and (c) reduce the gap between the minority subgroup MSE and the majority subgroup MSE across all thresholds. These aspects could be quantitatively captured by looking at (a) the area under the majority MSE vs. coverage curve, i.e., AUC ( $D = 0$ ), (b) the area under the minority MSE vs. coverage curve, i.e., AUC ( $D = 1$ ), and (c) the area under the absolute difference of the subgroup MSE vs coverage curves (AUADC) (Franc & Prusa, 2019; Lee et al., 2021) respectively. We provide these results in Table 2 and observe that our algorithms outperform the baselines across datasets in terms of AUC ( $D = 1$ ) and AUADC while being comparable in terms of AUC ( $D = 0$ ).

**Application to Causal Inference.** We provide an application of our work to fair-treatment effect estimation.



Table 2: AUC, AUC (D=0), AUC (D=1), and AUADC (averaged across 5 runs) for all algorithms and all datasets. Smaller values are better. B1, B2, A1, and A2 refer to Baseline 1, Baseline 2, Algorithm 1, and Algorithm 2. See Appendix C for standard deviations.

Dataset	Algo-rithm	AUC	AUC (D=0)	AUC (D=1)	AUADC
Insurance	B1	0.0371	0.0342	0.0442	0.0069
	A1	0.0195	0.0207	0.0167	0.0052
	B2	0.0142	0.0129	0.0175	0.0079
	A2	<b>0.0099</b>	<b>0.0087</b>	<b>0.0120</b>	<b>0.0051</b>
Crime	B1	<b>0.0075</b>	<b>0.0040</b>	0.0345	0.0309
	A1	0.0079	0.0045	<b>0.0296</b>	0.0298
	B2	0.0101	0.0060	0.0442	0.0272
	A2	0.0117	0.0082	0.0375	<b>0.0257</b>
IHDP (Treatment)	B1	0.3053	<b>0.2000</b>	0.3509	0.2266
	A1	<b>0.2435</b>	0.2024	<b>0.2849</b>	<b>0.2034</b>
IHDP (Control)	B1	0.2041	<b>0.2144</b>	0.1983	0.0495
	A1	<b>0.2017</b>	0.2169	<b>0.1877</b>	<b>0.0398</b>

Treatment effect estimation has been studied under the paradigm of *prediction with reject option* (Jesson et al., 2020), and we explore an additional dimension to this, i.e., the behavior across sub-populations. We follow the standard approach (Künzel et al., 2019) of viewing the dataset as two distinct datasets — one corresponding to treatment group and the other corresponding to control group — and apply our framework to these distinct datasets. We focus only on Algorithm 1 (Hill, 2011) and compare with Baseline 1 since the simulated target for the IHDP dataset perfectly fits the conditional Gaussian assumption (Hill, 2011). For treatment group, we provide the subgroup MSE vs coverage behavior for Baseline 1 in Figure 2e and Algorithm 1 in Figure 2f. We provide corresponding curves for the control group in Appendix C. Compared to Baseline 1 (Figure 2e), Algorithm 1 (Figure 2f) (a) shows a general trend of decreasing MSE with decrease in coverage for both subgroups and (b) achieves a better minority MSE for a fixed coverage. Table 2 suggests that our algorithm outperforms the baseline in terms of AUC (D = 1) and AUADC while being comparable in AUC (D = 0).

## 7. Concluding Remarks

We proposed a new fairness criterion, *monotonic selective risk* for selective regression, which requires the performance of each subgroup to improve with a decrease in coverage. We provided two conditions for the feature representation (sufficiency and calibrated for mean and variance) under which the proposed fairness criterion is met. We presented algorithms to enforce these conditions and demonstrated

mitigation of disparity in the performances across groups for three real-world datasets.

Monotonic selective risk is one criteria for fairness in selective regression. Developing and investigating other such criteria and understanding their relationship with monotonic selective risk is an important question for future research.

## Acknowledgements

We thank the anonymous reviewers for their comments and suggestions. This work was supported, in part, by the MIT-IBM Watson AI Lab, and its member companies Boston Scientific, Samsung, and Wells Fargo; and NSF under Grant No. CCF-1717610.

## References

- Agarwal, A., Beygelzimer, A., Dudík, M., Langford, J., and Wallach, H. A reductions approach to fair classification. In *International Conference on Machine Learning*, pp. 60–69. PMLR, 2018.
- Agarwal, A., Dudík, M., and Wu, Z. S. Fair regression: Quantitative definitions and reduction-based algorithms. In *International Conference on Machine Learning*, pp. 120–129. PMLR, 2019.
- Arjovsky, M., Bottou, L., Gulrajani, I., and Lopez-Paz, D. Invariant risk minimization. *arXiv preprint arXiv:1907.02893*, 2019.
- Bartlett, P. L. and Wegkamp, M. H. Classification with a reject option using a hinge loss. *Journal of Machine Learning Research*, 9(8), 2008.
- Bellamy, R. K., Dey, K., Hind, M., Hoffman, S. C., Houde, S., Kannan, K., Lohia, P., Martino, J., Mehta, S., Mojsilovic, A., et al. AI Fairness 360: An extensible toolkit for detecting, understanding, and mitigating unwanted algorithmic bias. *arXiv preprint arXiv:1810.01943*, 2018.
- Bengio, Y., Courville, A., and Vincent, P. Representation learning: A review and new perspectives. *IEEE transactions on pattern analysis and machine intelligence*, 35(8):1798–1828, 2013.
- Berk, R., Heidari, H., Jabbari, S., Joseph, M., Kearns, M., Morgenstern, J., Neel, S., and Roth, A. A convex framework for fair regression. *arXiv preprint arXiv:1706.02409*, 2017.
- Calders, T., Karim, A., Kamiran, F., Ali, W., and Zhang, X. Controlling attribute effect in linear regression. In *2013 IEEE 13th international conference on data mining*, pp. 71–80. IEEE, 2013.
- Calmon, F., Wei, D., Vinzamuri, B., Ramamurthy, K. N., and Varshney, K. R. Optimized pre-processing for discrimination prevention. In *Advances in Neural Information Processing Systems*, pp. 3992–4001, 2017.
- Castelnovo, A., Crupi, R., Greco, G., Regoli, D., Penco, I. G., and Cosentini, A. C. A clarification of the nuances in the fairness metrics landscape. *Scientific Reports*, 12(1):1–21, 2022.
- Chi, J., Tian, Y., Gordon, G. J., and Zhao, H. Understanding and mitigating accuracy disparity in regression. *arXiv preprint arXiv:2102.12013*, 2021.
- Chow, C. On optimum recognition error and reject tradeoff. *IEEE Transactions on information theory*, 16(1):41–46, 1970.
- Chow, C.-K. An optimum character recognition system using decision functions. *IRE Transactions on Electronic Computers*, (4):247–254, 1957.
- Chzhen, E. and Schreuder, N. An example of prediction which complies with demographic parity and equalizes group-wise risks in the context of regression. *arXiv preprint arXiv:2011.07158*, 2020.
- Chzhen, E., Denis, C., Hebiri, M., Oneto, L., and Pontil, M. Fair regression via plug-in estimator and recalibration with statistical guarantees. In *NeurIPS 2020-34th Conference on Neural Information Processing Systems*, 2020.
- Cleary, T. A. Test bias: Validity of the scholastic aptitude test for negro and white students in integrated colleges. *ETS Research Bulletin Series*, 1966(2):i–23, 1966.
- Corbett-Davies, S., Pierson, E., Feller, A., Goel, S., and Huq, A. Algorithmic decision making and the cost of fairness. In *Proceedings of the 23rd acm sigkdd international conference on knowledge discovery and data mining*, pp. 797–806, 2017.
- Creager, E., Jacobsen, J.-H., and Zemel, R. Environment inference for invariant learning. In *International Conference on Machine Learning*, pp. 2189–2200. PMLR, 2021.
- Dorie, V. Npci: Non-parametrics for causal inference. 2016. URL <https://github.com/vdorie/npci>.
- Fitzsimons, J., Ali, A., Osborne, M., and Roberts, S. Equality constrained decision trees: For the algorithmic enforcement of group fairness. *arXiv preprint arXiv:1810.05041*, 2018.
- Franc, V. and Prusa, D. On discriminative learning of prediction uncertainty. In *International Conference on Machine Learning*, pp. 1963–1971. PMLR, 2019.
- Gal, Y. Uncertainty in deep learning. 2016.
- Geifman, Y. and El-Yaniv, R. Selective classification for deep neural networks. In *Advances in neural information processing systems*, pp. 4878–4887, 2017.
- Geifman, Y. and El-Yaniv, R. Selectivenet: A deep neural network with an integrated reject option. In *International Conference on Machine Learning*, pp. 2151–2159. PMLR, 2019.
- Gölz, P., Kahng, A., and Procaccia, A. D. Paradoxes in fair machine learning. In *Advances in Neural Information Processing Systems*, pp. 8340–8350, 2019.

- Hall, P. and Carroll, R. J. Variance function estimation in regression: The effect of estimating the mean. *Journal of the Royal Statistical Society: Series B (Methodological)*, 51(1):3–14, 1989.
- Hardt, M., Price, E., and Srebro, N. Equality of opportunity in supervised learning. In *Advances in Neural Information Processing Systems 29*, pp. 3315–3323, Barcelona, Spain, December 2016.
- Hellman, M. E. The nearest neighbor classification rule with a reject option. *IEEE Transactions on Systems Science and Cybernetics*, 6(3):179–185, 1970.
- Herbei, R. and Wegkamp, M. H. Classification with reject option. *The Canadian Journal of Statistics/La Revue Canadienne de Statistique*, pp. 709–721, 2006.
- Hill, J. L. Bayesian nonparametric modeling for causal inference. *Journal of Computational and Graphical Statistics*, 20(1):217–240, 2011.
- Jesson, A., Mindermann, S., Shalit, U., and Gal, Y. Identifying causal-effect inference failure with uncertainty-aware models. *Advances in Neural Information Processing Systems*, 33, 2020.
- Jiang, W., Zhao, Y., and Wang, Z. Risk-controlled selective prediction for regression deep neural network models. In *2020 International Joint Conference on Neural Networks (IJCNN)*, pp. 1–8. IEEE, 2020.
- Jones, E., Sagawa, S., Koh, P. W., Kumar, A., and Liang, P. Selective classification can magnify disparities across groups. *arXiv preprint arXiv:2010.14134*, 2020.
- Kamishima, T., Akaho, S., and Sakuma, J. Fairness-aware learning through regularization approach. In *2011 IEEE 11th International Conference on Data Mining Workshops*, pp. 643–650. IEEE, 2011.
- Komiyama, J. and Shimao, H. Two-stage algorithm for fairness-aware machine learning. *arXiv preprint arXiv:1710.04924*, 2017.
- Komiyama, J., Takeda, A., Honda, J., and Shimao, H. Nonconvex optimization for regression with fairness constraints. In *International conference on machine learning*, pp. 2737–2746. PMLR, 2018.
- Krueger, D., Caballero, E., Jacobsen, J.-H., Zhang, A., Binas, J., Zhang, D., Le Priol, R., and Courville, A. Out-of-distribution generalization via risk extrapolation (rex). In *International Conference on Machine Learning*, pp. 5815–5826. PMLR, 2021.
- Künzel, S. R., Sekhon, J. S., Bickel, P. J., and Yu, B. Metalearners for estimating heterogeneous treatment effects using machine learning. *Proceedings of the national academy of sciences*, 116(10):4156–4165, 2019.
- Lantz, B. *Machine learning with R: expert techniques for predictive modeling*. Packt publishing ltd, 2019.
- Lee, J., Bu, Y., Sattigeri, P., Panda, R., Wornell, G., Karlinsky, L., and Feris, R. A maximal correlation approach to imposing fairness in machine learning. *arXiv preprint arXiv:2012.15259*, 2020.
- Lee, J. K., Bu, Y., Rajan, D., Sattigeri, P., Panda, R., Das, S., and Wornell, G. W. Fair selective classification via sufficiency. In *International Conference on Machine Learning*, pp. 6076–6086. PMLR, 2021.
- Lei, J. Classification with confidence. *Biometrika*, 101(4): 755–769, 2014.
- Liu, L. T., Simchowitz, M., and Hardt, M. The implicit fairness criterion of unconstrained learning. In *International Conference on Machine Learning*, pp. 4051–4060. PMLR, 2019.
- Louizos, C., Swersky, K., Li, Y., Welling, M., and Zemel, R. The variational fair autoencoder. *arXiv preprint arXiv:1511.00830*, 2015.
- Mary, J., Calauzenes, C., and El Karoui, N. Fairness-aware learning for continuous attributes and treatments. In *International Conference on Machine Learning*, pp. 4382–4391, 2019.
- Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K., and Galstyan, A. A survey on bias and fairness in machine learning. *ACM Computing Surveys (CSUR)*, 54(6):1–35, 2021.
- Nabi, R., Malinsky, D., and Shpitser, I. Learning optimal fair policies. In *International Conference on Machine Learning*, pp. 4674–4682. PMLR, 2019.
- Nadeem, M. S. A., Zucker, J.-D., and Hanczar, B. Accuracy-rejection curves (arcs) for comparing classification methods with a reject option. In *Machine Learning in Systems Biology*, pp. 65–81. PMLR, 2009.
- Oneto, L., Donini, M., and Pontil, M. General fair empirical risk minimization. In *2020 International Joint Conference on Neural Networks (IJCNN)*, pp. 1–8. IEEE, 2020.
- Pérez-Suay, A., Laparra, V., Mateo-García, G., Muñoz-Marí, J., Gómez-Chova, L., and Camps-Valls, G. Fair kernel learning. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pp. 339–355. Springer, 2017.

- Pleiss, G., Raghavan, M., Wu, F., Kleinberg, J., and Weinberger, K. Q. On fairness and calibration. In *Advances in Neural Information Processing Systems*, pp. 5680–5689, 2017.
- Raff, E., Sylvester, J., and Mills, S. Fair forests: Regularized tree induction to minimize model bias. In *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society*, pp. 243–250, 2018.
- Redmond, M. and Baveja, A. A data-driven software tool for enabling cooperative information sharing among police departments. *European Journal of Operational Research*, 141(3):660–678, 2002.
- Schreuder, N. and Chzhen, E. Classification with abstention but without disparities. *arXiv preprint arXiv:2102.12258*, 2021.
- Selbst, A. D., Boyd, D., Friedler, S. A., Venkatasubramanian, S., and Vertesi, J. Fairness and abstraction in sociotechnical systems. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*, pp. 59–68, 2019.
- Verma, S. and Rubin, J. Fairness definitions explained. In *2018 IEEE/ACM International Workshop on Software Fairness (Fairware)*, pp. 1–7. IEEE, 2018.
- Wiener, Y. and El-Yaniv, R. Pointwise tracking the optimal regression function. In *Advances in Neural Information Processing Systems*, pp. 2042–2050, 2012.
- Williamson, R. and Menon, A. Fairness risk measures. In *International Conference on Machine Learning*, pp. 6786–6797. PMLR, 2019.
- Zafar, M. B., Valera, I., Rodriguez, M. G., and Gummadi, K. P. Fairness constraints: Mechanisms for fair classification. In *Artificial Intelligence and Statistics*, pp. 962–970. PMLR, 2017.
- Zaoui, A., Denis, C., and Hebiri, M. Regression with reject option and application to knn. *Advances in Neural Information Processing Systems*, 33:20073–20082, 2020.
- Zemel, R., Wu, Y., Swersky, K., Pitassi, T., and Dwork, C. Learning fair representations. In *International Conference on Machine Learning*, pp. 325–333, 2013.

## Appendix

**Organization.** The Appendix is organized as follows. In Appendix A, we provide the proof of Theorem 4.1. In Appendix B, we provide more details for imposing calibration for mean and variance via the residual-based neural network as well as provide the pseudo-code, i.e., Algorithm 2. In Appendix C, we provide more experimental details and results.

### A. Proof of Theorem 4.1

We restate the Theorem below and then provide a proof.

**Theorem 4.1.** *Suppose the representation  $\Phi(X)$  is sufficient i.e.,  $Y \perp D | \Phi(X)$ . Let  $f(\Phi(X)) = \mathbb{E}[Y | \Phi(X)]$  and  $g(\Phi(X)) = \text{Var}[Y | \Phi(X)]$ . Then, for any  $d \in \mathcal{D}$  and any  $\tau < \tau'$ , we have  $\text{MSE}(f, g, \tau, d) < \text{MSE}(f, g, \tau', d)$ .*

*Proof.* First, let us simplify the expression for  $\text{MSE}(f, g, \tau, d)$ . We have

$$\begin{aligned} \text{MSE}(f, g, \tau, d) &= \frac{\mathbb{E}[(Y - f(\Phi(X)))^2 \cdot \mathbf{1}(g(\Phi(X)) \leq \tau) | D = d]}{\mathbb{E}[\mathbf{1}(g(\Phi(X)) \leq \tau) | D = d]} \\ &= \frac{\mathbb{E}[(Y - f(\Phi(X)))^2 \cdot \mathbf{1}(g(\Phi(X)) \leq \tau) | D = d]}{\mathbb{P}(g(\Phi(X)) \leq \tau | D = d)} \end{aligned} \quad (10)$$

Let  $c_{\tau,d} \triangleq \mathbb{P}(g(\Phi(X)) \leq \tau | D = d)$ . Using (10), we have

$$\begin{aligned} &\text{MSE}(f, g, \tau, d) - \text{MSE}(f, g, \tau', d) \\ &= \frac{1}{c_{\tau,d}} \mathbb{E}[(Y - f(\Phi(X)))^2 \cdot \mathbf{1}(g(\Phi(X)) \leq \tau) | D = d] - \frac{1}{c_{\tau',d}} \mathbb{E}[(Y - f(\Phi(X)))^2 \cdot \mathbf{1}(g(\Phi(X)) \leq \tau') | D = d] \\ &= \left( \frac{1}{c_{\tau,d}} - \frac{1}{c_{\tau',d}} \right) \mathbb{E}[(Y - f(\Phi(X)))^2 \cdot \mathbf{1}(g(\Phi(X)) \leq \tau) | D = d] \end{aligned} \quad (11)$$

$$- \frac{1}{c_{\tau',d}} \mathbb{E}[(Y - f(\Phi(X)))^2 \cdot \mathbf{1}(\tau < g(\Phi(X)) \leq \tau') | D = d] \quad (12)$$

Now, let us upper bound  $\mathbb{E}[(Y - f(\Phi(X)))^2 \cdot \mathbf{1}(g(\Phi(X)) \leq \tau) | D = d]$ . We have

$$\begin{aligned} &\mathbb{E}[(Y - f(\Phi(X)))^2 \cdot \mathbf{1}(g(\Phi(X)) \leq \tau) | D = d] \\ &\stackrel{(a)}{=} \mathbb{E}_{\Phi(X)|D}[\mathbf{1}(g(\Phi(X)) \leq \tau) \cdot \mathbb{E}_{Y|\Phi(X),D}[(Y - f(\Phi(X)))^2 | \Phi(X), D = d] | D = d] \\ &\stackrel{(b)}{=} \mathbb{E}_{\Phi(X)|D}[\mathbf{1}(g(\Phi(X)) \leq \tau) \cdot \mathbb{E}_{Y|\Phi(X),D}[(Y - \mathbb{E}[Y | \Phi(X)])^2 | \Phi(X), D = d] | D = d] \\ &\stackrel{(c)}{=} \mathbb{E}_{\Phi(X)|D}[\mathbf{1}(g(\Phi(X)) \leq \tau) \cdot \mathbb{E}_{Y|\Phi(X),D}[(Y - \mathbb{E}[Y | \Phi(X), D = d])^2 | \Phi(X), D = d] | D = d] \\ &\stackrel{(d)}{=} \mathbb{E}_{\Phi(X)|D}[\mathbf{1}(g(\Phi(X)) \leq \tau) \cdot \text{Var}[Y | \Phi(X), D = d] | D = d] \\ &\stackrel{(e)}{=} \mathbb{E}_{\Phi(X)|D}[\mathbf{1}(g(\Phi(X)) \leq \tau) \cdot \text{Var}[Y | \Phi(X)] | D = d] \\ &\stackrel{(f)}{=} \mathbb{E}_{\Phi(X)|D}[\mathbf{1}(g(\Phi(X)) \leq \tau) \cdot g(\Phi(X)) | D = d] \\ &\leq \tau \mathbb{E}_{\Phi(X)|D}[\mathbf{1}(g(\Phi(X)) \leq \tau) | D = d] = \tau \mathbb{P}(g(\Phi(X)) \leq \tau | D = d) = \tau c_{\tau,d} \end{aligned} \quad (13)$$

where (a) follows from the definition of conditional expectation and because  $\mathbf{1}(g(\Phi(X)) \leq \tau)$  is a constant conditioned on  $\Phi(X)$ , (b) follows because  $f(\Phi(X)) = \mathbb{E}[Y | \Phi(X)]$ , (c) follows because  $Y \perp D | \Phi(X) \implies \mathbb{E}[Y | \Phi(X)] = \mathbb{E}[Y | \Phi(X), D = d]$ , (d) follows from the definition of conditional variance, (e) follows because  $Y \perp D | \Phi(X) \implies \text{Var}[Y | \Phi(X)] = \text{Var}[Y | \Phi(X), D = d]$ , and (f) follows because  $g(\Phi(X)) = \text{Var}[Y | \Phi(X)]$ .

Now, let us lower bound  $\mathbb{E}[(Y - f(\Phi(X)))^2 \cdot \mathbf{1}(\tau < g(\Phi(X)) \leq \tau') | D = d]$ . Similar to above, we have

$$\begin{aligned} &\mathbb{E}[(Y - f(\Phi(X)))^2 \cdot \mathbf{1}(\tau < g(\Phi(X)) \leq \tau') | D = d] = \mathbb{E}_{\Phi(X)|D}[\mathbf{1}(\tau < g(\Phi(X)) \leq \tau') \cdot g(\Phi(X)) | D = d] \\ &> \tau \mathbb{E}_{\Phi(X)|D}[\mathbf{1}(\tau < g(\Phi(X)) \leq \tau') | D = d] \\ &= \tau (\mathbb{P}(g(\Phi(X)) \leq \tau' | D = d) - \mathbb{P}(g(\Phi(X)) \leq \tau | D = d)) \\ &= \tau (c_{\tau',d} - c_{\tau,d}) \end{aligned} \quad (14)$$

Plugging in (13) and (14) in (11) and (12), we have

$$\text{MSE}(f, g, \tau, d) - \text{MSE}(f, g, \tau', d) < \left( \frac{1}{c_{\tau, d}} - \frac{1}{c_{\tau', d}} \right) \cdot \tau c_{\tau, d} - \frac{1}{c_{\tau', d}} \tau \cdot (c_{\tau', d} - c_{\tau, d}) = 0$$

□

## B. More Details to Impose Calibration for Mean and Variance

In this section, we provide more details for imposing calibration for mean and variance via the residual-based neural network as well as provide the pseudo-code. As described in Section 5.2, in a residual-based neural network, once the mean-prediction network  $f$  is trained, the residuals i.e.,  $r_i = (y_i - f(\Phi_1(x_i); \theta_f))^2$  are used to train the variance-prediction network  $g$  by minimizing:

$$L_{S2}(\Phi_2, \theta_g) \triangleq \sum_{i=1}^n (r_i - g(\Phi_2(x_i); \theta_g))^2.$$

The feature representation  $\Phi_2$  is parameterized by  $\theta_{\Phi_2}$  and the variance-prediction network is supposed to learn the parameters  $\theta_{\Phi_2}$  and  $\theta_g$ .

To impose calibration under variance, we construct a contrastive loss similar to (7). Then, the regularizer can be written as

$$L_{R2}(\Phi_2) \triangleq \sum_{i=1}^n \left( (r_i - g(\Phi_2(x_i); w_g^{(\tilde{d}_i)}))^2 - (r_i - g(\Phi_2(x_i); w_g^{(d_i)}))^2 \right),$$

where  $\tilde{d}_i$  are drawn i.i.d. from the marginal distribution  $P_D$ , and for  $d \in \mathcal{D}$ ,

$$w_g^{(d)} = \arg \min_w \sum_{i: d_i=d} (r_i - g(\Phi_2(x_i); w))^2.$$

Summarizing, the overall objective for variance-prediction is

$$\min_{\theta_g, \Phi_2} L_{S2}(\Phi_2, \theta_g) + \lambda_2 L_{R2}(\Phi_2).$$

We provide a pseudo-code in Algorithm 2 where

$$L_{d1}(w_f) \triangleq \sum_{i: d_i=d} (y_i - f(\Phi_1(x_i); w_f))^2,$$

$$L_{d2}(w_g) \triangleq \sum_{i: d_i=d} (r_i - g(\Phi_2(x_i); w_g))^2.$$

## C. Additional experimental results

In this section, we provide more experimental details and results. We start by providing those experimental details that remain the same across the datasets. Next, we provide details that are specific to each dataset, i.e., Insurance, Crime, and IHDP. Finally, we provide more experimental results and some discussion.

### C.1. Experimental Details

In all of our experiments, we use two-layer neural networks. For all hidden layers, we use the *selu* activation function. For the output layer, we use a non-linear activation function only for the variance-prediction network associated with Algorithm 2 to ensure that the predictions of variance are non-negative. In particular, we use the *soft-plus* activation function for the variance-prediction network associated with Algorithm 2. In our implementation of Algorithm 1, we predict log-variance instead of variance and therefore stick to *linear* activation function.

We train all our neural networks with the Adam optimizer, a batch size of 128, and over 40 epochs. Further, we use a step learning rate scheduler with an initial learning rate of  $5 \times 10^{-3}$  and decay it by a factor of half after every two epochs. Finally, as described in Section 6, we set the regularizer  $\lambda = 1$  for all our experiments after observing that the performance of our algorithms is agnostic to the choice of  $\lambda$  as long as it is in a reasonable range, i.e.,  $\lambda \in [0.5, 3]$ .

**Algorithm 2** Residual-based neural network with calibration-based regularizer

---

**Input:** training samples  $\{(x_i, y_i, d_i)\}_{i=1}^n$ , regularizers  $\lambda_1$  and  $\lambda_2$   
**Draw:**  $\{\tilde{d}_1, \dots, \tilde{d}_n\}$  drawn i.i.d. from  $\hat{\mathbb{P}}_D$   
**Initialize:**  $\theta_f, \theta_g, \theta_{\Phi_1}, \theta_{\Phi_2}, w_f^{(d)}$ , and  $w_g^{(d)}$  with pre-trained models  
**Initialize:**  $n_d =$  number of samples in group  $d \forall d \in \mathcal{D}$   
**for** each training iteration **do**  
    **for** each batch **do**  
        **for**  $d = 1, \dots, |\mathcal{D}|$  **do**  
            # update subgroup-specific mean predictor  
             $w_f^{(d)} \leftarrow w_f^{(d)} - \frac{1}{n_d} \eta_f \nabla_{w_f} L_{d1}(w_f)$   
        **end for**  
    **end for**  
    **for** each batch **do**  
        # update feature extractor for mean predictor  
         $\theta_{\Phi_1} \leftarrow \theta_{\Phi_1} - \frac{1}{n} \eta \nabla_{\theta_{\Phi_1}} (L_{S1}(\Phi_1, \theta_f) + \lambda_1 L_{R1}(\Phi_1))$   
        # update mean predictor  
         $\theta_f \leftarrow \theta_f - \frac{1}{n} \eta \nabla_{\theta_f} L_{S1}(\Phi_1, \theta_f)$   
    **end for**  
**end for**  
Compute the residuals:  $r_i = (y_i - f(\Phi_1(x_i); \theta_f))^2$   
**for** each training iteration **do**  
    **for** each batch **do**  
        **for**  $d = 1, \dots, |\mathcal{D}|$  **do**  
            # update subgroup-specific variance predictor  
             $w_g^{(d)} \leftarrow w_g^{(d)} - \frac{1}{n_d} \eta_g \nabla_{w_g} L_{d2}(w_g)$   
        **end for**  
    **end for**  
    **for** each batch **do**  
        # update feature extractor for variance predictor  
         $\theta_{\Phi_2} \leftarrow \theta_{\Phi_2} - \frac{1}{n} \eta \nabla_{\theta_{\Phi_2}} (L_{S2}(\Phi_2, \theta_g) + \lambda_2 L_{R2}(\Phi_2))$   
        # update variance predictor  
         $\theta_g \leftarrow \theta_g - \frac{1}{n} \eta \nabla_{\theta_g} L_{S2}(\Phi_2, \theta_g)$   
    **end for**  
**end for**

---

### C.2. Insurance

The Insurance dataset<sup>5</sup> is a semi-synthetic dataset that was created using demographic statistics from the U.S. Census Bureau and approximately reflects real-world conditions. A few features in this dataset include the BMI, number of children, age, etc. We remove the sensitive attribute from the set of input features to preprocess the data. To reflect the real-world scenarios where the accuracy disparity is significant due to the small and imbalanced dataset, similar to (Chi et al., 2021), we randomly drop 50% of examples with  $D = 1$ . Further, we normalize the output annual medical expenses and the features: age and BMI. We use 3 neurons in the hidden layer for this dataset.

### C.3. Communities and Crime

The Communities and Crime dataset<sup>6</sup> contains socio-economic information of communities in the U.S. and their crime rates. A few features in this dataset include population for community, mean people per household, percentage of the population that is white, per capita income, number of police cars, etc. We remove the non-predictive attributes and the sensitive attribute from the set of input features during preprocessing. All attributes in the dataset have been curated and normalized

<sup>5</sup><https://github.com/stedy/Machine-Learning-with-R-datasets/blob/master/insurance.csv>

<sup>6</sup><https://archive.ics.uci.edu/ml/datasets/communities+and+crime>

to  $[0, 1]$ , so we do not perform any additional normalization. Finally, we replace the missing values with the mean values of the corresponding attributes similar to (Chi et al., 2021). We use 50 neurons in the hidden layer for this dataset.

#### C.4. IHDP

The IHDP dataset<sup>7</sup> is generated based on a randomized control trial targeting low-birth-weight, premature infants. The 25 features measure various aspects about the children and their mothers, e.g., child’s birth weight, child’s gender, mother’s age, mother’s education, an indicator for maternal alcohol consumption during pregnancy, etc. We remove the sensitive attribute from the set of input features to preprocess the data. Further, we normalize the output cognitive test score and the features: child’s birth weight, child’s head circumference at birth, number of weeks pre-term that the child was born, birth order, neo-natal health index, and mom’s age when she gave birth to the child. Following the norm in the causal inference community, a biased subset of the treated group is removed to create an imbalance leaving 139 samples in the treatment group and 608 samples in the control group. The target is typically simulated using the setting “A” of the NPCI package (Dorie, 2016). We use 20 neurons in the hidden layer for this dataset.

#### C.5. Standard deviations

In Table 3 below, we provide the standard deviations associated with AUC, AUC (D = 0), AUC (D = 1), and AUADC whose means were provided in Table 2 in Section 6.

Table 3: Mean  $\pm$  standard deviation (averaged across 5 runs) for AUC, AUC (D = 0), AUC (D = 1), and AUADC for all the algorithms and all the datasets.

Dataset	Algorithm	AUC	AUC (D = 0)	AUC (D = 1)	AUADC
Insurance	B1 i.e., Baseline 1	$0.0371 \pm 0.0255$	$0.0342 \pm 0.0197$	$0.0442 \pm 0.0218$	$0.0069 \pm 0.0050$
	A1 i.e., Algorithm 1	$0.0195 \pm 0.0059$	$0.0207 \pm 0.0050$	$0.0167 \pm 0.0075$	$0.0052 \pm 0.0031$
	B2 i.e., Baseline 2	$0.0142 \pm 0.0052$	$0.0129 \pm 0.0042$	$0.0175 \pm 0.0026$	$0.0079 \pm 0.0041$
	A2 i.e., Algorithm 2	$0.0099 \pm 0.0006$	$0.0087 \pm 0.0004$	$0.0120 \pm 0.0011$	$0.0051 \pm 0.0018$
Crime	B1 i.e., Baseline 1	$0.0075 \pm 0.0002$	$0.0040 \pm 0.0011$	$0.0345 \pm 0.0037$	$0.0309 \pm 0.0008$
	A1 i.e., Algorithm 1	$0.0079 \pm 0.0004$	$0.0045 \pm 0.0013$	$0.0296 \pm 0.0054$	$0.0298 \pm 0.0011$
	B2 i.e., Baseline 2	$0.0101 \pm 0.0019$	$0.0060 \pm 0.0017$	$0.0442 \pm 0.0022$	$0.0272 \pm 0.0013$
	A2 i.e., Algorithm 2	$0.0117 \pm 0.0017$	$0.0082 \pm 0.0012$	$0.0375 \pm 0.0019$	$0.0257 \pm 0.0028$
IHDP (Treatment)	B1 i.e., Baseline 1	$0.3053 \pm 0.0823$	$0.2000 \pm 0.0899$	$0.3509 \pm 0.0811$	$0.2266 \pm 0.0919$
	A1 i.e., Algorithm 1	$0.2435 \pm 0.0823$	$0.2024 \pm 0.0935$	$0.2849 \pm 0.0767$	$0.2034 \pm 0.0925$
IHDP (Control)	B1 i.e., Baseline 1	$0.2041 \pm 0.0138$	$0.2144 \pm 0.0101$	$0.1983 \pm 0.0125$	$0.0495 \pm 0.0053$
	A1 i.e., Algorithm 1	$0.2017 \pm 0.0170$	$0.2169 \pm 0.0133$	$0.1877 \pm 0.0129$	$0.0398 \pm 0.0073$

#### C.6. Overall MSE vs. coverage curves

In Section 6, we compared different algorithms in terms of how well they performed selective regression (i.e., with no consideration of fairness) by looking at the area under MSE vs. coverage curve (AUC). Here, we provide the MSE vs. coverage curves for the Insurance dataset in Figure 3a, the Crime dataset in Figure 3b, the IHDP (control) dataset in Figure 3c, and the IHDP (treatment) dataset in Figure 3d.

For the Insurance dataset, we see that Algorithm 1 and Algorithm 2 perform selective regression better than Baseline 1 and Baseline 2, respectively. This is also evident via the values of AUC in Table 2/3. For the Crime dataset, the MSE decreases with decrease in coverage as expected for all four algorithms. Further, the performances of Baseline 1 and Baseline 2 are slightly better than that of Algorithm 1 and Algorithm 2 respectively. This is also evident via the values of AUC in Table 2/3. For the IHDP (control) dataset and IHDP (treatment) dataset, we see that Algorithm 1 performs selective regression better than Baseline 1. This is also evident via the values of AUC in Table 2/3.

<sup>7</sup><https://github.com/AMLab-Amsterdam/CEVAE/tree/master/datasets/IHDP>



## Selective Regression Under Fairness Criteria

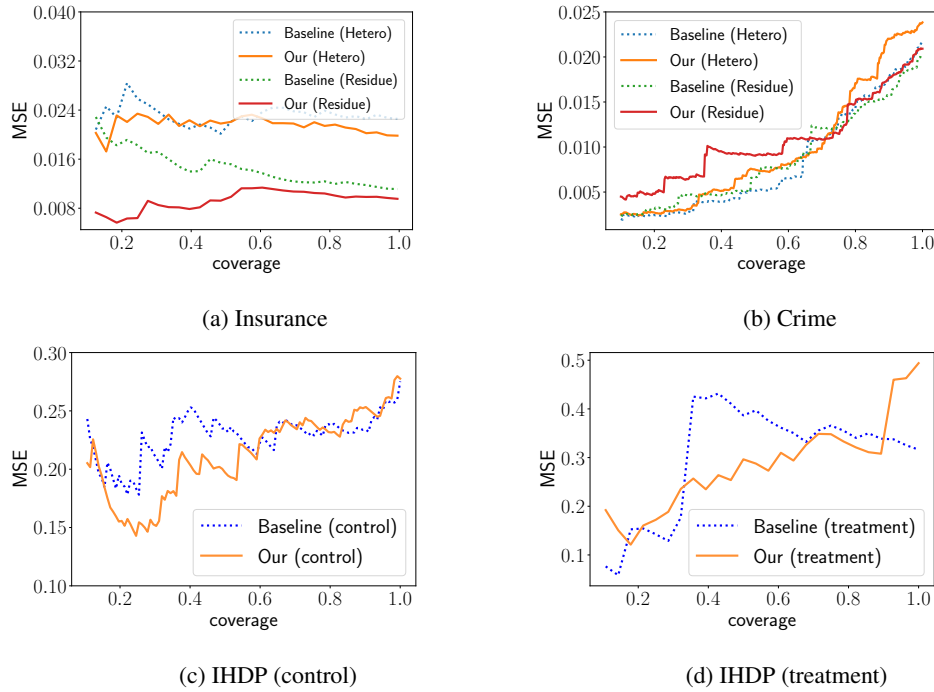


Figure 3: MSE vs. coverage for various datasets

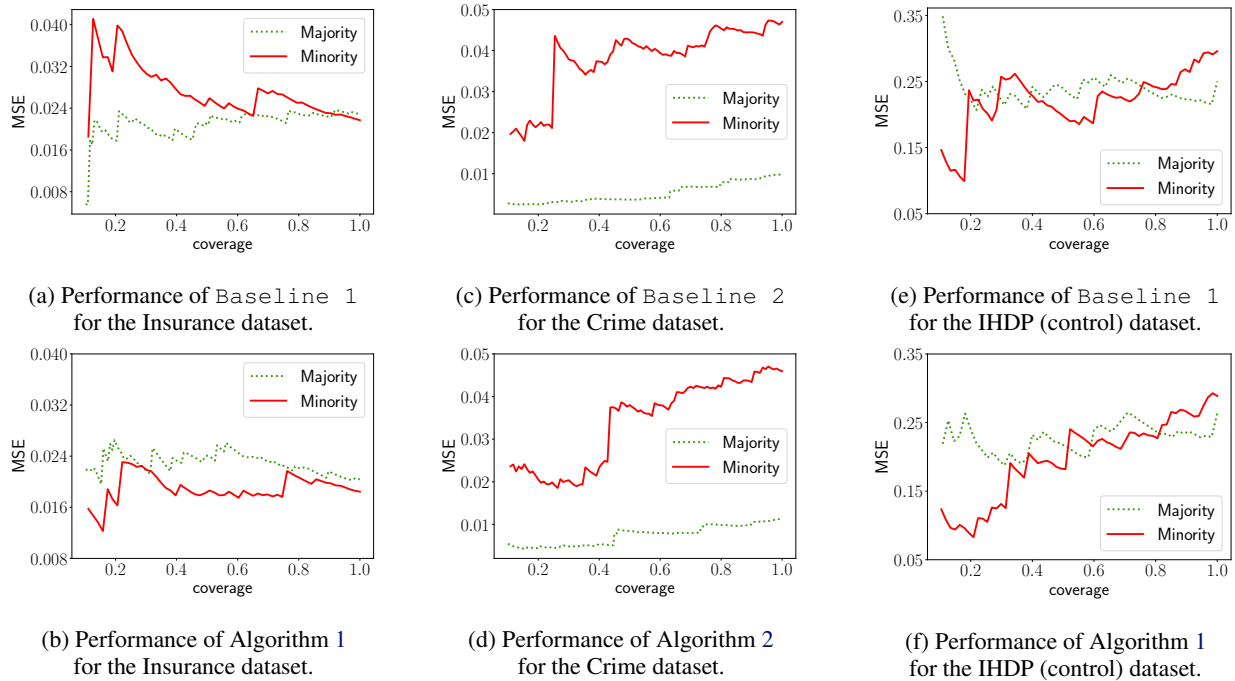


Figure 4: Subgroup MSE vs. coverage plots for various datasets comparing baselines (top) and our algorithms (bottom).

### C.7. Group-specific MSE vs. coverage curves.

In Section 6, we compared the different algorithms in terms of how well they perform fair selective regression by looking at the subgroup MSE vs. coverage curves in addition to AUC, AUC ( $D = 0$ ), AUC ( $D = 1$ ), and AUADC.

More specifically, we looked at Baseline 2 and Algorithm 2 for the Insurance dataset, Baseline 1 and Algorithm 1 for the Crime dataset, and Baseline 1 and Algorithm 1 for the IHDP (treatment) dataset. Here, we show the subgroup MSE vs. coverage curves for Baseline 1 and Algorithm 1 for the Insurance dataset (Figure 4a and 4b), Baseline 2 and Algorithm 2 for the Crime dataset (Figure 4c and 4d), and Baseline 1 and Algorithm 1 for the IHDP (control) dataset (Figure 4e and 4f).

For the Insurance dataset, we see that subgroup MSE for the minority subgroup increases with decrease in coverage for Baseline 1 (Figure 4a) as already described in Section 3.1. In contrast, the subgroup MSE for Algorithm 1 does not increase with decrease in coverage and stays relatively flat (Figure 4b). Further, for a particular coverage, Algorithm 1 achieves a better MSE for the minority subgroup, a comparable MSE for the majority subgroup, and reduces the gap between the subgroup curves than Baseline 1 (see the values of AUC ( $D = 0$ ), AUC ( $D = 1$ ), and AUADC in Table 2/3).

For the Crime dataset, we see that the subgroup MSE for Baseline 2 as well as Algorithm 2 tends to decrease with a decrease in coverage for both subgroups (Figure 4c and Figure 4d). However, for a particular coverage, Algorithm 2 achieves a better MSE for the minority subgroup, a comparable MSE for the majority subgroup, and reduces the gap between the subgroup curves than Baseline 2 (see the values of AUC ( $D = 0$ ), AUC ( $D = 1$ ), and AUADC in Table 2/3).


For the IHDP (control) dataset, we see that the subgroup MSE for Baseline 1 increases with decrease in coverage (Figure 4e). In contrast, the subgroup MSE for Algorithm 1 decreases with decrease in coverage (Figure 4f). Additionally, Algorithm 1 achieves a comparable MSE for the majority subgroup, and reduces the gap between the subgroup curves than Baseline 1 (see the values of AUC ( $D = 0$ ), AUC ( $D = 1$ ), and AUADC in Table 2/3).

### C.8. Group-specific MSE vs. coverage curves with three subgroups

In all of our experiments so far, we focused on the scenario where the sensitive attribute was binary. However, as we now demonstrate, our approach works equally well when the sensitive attribute can take more than two values.

More concretely, we use the Crime dataset to obtain 3 subgroups i.e., the sensitive attribute can take three values. This is possible since, race, the sensitive attribute, is reported as the population percentage of the black in the Crime dataset. We assign (a)  $D = 2$  if the population percentage of the black is more than or equal to 20, (b)  $D = 1$  if the population percentage of the black is less than 20 but more than or equal to 1, and (c)  $D = 0$  otherwise.

We show the performance of Baseline 1 and Algorithm 1 in Figure 5a and 5b respectively. As expected, Algorithm 1 ensures monotonic selective risk unlike Baseline 1 (see  $D = 2$ ).



(a) Performance of Baseline 1 for the Crime dataset.

(b) Performance of Algorithm 1 for the Crime dataset.

Figure 5: Subgroup MSE vs. coverage plots for the Crime dataset with three subgroups