Is My Prediction Arbitrary? Measuring Self-Consistency in Fair Classification

A. Feder Cooper [1]   Katherine Lee [1 2]   Solon Barocas [3 4]   Christopher De Sa [1]   Siddhartha Sen [3]   Baobao Zhang [5]

Abstract

Variance in predictions across different trained models is a significant, under-explored source of error in fair classification. Empirically, the variance on some instances is so large that decisions can be effectively *arbitrary*. To study this problem, we perform a large-scale empirical study and make four overarching contributions: We 1) Define a metric called *self-consistency*, derived from variance, which we use as a proxy for measuring and reducing arbitrariness; 2) Develop an ensembling algorithm that abstains from classification when a prediction would be arbitrary; 3) Conduct the largest to-date empirical study of the role of variance (*vis-a-vis* self-consistency and arbitrariness) in fair classification; and, 4) Release a toolkit that makes the US Home Mortgage Disclosure Act (HMDA) datasets easily usable for future research. Altogether, our empirical results reveal shocking insights about reproducibility. **Most fairness classification benchmarks are *close-to-fair* when taking into account the amount of arbitrariness present in predictions. Subgroup error rates are similar before we even try to apply common fairness interventions.**

1. Introduction

A goal of algorithmic fairness is to develop techniques that measure and mitigate discrimination in automated decision-making systems. In binary fair classification, this often involves training a model to satisfy a chosen *fairness metric*, which typically defines fairness as parity between model error rates for the different demographic subgroups in the dataset (Barocas et al., 2019). However, even if a model's classification decisions satisfy a particular fairness metric, it is not necessarily the case the model is equally confident in each prediction. This becomes clear if we measure
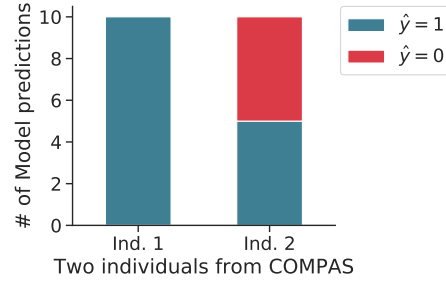

Figure 1. 10 logistic regression models trained on bootstrap replicates shows how models can be very consistent in predictions $\hat{y}$ for some individuals (Ind. 1) and random for others (Ind. 2).

variance between models. In this paper, we train multiple models by sampling subsets from the original dataset, measure the disagreement in predictions between models trained on each subset, and interpret this disagreement as a lack of confidence in the classification decision. Empirically, some instances can yield decisions that exhibit so much disagreement that they are effectively *arbitrary*.

We modify variance to define a metric of *self-consistency*, which we use as a quantitative proxy for measuring the arbitrariness of model outputs. To understand this intuitively, consider the following experiment: We fit 10 logistic regression models using the same learning process, which draws different subsamples of the training set from the COMPAS prison recidivism dataset (Larson et al., 2016; Friedler et al., 2019), and we compare the resulting classifications for two individuals in the test set. Figure 1 shows a difference in the consistency of predictions on Individual 1 vs. Individual 2. The 10 models completely agree to classify Individual 1 as "will recidivate" and disagree completely on whether to classify Individual 2 as "will" or "will not recidivate." If we were to pick one model at random to use in practice, there would be no effect on how Individual 1 is classified; yet, for Individual 2, the prediction is effectively random. The classification of Individual 2 is *arbitrary*. We can interpret Figure 1 to mean that the learning process that produced these predictions is not sufficiently confident to justify assigning Individual 2 *either decision outcome*.

Using just one model could subject individuals to consequential decisions for which the learning process is arbitrary (Citron and Pasquale, 2014; Cooper et al., 2022; Creel and Hellman, 2022; Black et al., 2022b). Further, this arbitrariness can also bring about discrimination if classification

---

[1]Department of Computer Science, Cornell University, Ithaca, NY, USA [2]Google DeepMind, San Francisco, CA, USA [3]Microsoft Research, New York, NY, USA [4]Department of Information Science, Cornell University, Ithaca, NY, USA [5]Maxwell School of Citizenship and Public Affairs, Syracuse University, Syracuse, NY, USA. Correspondence to: A. Feder Cooper <https://cacioepe.pe>.

decisions are *systematically more arbitrary* for certain demographic groups. However, popular fair classification metrics are commonly applied to evaluate the fairness of a *single model* (Hardt et al., 2016; Pleiss et al., 2017; Kleinberg et al., 2017; Chouldechova, 2017; Calders et al., 2009). Using a single model can mask the arbitrariness of its predictions. Instead, to reveal arbitrariness, we must examine *distributions of possible models for a given learning process*, like the one visualized in Figure 1. With this shift in frame, we ask:

> *What is the empirical role of arbitrariness in real-world fairness classification tasks?*

To study this question, we make four contributions: We

1. **Quantify arbitrariness.** We formalize a metric called *self-consistency*, derived from statistical variance, which we use as a quantitative proxy for arbitrariness. Self-consistency is a simple yet powerful tool for empirical analyses of fair classification (Section 3).

2. **Ensemble to improve self-consistency.** We extend bagging to allow for abstaining from prediction on instances for which self-consistency is low. This improves overall self-consistency and accuracy (Section 4).

3. **Perform a comprehensive empirical study of variance in fair classification.** We conduct the largest-to-date systematic study of the role of variance in empirical fair classification, through the lens of self-consistency and its relationship to arbitrariness. Our results reveal shocking insights concerning empirical reproducibility (Section 5). **Notably, most fairness classification benchmarks are close-to-fair when taking into account the amount of arbitrariness present in predictions. Subgroup error rates are similar *before* we even try to apply common pre-, in-, or post-processing fairness interventions.** In turn, our study raises crucial questions about future directions for research in fair classification (Section 6).

4. **Release a large-scale fairness dataset package.** We observe that variance, particularly in small datasets, can undermine the reliability of conclusions about fairness interventions. We therefore open-source a package that makes the large-scale US Home Mortgage Disclosure Act datasets (HMDA) easily usable for future research.

2. Preliminaries on Fair Binary Classification

To analyze arbitrariness in the context of algorithmic fairness, we first need to establish our background definitions for fair binary classification. **This material is likely familiar to most readers. We nevertheless highlight particular details important for understanding the empirical methods that enable our contributions.** We present the fair binary classification problem formulation (Section 2.1) and associated empirical approximations (Section 2.2) with an emphasis on the *distribution of possible models* that could be produced from training on different subsets of data

drawn from the same data distribution.

2.1. Problem formulation

Consider a distribution $q(\cdot)$, from which we can sample *examples* $(\boldsymbol{x}, \boldsymbol{g}, o)$; $\boldsymbol{x} \in \mathbb{X} \subseteq \mathbb{R}^m$ are *instances* with m features, and $\boldsymbol{g} \in \mathbb{G}$ is a group of *protected attributes* that we do not use for learning (e.g., race, gender).[1] The $o \in \mathbb{O}$ are the associated *observed labels*, and $\mathbb{O} \subseteq \mathbb{Y}$, where $\mathbb{Y} = \{0, 1\}$ is the label space. From $q(\cdot)$ we can sample training datasets $\{(\boldsymbol{x}, \boldsymbol{g}, o)\}_{i=1}^n$, with $\mathbb{D}$ representing the set of all n-sized datasets. To reason about the possible models of a hypothesis class $\mathbb{H}$ that could be learned from the different subsampled dataset, $\boldsymbol{D}_k \in \mathbb{D}$, we define a *learning process*:

Definition 2.1. A **learning process** is a randomized function that runs instances of a **training procedure** $\mathcal{A}$ on each $\boldsymbol{D}_k \in \mathbb{D}$ and a hypothesis class $\mathbb{H}$, in order to produce a deterministic **classifier** $h_{\boldsymbol{D}_k} \in \mathbb{H}$. A particular run $\mathcal{A}(\boldsymbol{D}_k) \to h_{\boldsymbol{D}_k}$, where $h_{\boldsymbol{D}_k} : \mathbb{X} \to \mathbb{Y}$ is a deterministic mapping from the instance space $\mathbb{X}$ to the label space $\mathbb{Y}$.

By running over different $\mathcal{A}(\boldsymbol{D}_k)$, a learning process produces a distribution over possible trained models, μ. Reasoning about μ, rather than individual models, enables us to contextualize the arbitrariness in the data, which is in turn captured by learned models (Section 3). Each particular model $h_{\boldsymbol{D}_k} \sim \mu$ deterministically produces classifications $\hat{y} = h_{\boldsymbol{D}_k}(\boldsymbol{x})$.[2] The classification rule is $h_{\boldsymbol{D}_k}(\boldsymbol{x}) = \mathbf{1}[r_{\boldsymbol{D}_k}(\boldsymbol{x}) \geq \tau]$, for some threshold τ, where regressor $r_{\boldsymbol{D}_k} : \mathbb{X} \to [0, 1]$ computes the probability of positive classification.

Analyzing error. Executing a training procedure $\mathcal{A}(\boldsymbol{D}_k)$ produces $h_{\boldsymbol{D}_k} \sim \mu$ by minimizing the *loss* of predictions $\hat{y}$ with respect to their associated observed labels o in $\boldsymbol{D}_k$. This loss is computed by a chosen *loss function* $\ell : \mathbb{Y} \times \mathbb{Y} \mapsto \mathbb{R}$. We compute predictions for a *test set* of fresh examples and calculate their loss. The loss is an estimate of the *error* of $h_{\boldsymbol{D}_k}$, which is dependent on the specific dataset $\boldsymbol{D}_k$ used for training. To generalize to the error of all possible models produced by a specific learning process (Definition 2.1), we consider the *expected error*, $\mathrm{Err}(\mathcal{A}, \mathbb{D}, (\boldsymbol{x}, \boldsymbol{g}, o)) = \mathbb{E}_{\boldsymbol{D}}[\ell(o, \hat{y})|\mathbf{x} = \boldsymbol{x}]$.

In fair classification, it is common to use *0-1 loss* $\triangleq \mathbf{1}[\hat{y} \neq o]$ or *cost-sensitive loss*, which assigns asymmetric costs C_{01} for false positives FP and C_{10} for false negatives FN (Agarwal et al., 2018; Elkan, 2001). These costs are related to the classifier threshold $\tau = \frac{C_{01}}{C_{01}+C_{10}}$, with $C_{01}, C_{10} \in \mathbb{R}^+$ (Appendix A.3). Common fair classification definitions, such as Equality of Opportunity (Hardt et al., 2016), further analyze error by computing disparities across group-specific error rates $\mathrm{FPR}_{\boldsymbol{g}}$ and $\mathrm{FNR}_{\boldsymbol{g}}$. For example, $\mathrm{FPR}_{\boldsymbol{g}} \triangleq$

---

[1]We examine the common setting in which $|\boldsymbol{g}| = 1$, and abuse notation by treating $\boldsymbol{g}$ like a scalar with $\mathbb{G} = \{0, 1\}$.

[2]Definition 2.1 does not include the data collection process or hyperparameter optimization, which can further introduce non-determinism, and are assumed to have been already completed.

$p_\mu[r_{\mathbf{D}}(\mathbf{x}) \geq \tau | o = 0, \mathbf{g} = \boldsymbol{g}] = p_\mu[\hat{y} = 1 | o = 0, \mathbf{g} = \boldsymbol{g}]$. Model-specific $\texttt{FPR}_{\boldsymbol{g}}$ and $\texttt{FNR}_{\boldsymbol{g}}$ are further-conditioned on the dataset used in training, i.e., $\mathbf{D} = \boldsymbol{D}_k$.

2.2. Empirical approximation of the formulation

We typically only have access to one dataset, not the data distribution $q(\cdot)$. In empirical work on fair classification, it is common to estimate expected error by performing *cross validation* (CV) on this dataset to produce a small handful of models (Chen et al., 2018; Jiang et al., 2020; Corbett-Davies et al., 2017, e.g.). **CV can be unreliable when there is high variance**; it can produce error estimates that are high variance, and does not reliably estimate expected error with respect to possible models μ (Section 5). For more details, see Efron and Tibshirani (1997; 1993); Wager (2020).

To get around these reliability issues, one can use the *bootstrap method*.[3] Bootstrapping splits the available data into train and test sets, and then simulates drawing different training datasets from a distribution by resampling the train set $\hat{\boldsymbol{D}}$ to generate replicates $\hat{\boldsymbol{D}}_1, \hat{\boldsymbol{D}}_2, \ldots, \hat{\boldsymbol{D}}_B := \hat{\mathbb{D}}$. We use these replicates $\hat{\mathbb{D}}$ to approximate the learning process on $\mathbb{D}$ (Definition 2.1). We evaluate the resulting $\hat{h}_{\hat{\boldsymbol{D}}_1}, \hat{h}_{\hat{\boldsymbol{D}}_2}, \ldots, \hat{h}_{\hat{\boldsymbol{D}}_B}$ by examining their predictions for the *same reserved test set*, which enables us to produce comparisons across test instances like in Figure 1 (Appendix A.4).

3. Variance, Self-Consistency, & Arbitrariness

From these preliminaries, we can now pin down arbitrariness more precisely. We develop a quantitative proxy for measuring arbitrariness, called *self-consistency* (Section 3.2), which is derived from a definition of statistical *variance* between different model predictions (Section 3.1). We then illustrate how self-consistency is a simple yet powerful tool for revealing the role of arbitrariness in fair classification (Section 3.3). In Section 4, we introduce an algorithm to improve self-consistency and compute self-consistency on popular fairness benchmarks in Section 5.

3.1. Arbitrariness resembles statistical variance

In Section 2, we discussed how common fairness definitions analyze error by computing the false positive rate (FPR) and false negative rate (FNR). Another common way to formalize error is as a mathematical decomposition of different statistical sources: *noise-*, *bias-*, and *variance*-induced error (Abu-Mostafa et al., 2012; Geman et al., 1992). In Appendix B.1, we informally describe noise and bias. To understand our metric for self-consistency (Section 3.2), we first describe how the arbitrariness in Figure 1 (almost, but not quite) resembles variance.

Informally, variance-induced error quantifies fluctuations in

individual example predictions for different models $h_{\boldsymbol{D}_k} \in \mathbb{H}$. Variance is the error in the learning process that comes from training on different datasets $\boldsymbol{D}_k \in \mathbb{D}$. In theory, we measure variance by imagining training all possible $h_{\boldsymbol{D}_k} \sim \mu$, testing them all on the same test instance $(\boldsymbol{x}, \boldsymbol{g})$, and then quantifying how much the resulting classifications for $(\boldsymbol{x}, \boldsymbol{g})$ deviate *from each other*. More formally,

Definition 3.1. For all pairs of possible models $h_{\boldsymbol{D}_i}, h_{\boldsymbol{D}_j} \sim \mu$ $(i \neq j)$, the **variance** for a fresh test instance $(\boldsymbol{x}, \boldsymbol{g})$ is
$$\mathrm{var}\big(\mathcal{A}, \mathbb{D}, (\boldsymbol{x}, \boldsymbol{g})\big) \triangleq \mathbb{E}_{h_{\boldsymbol{D}_i} \sim \mu, h_{\boldsymbol{D}_j} \sim \mu}\Big[\ell\big(h_{\boldsymbol{D}_i}(\boldsymbol{x}), h_{\boldsymbol{D}_j}(\boldsymbol{x})\big)\Big].$$

We can approximate variance directly by using the bootstrap method (Section 2.2, Appendix B.1). For 0-1 and cost-sensitive loss with costs $C_{01}, C_{10} \in \mathbb{R}^+$ (Section 2.1), we can generate B replicates to train B concrete models. For $B = B_0 + B_1 > 1$, where B_0 and B_1 respectively denote the number of 0- and 1-class predictions for $(\boldsymbol{x}, \boldsymbol{g})$,
$$\hat{\mathrm{var}}\big(\mathcal{A}, \hat{\mathbb{D}}, (\boldsymbol{x}, \boldsymbol{g})\big) := \frac{1}{B(B-1)} \sum_{i \neq j} \ell\big(\hat{h}_{\hat{\boldsymbol{D}}_i}(\boldsymbol{x}), \hat{h}_{\hat{\boldsymbol{D}}_j}(\boldsymbol{x})\big)$$
$$= \frac{(C_{01} + C_{10}) B_0 B_1}{B(B-1)}. \tag{1}$$
We derive (1) in Appendix B.2 and show that, for increasingly large B, $\hat{\mathrm{var}}$ is defined on $[0, \frac{C_{01} + C_{10}}{4} + \epsilon]$.

3.2. Defining self-consistency from variance

It is clear from above that, in general, variance (1) is unbounded. We can always increase the maximum possible $\hat{\mathrm{var}}$ by increasing the magnitude of our chosen C_{01} and C_{10}.[4] However, we can see from Figure 1 that the most important takeaway is the amount of (dis)agreement, reflected in the counts B_0 and B_1. Here, there is no notion of cost. So, variance (1) does not exactly measure what we want to capture. Instead, we want to focus unambiguously on the (dis)agreement part of variance, which we call *self-consistency of the learning process*:

Definition 3.2. For all pairs of possible models $h_{\boldsymbol{D}_i}, h_{\boldsymbol{D}_j} \sim \mu$ $(i \neq j)$, the **self-consistency of the learning process** for a fresh test instance $(\boldsymbol{x}, \boldsymbol{g})$ is
$$\mathrm{SC}\big(\mathcal{A}, \mathbb{D}, (\boldsymbol{x}, \boldsymbol{g})\big) \triangleq \mathbb{E}_{h_{\boldsymbol{D}_i} \sim \mu, h_{\boldsymbol{D}_j} \sim \mu}\Big[h_{\boldsymbol{D}_i}(\boldsymbol{x}) = h_{\boldsymbol{D}_j}(\boldsymbol{x})\Big]$$
$$= p_{h_{\boldsymbol{D}_i} \sim \mu, h_{\boldsymbol{D}_j} \sim \mu}\big(h_{\boldsymbol{D}_i}(\boldsymbol{x}) = h_{\boldsymbol{D}_j}(\boldsymbol{x})\big). \tag{2}$$
In words, (2) models the probability that two models produced by the same learning process on different training data subsets agree on their predictions for the same test instance.[5] Like variance, we can derive an empirical approximation of SC. Using the bootstrap method with $B = B_0 + B_1 > 1$,
$$\hat{\mathrm{SC}}\big(\mathcal{A}, \hat{\mathbb{D}}, (\boldsymbol{x}, \boldsymbol{g})\big) := \frac{1}{B(B-1)} \sum_{i \neq j} \mathbf{1}\Big[\hat{h}_{\hat{\boldsymbol{D}}_i}(\boldsymbol{x}) = \hat{h}_{\hat{\boldsymbol{D}}_j}(\boldsymbol{x})\Big]$$
$$= 1 - \frac{2 B_0 B_1}{B(B-1)}. \tag{3}$$

---

[3]Alternatively, we could use sampling, e.g. scalable exact MCMC (Zhang et al., 2020). We use optimization, as it is the standard tool (that allows use of standard models) in fair classification.

[4]Because $\tau = \frac{C_{01}}{C_{01} + C_{10}}$, for a given τ we can scale costs arbitrarily and have the same decision rule (Section 2.1). It is relative cost that affects the number of classifications B_0 and B_1, which in turn impact the comparison done in Figure 1.

[5](2) follows from the fact that, in our setup, a learning process is equally likely to draw to any datasets $\boldsymbol{D}_i \boldsymbol{D}_j \in \mathbb{D}$ (Appendix B.3).

Throughout, we use the shorthand *self-consistency*, but it is important to note that Definition 3.2 is a property of the distribution over possible models μ produced by the learning process, not of individual models. We summarize other important takeaways below:

Terminology. In naming our metric, we intentionally evoke related notions of "consistency" in logic and the law (Fuller (1965); Stalnaker (2006), Appendix B.3.1).

Interpretation. Definition 3.2 is defined on $[0.5, 1]$, which coheres with the intuition in Figure 1: 0.5 and 1 respectively reflect minimal (Individual 2) and maximal (Individual 1) SC. SC, unlike FPR and FNR (Section 2.1), does *not* depend on the observed label o. It captures the learning process's confidence in a classification $\hat{y}$, but says nothing directly about $\hat{y}$'s accuracy. By construction, **low self-consistency indicates high variance, and vice versa**. We derive empirical $\hat{SC}$ (3) from $\hat{var}$ (1) by leveraging observations about the definition of $\hat{var}$ for 0-1 loss (Appendix B.3).

Empirical focus. Since SC depends on the particular data subsets used in training, conclusions about its relevance will vary according to task. This is why we take an empirical approach for our main results — of running a large-scale study on many different datasets to extract more general observations about $\hat{SC}$'s practical effects (Section 5). We typically use $B = 101$, which yields a $\hat{SC}$ range of $[\approx 0.495, 1]$ in practice.[6] While there are no costs C_{01}, C_{10} in computing (3), they still effect empirical measurements of $\hat{SC}$. Because C_{01} and C_{10} affect τ (Section 2.1), they impact the concrete number of B_0 and B_1, and thus $\hat{SC}$, measured empirically.

Relationship to fairness concepts. Self-consistency is qualitatively different from traditional fairness metrics. As stated above, unlike FPR and FNR, SC does not depend on observed label o. This has two important implications. First, while the calibration metric in algorithmic fairness also measures a notion of confidence, it is fundamentally different. Calibration reflects confidence with respect to a model predicting o, but says nothing about the relative confidence in predictions $\hat{y}$ produced by the learning process (Pleiss et al., 2017). Second, a common assumption in algorithmic fairness is that there is *label bias* — that unfairness is due in part to discrimination reflected in recorded decisions o (Friedler et al., 2016; Cooper and Abrams, 2021). It is arguably a nice side effect that self-consistency does not depend on o; however, it is also possible to be perfectly self-consistent and inaccurate (e.g., $\hat{y}_k \neq o$, $\forall k$; Section 6).

3.3. Illustrating self-consistency in practice

$\hat{SC}$ enables us to empirically evaluate arbitrariness in classification tasks. It is straightforward to compute $\hat{SC}$ (3) with respect to multiple test instances $(\boldsymbol{x}, \boldsymbol{g})$ — for all instances in a

---

[6]Efron and Tibshirani (1993) recommend $B \in \{50 \ldots 200\}$.

test set or for all instances conditioned on membership in $\boldsymbol{g}$.

Therefore, beyond visualizing $\hat{SC}$ for individuals (Figure 1), we can also do so across sets of individuals. We plot the cumulative distribution (CDF) of $\hat{SC}$ over the groups $\boldsymbol{g}$ in the test set. In Figure 2, we provide illustrative examples from two of the most common fair classification benchmarks (Fabris et al., 2022), COMPAS and Old Adult using random forests (RFs). The x-axis shows the range of self-consistency (for $B = 101$, $[\approx 0.495, 1]$), and the y-axis shows the cumulative proportion of the test set that attains up to the x-amount of $\hat{SC}$. We split the available data into train and test sets, and bootstrap the train set $B = 101$ times to train models $\hat{h_1}, \hat{h_2}, \ldots, \hat{h_{101}}$ (Section 2.2). We repeat this process on 10 train/test splits, and the resulting confidence intervals (shown in inset) indicate that our $\hat{SC}$ estimates are stable. We group observations regarding these examples into two overarching categories:

Individual arbitrariness. Both CDFs show that $\hat{SC}$ varies drastically across test instances. For RFs on the COMPAS dataset, about one-half of instances are under 70% self-consistent. **Nearly one-quarter of test instances are effectively 50% self-consistent; they resemble Individual 2 in Figure 1, meaning that their predictions are essentially arbitrary.** These differences in $\hat{SC}$ across the test set persist even though the 101 models exhibit relatively small average disparities $\Delta E\hat{r}r$, $\Delta F\hat{P}R$, and $\Delta F\hat{N}R$ (Figure 2a, bottom; Section 5.2). This supports our motivating claim: It is possible to come close to satisfying fairness metrics, while the learning process exhibits very different levels of confidence for the underlying classifications that inform those metrics (Section 1).

Systematic arbitrariness. We can also highlight $\hat{SC}$ according to groups $\boldsymbol{g}$. For example, the $\hat{SC}$ plot for Old Adult shows that it is possible for the degree of arbitrariness to be *systematically worse* for a particular demographic $\boldsymbol{g}$ (Figure 2b). While the lack of $\hat{SC}$ is not as extreme as it is for COMPAS (Figure 2a) — the majority of test instances exhibit over 90% $\hat{SC}$ — there is more arbitrariness in the Male subgroup. We can quantify such *systematic arbitrariness* using a measure of distance between probability distributions such as the Wasserstein-1 distance ($\mathcal{W}_1$), which has a closed form for CDFs (Ramdas et al., 2015). The $\mathcal{W}_1$ distance has an intuitive interpretation for measuring systematic arbitrariness: It computes the total disparity in SC by examining all possible SC levels κ at once (Appendix B.3.1). Formally, for two groups $\boldsymbol{g} = 0$ and $\boldsymbol{g} = 1$ with respective SC CDFs F_0 and F_1,

$$\mathcal{W}_1 \triangleq \int_{\mathbb{R}} |F_0(\kappa) - F_1(\kappa)| \, d\kappa. \quad (4)$$

For Old Adult, $\hat{\mathcal{W}}_1 = 0.127$; for COMPAS, which visually does not indicate systematic arbitrariness, $\hat{\mathcal{W}}_1 = 0.007$.

	$\Delta \hat{\mathrm{Err}}$	$\Delta \hat{\mathrm{FPR}}$	$\Delta \hat{\mathrm{FNR}}$
	$1.0 \pm 0.9\%$	$2.0 \pm 0.5\%$	$0.9 \pm 0.3\%$
	$\hat{\mathrm{Err}}$	$\hat{\mathrm{FPR}}$	$\hat{\mathrm{FNR}}$
Total	$36.6 \pm 0.5\%$	$17.3 \pm 0.8\%$	$19.3 \pm 0.7\%$
NW	$36.9 \pm 0.5\%$	$18.0 \pm 0.7\%$	$19.0 \pm 0.8\%$
W	$35.9 \pm 1.3\%$	$16.0 \pm 1.2\%$	$19.9 \pm 1.1\%$

	$\Delta \hat{\mathrm{Err}}$	$\Delta \hat{\mathrm{FPR}}$	$\Delta \hat{\mathrm{FNR}}$
	$12.2 \pm 0.0\%$	$6.0 \pm 0.2\%$	$6.3 \pm 0.2\%$
	$\hat{\mathrm{Err}}$	$\hat{\mathrm{FPR}}$	$\hat{\mathrm{FNR}}$
Total	$17.3 \pm 0.3\%$	$7.7 \pm 0.3\%$	$9.6 \pm 0.1\%$
F	$9.0 \pm 0.3\%$	$3.7 \pm 0.1\%$	$5.3 \pm 0.3\%$
M	$21.2 \pm 0.3\%$	$9.7 \pm 0.3\%$	$11.6 \pm 0.1\%$

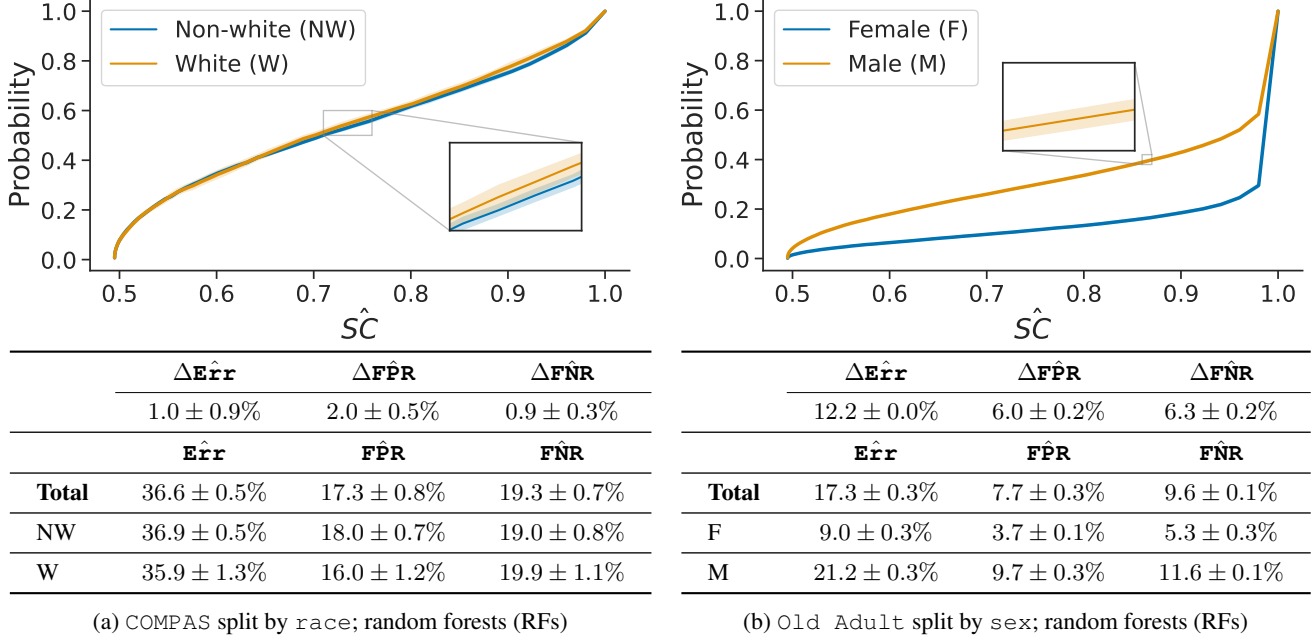(a) `COMPAS` split by `race`; random forests (RFs)   (b) `Old Adult` split by `sex`; random forests (RFs)

Figure 2. $\hat{\mathrm{SC}}$ CDFs for `COMPAS` (2a) and `Old Adult` (2b). We train random forests ($B = 101$ replicates), and repeat with 10 train/test splits to produce confidence intervals (which are very tight). $\hat{\mathrm{SC}}$ is effectively identical across subgroups g in `COMPAS`; `Old Adult` exhibits **systematic** differences in **arbitrariness** across g. Tables show mean $\pm$ STD of the relative disparities $\Delta \hat{\mathrm{Err}}, \Delta \hat{\mathrm{FPR}}, \Delta \hat{\mathrm{FNR}}$, and $\Delta \hat{\mathrm{SC}}$, e.g., $\Delta \hat{\mathrm{Err}} = |\hat{\mathrm{Err}}_0 - \hat{\mathrm{Err}}_1|$ (top); and, the absolute $\hat{\mathrm{Err}}, \hat{\mathrm{FPR}}, \hat{\mathrm{FNR}}$, and $\hat{\mathrm{SC}}$, also broken down by g (bottom) (Appendix E).

4. Accounting for Self-Consistency

By definition, low $\hat{\mathrm{SC}}$ signals that there is high $\hat{\mathrm{var}}$ (Section 3.2). It is therefore a natural idea to use variance reduction techniques to improve $\hat{\mathrm{SC}}$ and thus reduce arbitrariness.

As a starting point for improving $\hat{\mathrm{SC}}$, we perform variance reduction with Breiman (1996)'s *bootstrap aggregation*, or *bagging* ensembling algorithm. Bagging involves bootstrapping to produce a set of B models (Section 2.2), and then, for each test instance, producing an aggregated prediction $\hat{y}_A$, which takes the majority vote of the $\hat{y}_1, \ldots, \hat{y}_B$ classifications. This procedure is practically effective for classifiers with high variance (Breiman, 1996; 1998). However, by taking the majority vote, bagging embeds the idea that having slightly-better-than-random classifiers is sufficient for improving ensembled predictions, $\hat{y}_A$. Unfortunately, there exist instances like Individual 2 (Figure 1), where the classifiers in the ensemble are evenly split between classes. This means that bagging alone cannot overcome arbitrariness.

4.1. Self-consistent ensembling with abstention

To remedy this, we add the option to abstain from prediction if $\hat{\mathrm{SC}}$ is low (Algorithm 1).[7] We bootstrap as usual, but only produce a prediction $\hat{y} \in [0, 1]$ for an instance $\boldsymbol{x}$, if $\boldsymbol{x}$ surpasses a user-specified minimum level of $\hat{\mathrm{SC}}$ κ; otherwise, if an instance fails to achieve a $\hat{\mathrm{SC}}$ of at least κ, we `Abstain`

---

[7] A minor adjustment to (3) accounts for abstentions; a simple proof follows that Algorithm 1 improves $\hat{\mathrm{SC}}$ (Appendix D).

from predicting. Based on this strategy, we end up dividing the test set into two subsets: We group together the instances we `Abstain` on in an *abstention set* and those we predict on in a *prediction set*. This method improves self-consistency through two complementary mechanisms: 1) Variance reduction (inherited from bagging) and 2) abstaining from instances that exhibit low $\hat{\mathrm{SC}}$ (thereby raising the overall amount of $\hat{\mathrm{SC}}$ for the prediction set).

Further, since variance is a component of error (Section 3), variance reduction also tends to improve accuracy (Breiman, 1996). This leads to an important observation: The abstention set, by definition, exhibits high variance; we can therefore expect it to exhibit higher error than the prediction set (Section 5, Appendix E). So, while at first glance it may seem odd that our solution for arbitrariness is to *not predict*, it is worth noting that **we often would have predicted incorrectly on a large portion of the abstention set** (Appendix D). In practice, we test two versions of our method:

Simple ensembling. We run Algorithm 1 to build ensembles of typical hypothesis classes. For example, running with $B = 101$ decision trees and $\kappa = 0.75$ produces a bagged classifier that contains 101 underlying decision trees, for which the bagged classifier abstains from predicting on test instances that exhibit less than 75% $\hat{\mathrm{SC}}$. If overall $\hat{\mathrm{SC}}$ is low, then simple ensembling will lead to a large number of abstentions. For example, almost half of all test instances in `COMPAS` using random forests would fail to surpass the threshold $\kappa = 0.75$ (Figure 2a). The potential

---

Algorithm 1 Self-Consistent Ensembling with Abstention

---

input training dataset $(\boldsymbol{X}, \boldsymbol{o})$, $\mathcal{A}$, B, $\hat{\text{SC}}$ $\kappa \in [0.5, 1]$, $\boldsymbol{x}_{\text{test}}$

output $\hat{y}$ with $\hat{\text{SC}} \geq \kappa$ or `Abstain`

1: **function** BagConsistently$\big((\boldsymbol{X}, \boldsymbol{o}), \mathcal{A}, B, \kappa, \boldsymbol{x}_{\text{test}}\big)$
2:     $\hat{y}_A := \text{list}()$ ▷ To store ensemble predictions
3:     **for** $1 \ldots B$ **do**
4:        $\boldsymbol{D}_B \leftarrow \text{Bootstrap}\big((\boldsymbol{X}, \boldsymbol{o})\big)$
         ▷ $\mathcal{A}$ may produce bagged models $\hat{h}_{\boldsymbol{D}_B}$ using $\boldsymbol{D}_B$
5:        $\hat{h}_{\boldsymbol{D}_B} \leftarrow \mathcal{A}(\boldsymbol{D}_B)$
6:        $\hat{y}_A.\text{append}\big(\hat{h}_{\boldsymbol{D}_B}(\boldsymbol{x}_{\text{test}})\big)$ ▷ $\hat{y}_1, \ldots, \hat{y}_B$
7:     **end for**
8:     **return** Aggregate$(\hat{y}_A, \kappa)$
9: **end function**
10: ▷ Ensures a super-majority to predict, otherwise abstains
11: **function** Aggregate$\big(\hat{y}_1, \ldots, \hat{y}_B, \kappa\big)$
12:     **if** SelfConsistency$(\hat{y}_1, \ldots, \hat{y}_B) \geq \kappa$ **then**
13:        **return** $\arg\max_{y' \in \mathbb{Y}} \left[\sum_{i=1}^{B} \mathbb{1}[y' = \hat{y}_i] \right]$
14:     **end if**
15:     **return** `Abstain`
16: **end function**

---

for large abstention sets informs our second approach:

Super ensembling. When there is low $\hat{\text{SC}}$ (and thus high $\hat{\text{var}}$) it can be beneficial to do an initial pass of variance reduction. That is, we produce bagged classifiers using traditional bagging but without abstaining, and *then* apply Algorithm 1, using those bagged classifiers as the underlying models $\hat{h}$. The first round of bagging raises the overall $\hat{\text{SC}}$ before the second round, which is when we decide whether to `Abstain`. We therefore expect this approach to abstain less; however, it may potentially incur higher error, if simple-majority-vote bagging chooses $\hat{y} \neq o$ (Appendix D).

We also experiment with an Aggregate rule that averages the output probabilities of the underlying regressors $r_{\boldsymbol{D}_k}$, and then applies threshold τ to produce ensembled predictions. We do not observe major differences in results.

5. Experiments

We release an extensible package of different Aggregate methods, with which **we trained and compared several million different models** over the course of our study. We include results covering common fair classification datasets and models: `COMPAS`, `Old Adult`, `German Credit`, `Taiwan Credit`, and 3 large-scale `New Adult – CA` tasks on logistic regression, decision trees, random forests, MLPs, and SVMs (Appendix E). Ultimately, when applying Algorithm 1 to these common tasks, **we typically do not find evidence of significant fairness metric violations. By accounting for arbitrariness, we observe close-to-fairness in nearly every task, _without_ applying common fairness-improving interventions** (Section 5.2, Appendix E). A possible explanation is that most fairness benchmark tasks are small — fewer than $25,000$ examples — and this finding might not generalize to larger datasets.

Releasing an `HMDA` toolkit. We clean and release a larger, more diverse, and newer dataset for investigating algorithmic fairness.[8] We examine the `NY` and `TX 2017` subsets of the the Home Mortgage Data Disclosure Act (`HMDA`) 2007-2017 dataset (Federal Financial Institutions Examination Council, 2017), which have $244,107$ and $576,978$ examples, respectively. We process all `HMDA` datasets and release them with a standalone, easy-to-use software package.

Presentation. To visualize Algorithm 1, we plot the CDFs of the $\hat{\text{SC}}$ of the underlying models used in each ensembling method. We simultaneously plot the results of **simple ensembling** (dotted curves) and **super ensembling** (solid curves). Instances to the left of the vertical line (the minimum $\hat{\text{SC}}$ threshold κ) form the abstention set. We also provide corresponding mean $\pm$ STD fairness and accuracy metrics for individual models (our **baseline**) and for both **simple** and **super**. For ensembling methods, we report these metrics on the prediction set, along with the **abstention rate ($\hat{\text{AR}}$)**.

We present our results in two sections: We validate the effectiveness of both ensembling variants (Section 5.1); we then explain how our results reveal shocking insights about reproducibility in fairness research (Section 5.2). For all experiments, we illustrate Algorithm 1 with $\kappa = 0.75$, but κ is task-dependent in practice. We detail our setup, hyperparameter optimization, and full results in Appendix E.

5.1. Validating Algorithm 1

We highlight results for two illustrative examples: `Old Adult` (Section 3.3) and `HMDA-NY-2017`, for `ethnicity` (Hispanic or Latino (HL), Non-Hispanic or Latino (NHL)). We plot $\hat{\text{SC}}$ CDFs and show $\hat{\text{FNR}}$ metrics using random forests (RFs; Figure 3). For `Old Adult`, the expected disparity of the RF baseline is $\Delta\hat{\text{FNR}} = 6.3\%$. The light set of curves plots the underlying $\hat{\text{SC}}$ for these RFs (Figure 3a). When we apply **simple** to these RFs, overall $\hat{\text{Err}}$ decreases (Appendix E), shown in part by the decrease in $\hat{\text{FNR}}_F$ and $\hat{\text{FNR}}_M$. Fairness also improves: $\Delta\hat{\text{FNR}}$ decreases to 4.1%. However, the corresponding $\hat{\text{AR}}$ is quite high, especially for the `Male` subgroup ($\boldsymbol{g} = \text{M}$, Figure 4).

As expected, **super** improves overall $\hat{\text{SC}}$ through a first pass of variance reduction (Section 4). The $\hat{\text{SC}}$ CDF curves are brought down, indicating a lower proportion of the test set exhibits low $\hat{\text{SC}}$. Abstention rate $\hat{\text{AR}}$ is lower and more equal; however, error, while still lower than the baseline RFs, has gone up for all metrics. There is also a decrease in systematic arbitrariness (Section 3.3): The dark gray area for **super** ($\hat{\mathcal{W}}_1 = .014$) is smaller than the light gray area for **simple** ($\hat{\mathcal{W}}_1 = .063$) (Appendices B.3.1 & E.4.7).
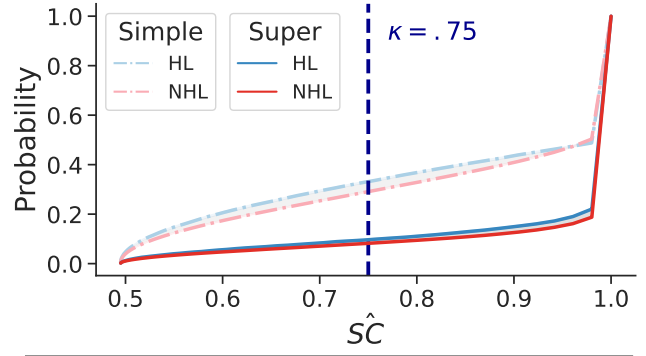
For `HMDA`, **simple** similarly improves $\hat{\text{FNR}}$, but has a less

---

[8]It is repeatedly argued that the field needs such datasets (Ding et al., 2021, e.g.). `HMDA` is less commonly used, possibly because the ≈ 140 million examples are only available in multi-GiB files.
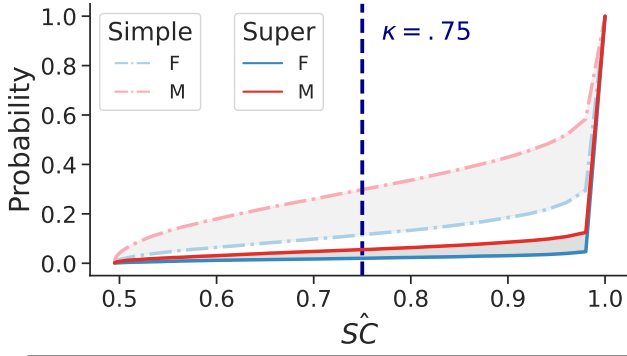
	Baseline	Simple	Super
$\Delta\hat{\text{FNR}}$	$6.3 \pm .2\%$	$4.1 \pm .2\%$	$5.8 \pm .1\%$
$\hat{\text{FNR}}_F$	$5.3 \pm .3\%$	$3.5 \pm .1\%$	$4.9 \pm .2\%$
$\hat{\text{FNR}}_M$	$11.6 \pm .1\%$	$7.6 \pm .3\%$	$10.7 \pm .3\%$

(a) `Old Adult` split by `sex`; random forests (RFs)

	Baseline	Simple	Super
$\Delta\hat{\text{FNR}}$	$0.7 \pm .1\%$	$1.1 \pm .2\%$	$2.2 \pm .2\%$
$\hat{\text{FNR}}_{HL}$	$10.1 \pm .2\%$	$3.3 \pm .3\%$	$8.0 \pm .3\%$
$\hat{\text{FNR}}_{NHL}$	$9.4 \pm .1\%$	$2.2 \pm .1\%$	$5.8 \pm .1\%$

(b) `HMDA-NY-2017` split by `ethnicity`; random forests (RFs)

Figure 3. Algorithm 1: **Simple** and **super ensembling** for `Old Adult` (3a) and `HMDA-NY-2017` (3b). Tables each show $\hat{\text{FNR}}$ metrics (mean $\pm$ STD) for individual models (**Baseline**) and each ensembling method's prediction set; $B = 101$, 10 train/test splits (Appendix E). To highlight systematic arbitrariness (Section 3.3), we shade in gray the area between group-specific $\hat{\text{SC}}$ CDFs for each method. An initial pass of variance reduction in **super** significantly decreases the systematic arbitrariness in `Old Adult`.

beneficial effect on fairness ($\Delta\hat{\text{FNR}}$). However, note that since the baseline is the empirical expected error over thousands of RF models, the specific $\Delta\hat{\text{FNR}}$ is not necessarily attainable by any individual model. In this respect, **simple** has the benefit of obtaining this disparity reliably in practice: $\Delta\hat{\text{FNR}} = 1.3\%$ is the mean over 10 simple ensembles, and, notably, is extremely low, even without applying traditional fairness techniques. Similar to `Old Adult`, **simple** exhibits high $\hat{\text{AR}}$, which decreases with **super** at the cost of higher error. $\hat{\text{FNR}}$ still improves for both g in comparison to the baseline, but the benefits are unequally applied: $\hat{\text{FNR}}_W$

(a) `Old Adult`, $g = $ `sex`

(b) `HMDA-NY-2017`, $g = $ `ethnicity`

Figure 4. Group-specific abstention rates $\hat{\text{AR}}_g$ for each algorithm. **Super** abstains less than **simple**, and the disparity in $\hat{\text{AR}}$ between subgroups in both datasets is similar.

has a larger benefit, so $\Delta\hat{\text{FNR}}$ increases slightly.

Abstention set error. As an example, the average total $\hat{\text{Err}}$ in the `Old Adult` **simple** abstention set is close to 40% — compared to 17% for the RF baseline, and 8% for **simple** and 14% for **super** prediction sets. Thus, beyond reducing arbitrariness, we abstain from predicting for many instances for which we also would have been inaccurate.

Trade-off between abstention and error. Algorithm 1 identifies low $\hat{\text{SC}}$ instances for which ML prediction does a poor job. Nevertheless, it may be infeasible for some applications to tolerate a high $\hat{\text{AR}}$. Thus the choice of κ and ensembling method should be considered a context-dependent choice (Section 6).

Unequal abstention rates. When there is a high degree of systematic arbitrariness, $\hat{\text{AR}}$ can vary a lot by g (Figure 4a). With respect to improving $\hat{\text{SC}}$ and error, this may be a reasonable outcome. Further, we rarely observe systematic arbitrariness; unequal $\hat{\text{AR}}$ is uncommon in practice (Section 6).

5.2. A problem of fairness

We additionally experiment on `COMPAS`, one of the three most common fairness datasets (Fabris et al., 2022). Algorithm 1 is similarly very effective at reducing arbitrariness (Figure 5), and is able to obtain state-of-the-art accuracy (Lin et al., 2020) with $\Delta\hat{\text{FPR}}$ between $1.8 - 3\%$; analogous results for `German Credit` indicate statistical equivalence in fairness metrics (Appendix E).

These low-single-digit fairness disparities do not cohere with much of the literature on fair classification, which often reports much larger fairness violations (Larson et al., 2016,
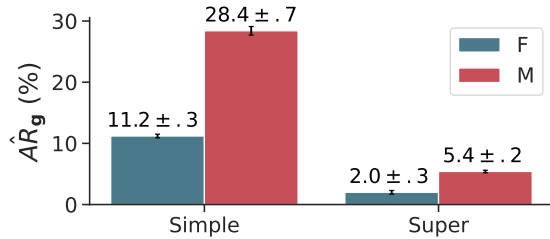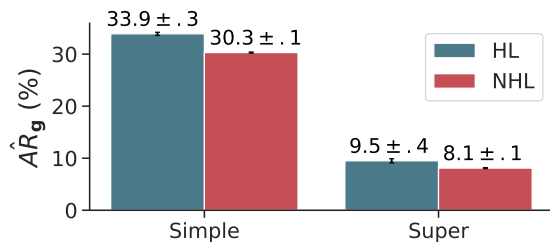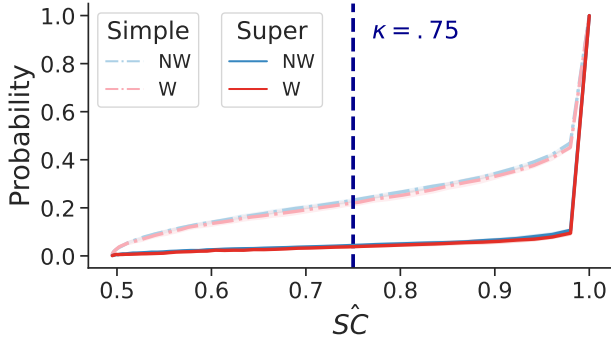
	Baseline	Simple	Super
$\Delta\hat{\text{FPR}}$	$2.1 \pm 0.0\%$	$3.0 \pm 0.0\%$	$1.8 \pm .2\%$
$\hat{\text{FPR}}_{\text{NW}}$	$14.7 \pm 1.3\%$	$11.4 \pm 1.0\%$	$12.9 \pm .8\%$
$\hat{\text{FPR}}_{\text{W}}$	$12.6 \pm 1.3\%$	$8.4 \pm 1.0\%$	$11.1 \pm .6\%$

Figure 5. Algorithm 1: COMPAS, logistic regression, $B = 101$, 10 train/test splits. Table shows mean $\hat{\text{FPR}} \pm$ STD for individual models (**Baseline**) and each ensembling method's prediction set. The $\hat{SC}$ CDFs are effectively identical across subgroups.

notably). However, as noted in Section 2, most work on fair classification examines individual models, selected via cross-validation with a handful of random seeds. Our results suggest that selecting between a few individual models in fair classification experiments is unreliable (Wager, 2020). When we more rigorously estimate expected error and ensemble, we have difficulty reproducing unfairness in practice. Variance in the underlying models in μ seems to be the culprit. The individual models we train on these tasks exhibit radically different subgroup-specific error rates. Our strategy of shifting focus to the overall behavior of μ provides a solution: We not only mitigate arbitrariness, we **also improve accuracy *and* usually average away most underlying, individual model unfairness** (Appendix E.5).

6. Discussion and Related Work

In this paper, we advocate for a shift in thinking about *individual* models to the *distribution over possible models* in fair classification tasks, as this makes visible latent arbitrariness in underlying model decisions. We suggest a metric of *self-consistency* as a proxy for arbitrariness (Section 3) and an intuitive extension of the classic bagging algorithm to mitigate it (Section 4), which we observe to be tremendously effective with respect to improving $\hat{SC}$, accuracy, and fairness metrics (Section 5, Appendix E).

Our findings contradict accepted truths in algorithmic fairness. For example, while much work posits that there is an inherent analytical trade-off between fairness and accuracy (Corbett-Davies et al., 2017; Menon and Williamson, 2018), our empirical studies complement prior work that disputes the practical relevance of this formulation (Rodolfa et al., 2021). We show it is typically possible to achieve higher accuracy while retaining close-to-fairness — and

to do so **without** using fairness-focused interventions.

Other research also calls attention to the need for metrics beyond fairness and accuracy. *Model multiplicity* reasons about sets of models that have similar accuracy (Breiman, 2001), but differ in underlying properties due to variance in decision rules (Black et al., 2022b; Marx et al., 2020; Watson-Daniels et al., 2022). This work still emphasizes selecting in an *individual* model. Instead, our work uses a *distribution over possible models* (with no claims about model accuracy) to reason about arbitrariness. Some related work considers the role of uncertainty and variance in fairness (Ali et al., 2021; Bhatt et al., 2021; Chen et al., 2018; Khan et al., 2023). Notably, Black et al. (2022a) concurrently investigate abstention-based ensembling, employing a strategy that does not fundamentally address the arbitrariness we describe (Appendix C). Also concurrently, Ko et al. (2023) build on prior work that studies fairness and variance in deep learning tasks (Forde et al., 2021; Qian et al., 2021), and find that fairness emerges in deep ensembles.

Most importantly, we take a comprehensive empirical approach missing from prior work. We uncover alarming results: **Almost all tasks and settings demonstrate close-to or complete statistical equality in fairness metrics, after accounting for arbitrariness** (Appendix E.4). Old Adult (Figure 3a) is one of two exceptions. Further, we do not find many tasks that exhibit high systematic arbitrariness and, even when we do, we can substantially improve it. These result hold for larger, newer datasets like HMDA, which we clean and release, and New Adult (Ding et al., 2021). Altogether, our findings indicate that **variance is undermining the reliability of conclusions in fair classification experiments**. It is worth revisiting all prior fair classification experiments that depend on CV or few models.

So, what does this mean for fairness research? While the field has put forth numerous theoretical results about (un)fairness regarding single models — impossibility of satisfying multiple metrics (Kleinberg et al., 2017), post-processing individual models to achieve a particular metric (Hardt et al., 2016) — these results seem to miss the point. By examining individual models, arbitrariness remains latent; when we account for arbitrariness in practice, most problems of empirical unfairness go away.

We are not suggesting that there are no reasons to be concerned with the fairness of machine learning models. We are not even challenging the idea that actual violations of standard fairness metrics should be of concern. Instead, we are suggesting that common formalisms for measuring fairness can lead to false conclusions about the degree to which such violations are happening in practice. Worse, they can conceal the tremendous amount of arbitrariness that should really be the issue of concern.

Lastly, it has not escaped our notice that, altogether, our results signal severe limits to prediction in social settings (Narayanan and Salganik, 2020). It is true that our method performs reasonably well with respect to both fairness and accuracy metrics; however, arbitrariness is such a rampant problem, it is arguably unreasonable to assign these metrics much value in practice.

Acknowledgements

References

Daniel Abadi. Consistency Tradeoffs in Modern Distributed Database System Design: CAP is Only Part of the Story. *Computer*, 45(2):37–42, February 2012.

Yaser S. Abu-Mostafa, Malik Magdon-Ismail, and Hsuan-Tien Lin. *Learning From Data: A Short Course*. AMLbook, USA, 2012.

Alekh Agarwal, Alina Beygelzimer, Miroslav Dudik, John Langford, and Hanna Wallach. A Reductions Approach to Fair Classification. In Jennifer Dy and Andreas Krause, editors, *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 60–69. PMLR, 10–15 Jul 2018.

Junaid Ali, Preethi Lahoti, and Krishna P. Gummadi. Accounting for Model Uncertainty in Algorithmic Discrimination. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*, AIES '21, page 336–345, New York, NY, USA, 2021. Association for Computing Machinery. ISBN 9781450384735. doi: 10.1145/3461702.3462630. URL https://doi.org/10.1145/3461702.3462630.

Solon Barocas, Moritz Hardt, and Arvind Narayanan. *Fairness and Machine Learning: Limitations and Opportunities*. fairmlbook.org, 2019. http://www.fairmlbook.org.

Umang Bhatt, Javier Antorán, Yunfeng Zhang, Q. Vera Liao, Prasanna Sattigeri, Riccardo Fogliato, Gabrielle Melançon, Ranganath Krishnan, Jason Stanley, Omesh Tickoo, Lama Nachman, Rumi Chunara, Madhulika Srikumar, Adrian Weller, and Alice Xiang. Uncertainty as a Form of Transparency: Measuring, Communicating, and Using Uncertainty. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*, page 401–413. Association for Computing Machinery, New York, NY, USA, 2021.

Emily Black, Klas Leino, and Matt Fredrikson. Selective Ensembles for Consistent Predictions. In *International Conference on Learning Representations*, 2022a.

Emily Black, Manish Raghavan, and Solon Barocas. Model Multiplicity: Opportunities, Concerns, and Solutions. In *2022 ACM Conference on Fairness, Accountability, and Transparency*, FAccT '22, page 850–863, New York, NY, USA, 2022b. Association for Computing Machinery. ISBN 9781450393522. doi: 10.1145/3531146.3533149. URL https://doi.org/10.1145/3531146.3533149.

Leo Breiman. Bagging Predictors. *Machine Learning*, 24(2):123–140, August 1996. ISSN 0885-6125. doi: 10.1023/A:1018054314350.

Leo Breiman. Arcing classifiers. *Annals of Statistics*, 26: 801–823, 1998.

Leo Breiman. Statistical Modeling: The Two Cultures. *Statistical Science*, 16(3):199–215, 2001. ISSN 08834237.

Toon Calders, Faisal Kamiran, and Mykola Pechenizkiy. Building Classifiers with Independency Constraints. In *2009 IEEE International Conference on Data Mining Workshops*, pages 13–18, 2009.

Irene Chen, Fredrik D Johansson, and David Sontag. Why Is My Classifier Discriminatory? In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018.

Alexandra Chouldechova. Fair Prediction with Disparate Impact: A Study of Bias in Recidivism Prediction Instruments. *Big Data*, 5(2):153–163, 2017.

Danielle Keats Citron and Frank A. Pasquale. The Scored Society: Due Process for Automated Predictions. *Washington Law Review*, 89:655–690, 2014.

A. Feder Cooper and Ellen Abrams. Emergent Unfairness in Algorithmic Fairness-Accuracy Trade-Off Research. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*, page 46–54, New York, NY, USA, 2021. Association for Computing Machinery. ISBN 9781450384735.

A. Feder Cooper, Karen Levy, and Christopher De Sa. Accuracy-Efficiency Trade-Offs and Accountability in Distributed ML Systems. In *Equity and Access in Algorithms, Mechanisms, and Optimization*, New York, NY, USA, 2021a. Association for Computing Machinery. ISBN 9781450385534. URL https://doi.org/10.1145/3465416.3483289.

A. Feder Cooper, Yucheng Lu, Jessica Zosa Forde, and Christopher De Sa. Hyperparameter Optimization Is Deceiving Us, and How to Stop It. In *Advances in Neural Information Processing Systems*, volume 34. Curran Associates, Inc., 2021b.

A. Feder Cooper, Emanuel Moss, Benjamin Laufer, and Helen Nissenbaum. Accountability in an Algorithmic Society: Relationality, Responsibility, and Robustness in Machine Learning. In *2022 ACM Conference on Fairness, Accountability, and Transparency*, FAccT '22, page 864–876, New York, NY, USA, 2022. Association for Computing Machinery. ISBN 9781450393522. doi: 10.1145/3531146.3533150.

Sam Corbett-Davies, Emma Pierson, Avi Feller, Sharad Goel, and Aziz Huq. Algorithmic decision making and the cost of fairness. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '17, page 797–806, New York, NY, USA, 2017. Association for Computing Machinery. ISBN 9781450348874.

Kathleen Creel and Deborah Hellman. The Algorithmic Leviathan: Arbitrariness, Fairness, and Opportunity in Algorithmic Decision-Making Systems. *Canadian Journal of Philosophy*, 52(1):26–43, 2022.

Frances Ding, Moritz Hardt, John Miller, and Ludwig Schmidt. Retiring Adult: New Datasets for Fair Machine Learning. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, volume 34, pages 6478–6490. Curran Associates, Inc., 2021.

Pedro Domingos. A Unified Bias-Variance Decomposition and Its Applications. In *Proceedings of the Seventeenth International Conference on Machine Learning*, ICML '00, page 231–238, San Francisco, CA, USA, 2000a. Morgan Kaufmann Publishers Inc. ISBN 1558607072.

Pedro Domingos. A unified bias-variance decomposition. Technical report, University of Washington, 2000b. URL https://homes.cs.washington.edu/~pedrod/bvd.pdf.

B. Efron. Bootstrap Methods: Another Look at the Jackknife. *The Annals of Statistics*, 7(1):1 – 26, 1979. doi: 10.1214/aos/1176344552.

Bradley Efron and Robert Tibshirani. Improvements on Cross-Validation: The 632+ Bootstrap Method. *Journal of the American Statistical Association*, 92(438):548–560, 1997. doi: 10.1080/01621459.1997.10474007.

Bradley Efron and Robert J. Tibshirani. *An Introduction to the Bootstrap*. Number 57 in Monographs on Statistics and Applied Probability. Chapman & Hall/CRC, Boca Raton, Florida, USA, 1993.

Charles Elkan. The Foundations of Cost-Sensitive Learning. In *Proceedings of the 17th International Joint Conference on Artificial Intelligence - Volume 2*, IJCAI'01, page 973–978, San Francisco, CA, USA, 2001. Morgan Kaufmann Publishers Inc.

Alessandro Fabris, Stefano Messina, Gianmaria Silvello, and Gian Antonio Susto. Algorithmic Fairness Datasets: the Story so Far. *Data Mining and Knowledge Discovery*, 36(6):2074–2152, September 2022. doi: 10.1007/s10618-022-00854-z.

Federal Financial Institutions Examination Council. HMDA Data Publication, 2017. URL https://www.consumerfinance.gov/data-research/hmda/historic-data/. Released due to the Home Mortgage Disclosure Act.

Jessica Zosa Forde, A. Feder Cooper, Kweku Kwegyir-Aggrey, Chris De Sa, and Michael L. Littman. Model Selection's Disparate Impact in Real-World Deep Learning Applications, 2021. URL https://arxiv.org/abs/2104.00606.

Sorelle A. Friedler, Carlos Scheidegger, and Suresh Venkatasubramanian. On the (im)possibility of fairness, 2016.

Sorelle A. Friedler, Carlos Scheidegger, Suresh Venkatasubramanian, Sonam Choudhary, Evan P. Hamilton, and Derek Roth. A Comparative Study of Fairness-Enhancing Interventions in Machine Learning. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*, FAT* '19, page 329–338, New York, NY, USA, 2019. Association for Computing Machinery. ISBN 9781450361255. doi: 10.1145/3287560.3287589.

Lon L. Fuller. *The Morality of Law*. Yale University Press, New Haven, 1965. ISBN 9780300010701.

Stuart Geman, Elie Bienenstock, and René Doursat. Neural Networks and the Bias/Variance Dilemma. *Neural Comput.*, 4(1):1–58, January 1992. ISSN 0899-7667. doi: 10.1162/neco.1992.4.1.1.

Ulrike Grömping. South German Credit Data: Correcting a Widely Used Data Set. Technical report, Beuth University of Applied Sciences Berlin, 2019.

Moritz Hardt, Eric Price, Eric Price, and Nati Srebro. Equality of Opportunity in Supervised Learning. In D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc., 2016.

Trevor Hastie, Robert Tibshirani, and Jerome Friedman. *The Elements of Statistical Learning: Data Mining, Inference and Prediction*. Springer, USA, 2 edition, 2009.

Jaakko Hintikka. *Knowledge and Belief.* Cornell University Press, 1962.

Ray Jiang, Aldo Pacchiano, Tom Stepleton, Heinrich Jiang, and Silvia Chiappa. Wasserstein Fair Classification. In Ryan P. Adams and Vibhav Gogate, editors, *Proceedings of The 35th Uncertainty in Artificial Intelligence Conference*, volume 115 of *Proceedings of Machine Learning Research*, pages 862–872. PMLR, 22–25 Jul 2020.

Falaah Arif Khan, Denys Herasymuk, and Julia Stoyanovich. On Fairness and Stability: Is Estimator Variance a Friend or a Foe?, 2023.

Jon M. Kleinberg, Sendhil Mullainathan, and Manish Raghavan. Inherent Trade-Offs in the Fair Determination of Risk Scores. In Christos H. Papadimitriou, editor, *8th Innovations in Theoretical Computer Science Conference, ITCS 2017, January 9-11, 2017, Berkeley, CA, USA*, volume 67 of *LIPIcs*, pages 43:1–43:23. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2017.

Wei-Yin Ko, Daniel D'souza, Karina Nguyen, Randall Balestriero, and Sara Hooker. FAIR-Ensemble: When Fairness Naturally Emerges From Deep Ensembling, 2023.

Ron Kohavi. Scaling up the Accuracy of Naive-Bayes Classifiers: A Decision-Tree Hybrid. In *Proceedings of the Second International Conference on Knowledge Discovery and Data Mining*, KDD'96, page 202–207. AAAI Press, 1996.

Eun Bae Kong and Thomas G. Dietterich. Error-Correcting Output Coding Corrects Bias and Variance. In Armand Prieditis and Stuart Russell, editors, *Machine Learning, Proceedings of the Twelfth International Conference on Machine Learning, Tahoe City, California, USA, July 9-12, 1995*, pages 313–321, 1995.

Jeff Larson, Surya Mattu, Lauren Kirchner, and Julia Angwin. How We Analyzed the COMPAS Recidivism Algorithm. Technical report, ProPublica, May 2016. URL https://www.propublica.org/article/how-we-analyzed-the-compas-recidivism-algorithm.

Zhiyuan Jerry Lin, Jongbin Jung, Sharad Goel, and Jennifer Skeem. The limits of human predictions of recidivism. *Science Advances*, 6(7):eaaz0652, 2020. doi: 10.1126/sciadv.aaz0652.

Paul Mackun, Joshua Comenetz, and Lindsay Spell. Around Four-Fifths of All U.S. Metro Areas Grew Between 2010 and 2020. Technical report, United States Census Bureau, August 2021. URL https://www.census.gov/library/stories/2021/08/more-than-half-of-united-states-counties-were-smaller-in-2020-than-in-2010.html.

Charles Marx, Flavio Calmon, and Berk Ustun. Predictive Multiplicity in Classification. In Hal Daumé III and Aarti Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 6765–6774. PMLR, 13–18 Jul 2020. URL https://proceedings.mlr.press/v119/marx20a.html.

Aditya Krishna Menon and Robert C. Williamson. The cost of fairness in binary classification. In Sorelle A. Friedler and Christo Wilson, editors, *Proceedings of the 1st Conference on Fairness, Accountability and Transparency*, volume 81 of *Proceedings of Machine Learning Research*, pages 107–118. PMLR, 23–24 Feb 2018.

Arvind Narayanan and Matt Salganik. Limits to prediction: pre-read. Technical report, Princeton University, 2020. URL https://msalganik.github.io/cos597E-soc555_f2020/class-materials/2020-09-03-background/limits-to-prediction-pre-read.pdf.

Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, Jake Vanderplas, Alexandre Passos, David Cournapeau, Matthieu Brucher, Matthieu Perrot, and Édouard Duchesnay. Scikit-Learn: Machine Learning in Python. *J. Mach. Learn. Res.*, 12:2825–2830, November 2011. ISSN 1532-4435.

Geoff Pleiss, Manish Raghavan, Felix Wu, Jon Kleinberg, and Kilian Q Weinberger. On Fairness and Calibration. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017.

Shangshu Qian, Hung Pham, Thibaud Lutellier, Zeou Hu, Jungwon Kim, Lin Tan, Yaoliang Yu, Jiahao Chen, and Sameena Shah. Are My Deep Learning Systems Fair? An Empirical Study of Fixed-Seed Training. In *Advances*

in Neural Information Processing Systems, volume 34, Red Hook, NY, USA, 2021. Curran Associates, Inc.

Aaditya Ramdas, Nicolas Garcia, and Marco Cuturi. On Wasserstein Two Sample Testing and Related Families of Nonparametric Tests, 2015. URL https://arxiv.org/abs/1509.02237.

K.T. Rodolfa, H. Lamba, and R Ghani. Empirical observation of negligible fairness–accuracy trade-offs in machine learning for public policy. *Nature Machine Intelligence*, 3, 2021. doi: https://doi.org/10.1038/s42256-021-00396-x.

Raymond M. Smullyan. Logicians Who Reason about Themselves. In *Proceedings of the 1986 Conference on Theoretical Aspects of Reasoning about Knowledge*, TARK '86, page 341–352, San Francisco, CA, USA, 1986. Morgan Kaufmann Publishers Inc. ISBN 0934613049.

Robert Stalnaker. On Logics of Knowledge and Belief. *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition*, 128(1):169–199, 2006. ISSN 00318116, 15730883. URL http://www.jstor.org/stable/4321718.

Brian Z. Tamanaha. *On the Rule of Law: History, Politics, Theory*. Cambridge University Press, Cambridge, 2004.

Stefan Wager. Cross-Validation, Risk Estimation, and Model Selection: Comment on a Paper by Rosset and Tibshirani. *Journal of the American Statistical Association*, 115(529):157–160, 2020. doi: 10.1080/01621459.2020.1727235.

Jamelle Watson-Daniels, David C. Parkes, and Berk Ustun. Predictive Multiplicity in Probabilistic Classification, 2022.

I.C. Yeh and C.H. Lien. The comparisons of data mining techniques for the predictive accuracy of probability of default of credit card clients. *Expert Systems with Applications*, 36:2473–2480, 2009.

Ruqi Zhang, A Feder Cooper, and Christopher De Sa. AMAGOLD: Amortized Metropolis Adjustment for Efficient Stochastic Gradient MCMC. In *International Conference on Artificial Intelligence and Statistics*, 2020.

Appendix Overview

This supplemental material goes into significantly more detail regarding the results presented in this paper. The organization is as follows:

A. Extended Preliminaries

A.1. Notes on notation and on our choice of terminology

Terminology. Traditionally, what we term "observed labels" o are often referred to instead as the "ground truth" or "correct" labels (Abu-Mostafa et al., 2012; Hastie et al., 2009; Kong and Dietterich, 1995, e.g.). We avoid this terminology because, as the work on label bias has explained, these labels are often unreliable or contested (Friedler et al., 2016; Cooper and Abrams, 2021).

Sets, random variables, and instances. We use bold non-italics letters to denote random variables (e.g., $\mathbf{x}$, $\mathbf{D}$), capital block letters to denote sets (e.g., $\mathbb{X}$, $\mathbb{Y}$), lower case italics letters to denote scalars (e.g., o), bold italics lower case letters to denote vectors (e.g., $\boldsymbol{x}$), and bold italics upper case to denote matrices (e.g., $\boldsymbol{D}_k$). For a complete example, $\boldsymbol{x}$ is an arbitrary instance's feature vector, $\mathbb{X}$ is the set representing the space of instances $\boldsymbol{x}$ ($\boldsymbol{x} \in \mathbb{X}$), and $\mathbf{x}$ is the random variable that can take on specific values of $\boldsymbol{x} \in \mathbb{X}$. We use this notation consistently, and thus do not always define all symbols explicitly.

A.2. Constraints on our setup

Our setup, per our definition of the learning process (Definition 2.1) is deliberately limited to studying the effects of variance due to changes in the underlying training dataset, with such datasets drawn from the same distribution. Of course, variance-induced error can have other sources in the training pipeline. For example, stochastic optimization methods, such as SGD and Adam, can cause fluctuations in test error; as, too, can choices in hyperparameter configurations (Cooper et al., 2021b). While each of these decision points is worthy of investigation with respect to their impact on fair classification outcomes, we aim to fix as many sources of randomness as possible in order to highlight the particular kind of arbitrariness that we describe in Sections 1 and 3. As such, we use the Limited-memory BFGS solver and fix our hyperparameters based on the results of an initial search (Section 5), for which we selected a search space through consulting related work such as Chen et al. (2018).

A.3. Costs and the classification decision threshold

For reference, we provide a bit more of the basic background regarding the relationship between the classification decision threshold τ and costs of false positives FP (C_{01}) and false negatives FN (C_{10}). We visualize the loss as follows:

Table 1. Confusion matrix for cost-sensitive loss ℓ, adapted from Elkan (2001).

	$\hat{y} = 0$	$\hat{y} = 1$
$o = 0$	TN: 0	FP: C_{01}
$o = 1$	FN: C_{10}	TP: 0

0-1 loss treats the cost of different types of errors equally $C_{01} = C_{10} = 1$); false positives and false negatives are quantified as equivalently bad – they are *symmetric*; the case for which $C_{01} \neq C_{10}$ is *asymmetric* or *cost-sensitive*.

Altering the asymmetric of costs shifts the classification decision threshold τ applied to the underlying regressor $r_{\boldsymbol{D}_k}$. We can see this by examining the behavior of $r_{\boldsymbol{D}_k}$ that we learn. $r_{\boldsymbol{D}_k}$ estimates the probability of a each label given $\boldsymbol{x}$ (since we do not learn using $\boldsymbol{g}$), i.e., that we develop a good approximation of the distribution $p(\mathbf{y}|\mathbf{x})$. Ideally, $r_{\boldsymbol{D}_k}$ will be similar to the Bayes optimal classifier (for which the classification rule produces classifications y^* that yield the smallest weighted sum of the loss, where the weights are the probabilities of a particular label $\mathbf{y} = i$ for a given $(\boldsymbol{x}, \boldsymbol{g})$, i.e., sums over

$$p(\mathbf{y} = i|\mathbf{x} = \boldsymbol{x})\,\ell(i, y'). \tag{5}$$

For binary classification, the terms of (5) in the sum for a particular y' yield two cases:

- $i = y'$: By definition, $\ell(i, y') = 0$; therefore, (5) $= 0$.

- $i \neq y'$: By definition, $\ell(i, y') = C_{01}$ or $\ell(i, y') = C_{10}$. So, (5) will weight the cost by the probability $p(\mathbf{y} = i|\mathbf{x} = \boldsymbol{x})$.

We can therefore break down the Bayes optimal classifier into the following decision rule, which we hope to approximate through learning. For an arbitrary $(\boldsymbol{x}, \boldsymbol{g})$ and $\mathbb{Y} = \{0, 1\}$,

$$\min \Big(\overbrace{\underbrace{p(\mathbf{y}=0|\mathbf{x}=\boldsymbol{x})}_{\text{Probability of FP}} \times C_{01} + \underbrace{p(\mathbf{y}=1|\mathbf{x}=\boldsymbol{x})}_{\text{Probability of TP}} \times 0}^{\text{Weighted cost of predicting positive (1) class}}, \ \overbrace{\underbrace{p(\mathbf{y}=0|\mathbf{x}=\boldsymbol{x})}_{\text{Probability of TN}} \times 0 + \underbrace{p(\mathbf{y}=1|\mathbf{x}=\boldsymbol{x})}_{\text{Probability of FN}} \times C_{10}}^{\text{Weighted cost of predicting negative (0) class}} \Big)$$

$$= \min \Big(\underbrace{p(\mathbf{y}=0|\mathbf{x}=\boldsymbol{x})}_{\text{Probability of FP}} \times C_{10}, \ \underbrace{p(\mathbf{y}=1|\mathbf{x}=\boldsymbol{x})}_{\text{Probability of FN}} \times C_{10} \Big)$$

That is, to predict label 1, the cost of mis-predicting 1 (i.e., the cost of a false positive FP) must be smaller than the cost of mis-predicting 0 (i.e, the cost of a false negative FN). In binary classification $p(\mathbf{y}|\mathbf{x}=\boldsymbol{x}) = p(\mathbf{y}=1|\mathbf{x}=\boldsymbol{x}) + p(\mathbf{y}=0|\mathbf{x}=\boldsymbol{x}) = 1$. So, we can assign $p(\mathbf{y}=1|\mathbf{x}=\boldsymbol{x}) = \tau$ and $p(\mathbf{y}=0|\mathbf{x}=\boldsymbol{x}) = 1 - \tau$, and rewrite the above as

$$\min \Big((1-\tau)\, C_{01}, \ \ \tau C_{10} \Big). \tag{6}$$

The decision boundary is the case for which both of the arguments to $\min$ in (6) are equivalent (i.e., the costs of predicting a false positive and a false negative are equal), i.e.,

$$(1-\tau)\, C_{01} = \tau C_{10} \Rightarrow \tau = \frac{C_{01}}{C_{01} + C_{10}}, \ \text{so,}$$

$$h_{\boldsymbol{D}_k}(\boldsymbol{x}) = \mathbf{1}[r_{\boldsymbol{D}_k}(\boldsymbol{x}) \geq \tau] = \begin{cases} 1, & \text{if } p(\mathbf{y}=1|\mathbf{x}=\boldsymbol{x}) \geq \tau = \frac{C_{01}}{C_{01}+C_{10}} \\ 0, & \text{otherwise.} \end{cases}$$

For 0-1 loss, in which $C_{01} = C_{10} = 1$, τ evaluates to $\frac{1}{2}$. If we want to model asymmetric costs, then we need to change this decision threshold to account for which type of error is more costly. For example, let us say that false negatives are more costly than false positives, with $C_{01} = 1$ and $C_{10} = 3$. This leads to a threshold of $\frac{1}{4}$, which biases $h_{\boldsymbol{D}_k}$ toward choosing the (generally cheaper to predict/more conservative) positive class.

A.4. The bootstrap method

In the bootstrap method, we treat each dataset $\hat{\boldsymbol{D}}_k \in \hat{\mathbb{D}}$ as equally likely. For each set aside test example $(\boldsymbol{x}, \boldsymbol{g}, o)$, we can approximate $\mathrm{Err}(\mathcal{A}, \mathbb{D}, (\boldsymbol{x}, \boldsymbol{g}, o))$ empirically by computing

$$\hat{\mathrm{Err}}\big(\mathcal{A}, \hat{\mathbb{D}}, (\boldsymbol{x}, \boldsymbol{g}, o)\big) = \frac{1}{B} \sum_{i=1}^{B} \ell\big(o, \hat{h}_{\hat{\boldsymbol{D}}_i}(\boldsymbol{x})\big) \tag{7}$$

for a concrete number of replicates B. We estimate overall error $\hat{\mathrm{Err}}(\mathcal{A}, \hat{\mathbb{D}})$ for the test set by additionally summing over each example instance $(\boldsymbol{x}, \boldsymbol{g}, o)$, which we can further delineate into $\hat{\mathrm{FPR}}$ and $\hat{\mathrm{FNR}}$, or into group-specific $\hat{\mathrm{Err}}_{\boldsymbol{g}}$, $\hat{\mathrm{FPR}}_{\boldsymbol{g}}$, and $\hat{\mathrm{FNR}}_{\boldsymbol{g}}$ by computing separate averages according to $\boldsymbol{g}$.

The bootstrap method exhibits less variance than cross-validation, but can be biased — in particular, pessimistic — with respect to estimating expected error. To reduce this bias, one can follow our setup in Definition 2.1, which splits into train and test sets before resampling. For more information comparing the two methods, see Efron and Tibshirani (1997; 1993). Further, recent work shows that, in relation to studying individual models, CV is in fact asymptotically uninformative regarding expected error (Wager, 2020).

B. Additional Details on Variance and Self-Consistency

In this appendix, we provide more details on other types of statistical error (Appendix B.1), on variance (Appendix B.2) and self-consistency (Appendix B.3). Following this longer presentation of our metrics, we then provide some additional

information on other definitions of variance that have been used in work on fair classification, and contextualize issues with these definitions that encouraged us to deviate from them in order to derive our definition of self-consistency (Appendix C).

B.1. Other statistical sources of error

Noise. Noise is traditionally understood as *irreducible error*; it is due to inherent randomness in the data, which cannot be captured perfectly accurately by a deterministic decision rule h_{D_k}. Notably, noise is an aspect of the data collection pipeline, not the learning process (Definition 2.1). It is *irreducible* in the sense that it does not depend on our choice of training procedure $\mathcal{A}$ or how we draw datasets for training from $\mathbb{D}$, either in theory or in practice. Heteroskedastic noise across demographic groups is often hypothesized to be a source of unfairness in machine learning (Cooper and Abrams, 2021; Chen et al., 2018). Importantly, albeit somewhat confusingly, this is commonly referred to as label bias, where "bias" connotes discrimination, as opposed to the statistical bias that we mention here.

Unlike noise, bias and variance are traditionally understood as sources of epistemic uncertainty. These sources of error are *reducible* because they are contingent on the modeling choices we make in the learning process; if we knew how to model the task at hand more effectively, in principle, we could reduce bias and variance error.

Bias. Within the amount of reducible error, bias reflects the error associated with the chosen hypothesis class $\mathbb{H}$, and is therefore governed by decisions concerning the training procedure $\mathcal{A}$ in the learning process (Definition 2.1). This type of error is persistent because it takes effect at the level of possible models in $\mathbb{H}$; in expectation, all models $h_{D_k} \in \mathbb{H}$ have the same amount of bias-induced error.

Whereas variance depends on stochasticity in the underlying training data, noise and bias error are traditionally formulated in relation to the Bayes optimal classifier — the best possible classifier that machine learning could produce for a given task (Geman et al., 1992; Abu-Mostafa et al., 2012; Domingos, 2000a). Since the Bayes optimal classifier is typically not available in practice, we often cannot estimate noise or bias directly in experiments.

Of the three types of statistical error, it is only variance that seems to reflect the intuition in Figure 1 concerning the behavior of different possible models h_{D_k}. This is because noise is a property of the data distribution; for a learning process (Definition 2.1), in expectation we can treat noise error as constant. Bias can similarly be treated as constant for the learning process: It is a property of the chosen hypothesis class $\mathbb{H}$, and thus is in expectation the same for each $h_{D_k} \in \mathbb{H}$. In Figure 1, we are keeping the data distribution constant and $\mathbb{H}$ constant; we are only changing the underlying subset of training data to produce different models h_{D_k}.

B.2. Our variance definition

We first provide a simple proof that explains the simplified version for our empirical approximation for variance in (1).

Proof. For the models $\{h_{D_b}\}_{b=1}^{B}$ that we produce, we denote $\hat{\mathbb{Y}}$ to be the multiset of their predictions on $(\boldsymbol{x}, \boldsymbol{g})$. $|\hat{\mathbb{Y}}| = B = B_0 + B_1$, where B_0 and B_1 represent the counts of 0 and 1-predictions, respectively. We also set the cost of false positives to be $\ell(0,1) = C_{01}$ and the cost of false negatives to be $\ell(1,0) = C_{10}$.

Looking at the sum in $\hat{\text{var}}$ (i.e., $\sum_{i \neq j}$), each of the B_0 0-predictions will get compared to the other $B_0 - 1$ 0-predictions and to the B_1 1-predictions. By the definition of ℓ, each of the $B_0 - 1$ computations of $\ell(0,0)$ evaluates to 0 and each of the B_1 computations of $\ell(0,1)$ evaluates to C_{01}. Therefore, the B_0 0-predictions contribute

$$B_0 \times \left[\left(0 \times (B_0 - 1)\right) + C_{01} \times B_1 \right] = C_{01} B_0 B_1$$

to the sum in $\hat{\text{var}}$, and, by similar reasoning, $B_1 \times \left[\left(0 \times (B_1 - 1)\right) + C_{10} \times B_0 \right] = C_{10} B_0 B_1$. It follows that the total sum in $\hat{\text{var}}$ is

$$\sum_{i \neq j} \ell\left(\hat{h}_{\hat{\boldsymbol{D}}_i}(\boldsymbol{x}), \hat{h}_{\hat{\boldsymbol{D}}_j}(\boldsymbol{x})\right) = (C_{01} + C_{10}) B_0 B_1. \text{ Therefore}$$

$$\overbrace{\frac{1}{B(B-1)} \sum_{i \neq j} \ell\left(\hat{h}_{\hat{\boldsymbol{D}}_i}(\boldsymbol{x}), \hat{h}_{\hat{\boldsymbol{D}}_j}(\boldsymbol{x})\right)}^{\hat{\text{var}}\left(\mathcal{A}, \hat{\mathbb{D}}, (\boldsymbol{x}, \boldsymbol{g})\right)} = \overbrace{\frac{(C_{01} + C_{10}) B_0 B_1}{B(B-1)}}^{(1)}$$

$\square$

The effect of τ on variance. As discussed in Appendix A.3, C_{01} and C_{10} can be related to changing τ applied to $r_{\boldsymbol{D}_k}$ to produce classifier $h_{\boldsymbol{D}_k}$. We analyze the range of minimal and maximal empirical variance by examining the behavior of $B \to \infty$, i.e.,

$$\lim_{B \to \infty} \frac{(C_{01} + C_{10})B_0 B_1}{B(B-1)}. \tag{8}$$

Minimal variance. Either B_0 or B_1 (exclusively, since $B_0 + B_1 > 1$) will be $= 0$, with the other being $= B$, making (8) equivalent to

$$\lim_{B \to \infty} \frac{(C_{01} + C_{10}) \times 0}{B(B-1)} = 0, \text{regardless of the value of } C_{01} + C_{10}.$$

Maximal variance. B_0 will represent half of B, with B_1 representing the other half. More particularly, $B_0 = \frac{B}{2}$ and $B_1 = \frac{B}{2}$; or, without loss of generality, $B_0 = \frac{B-1}{2}$ and $B_1 = \frac{B+1}{2}$. This means that

$$\frac{(C_{01} + C_{10})B_0 B_1}{B(B-1)} = \frac{(C_{01} + C_{10})(\frac{B}{2})^2}{B(B-1)} \qquad \left(\text{Or,} = \frac{(C_{01} + C_{10})(\frac{B-1}{2})(\frac{B+1}{2})}{B(B-1)} \right)$$

$$= \frac{(C_{01} + C_{10})(\frac{B^2}{4})}{B^2 - B} \qquad \left(\text{Or,} = \frac{(C_{01} + C_{10})(\frac{(B^2-1)}{4})}{B(B-1)}; \text{it will not matter in the limit} \right)$$

$$= \frac{(C_{01} + C_{10})B^2}{4B^2 - 4B}.$$

And, therefore,

$$\lim_{B \to \infty} \frac{(C_{01} + C_{10})B^2}{4B^2 - 4B} = \frac{C_{01} + C_{10}}{4}. \tag{9}$$

It follows analytically that variance will be in the range $[0, \frac{C_{01}+C_{10}}{4})$. However, empirically, for concrete B, $\mathrm{v\hat{a}r}(\mathcal{A}, \hat{\mathbb{D}}, (\boldsymbol{x}, \boldsymbol{g})) \to [0, \frac{C_{01}+C_{10}}{4} + \epsilon]$, for smaller positive ϵ as the number of models B increases. The maximal variance will better approximate $\frac{C_{01}+C_{10}}{4}$ as B gets larger, but will be $> \frac{C_{01}+C_{10}}{4}$. For example, for 0-1 loss $\frac{C_{01}+C_{10}}{4} = \frac{2}{4} = 0.5$. For $B = 100$, the maximal $\mathrm{v\hat{a}r}(\mathcal{A}, \hat{\mathbb{D}}, (\boldsymbol{x}, \boldsymbol{g})) = \frac{2 \times 50 \times 50}{100 \times 99} = \frac{50}{99} \approx .505$.

B.3. Deriving self-consistency from variance

In this appendix, we describe the relationship between variance (Definition 3.1) and self-consistency (Definition 3.2) in more detail, and show that $\mathrm{s\hat{c}}(\mathcal{A}, \{\boldsymbol{D}_b\}_{b=1}^B, (\boldsymbol{x}, \boldsymbol{g})) \to [0.5 - \epsilon, 1]$ for small positive ϵ as the number of models B increases.

Proof. Note that, by the definition of 0-1 loss, $C_{01} = C_{10} = 1$, so

$$\mathrm{v\hat{a}r}(\mathcal{A}, \hat{\mathbb{D}}, (\boldsymbol{x}, \boldsymbol{g}))_{\text{0-1}} = \frac{1}{B(B-1)} \sum_{i \neq j} \mathbf{1}[h_{\boldsymbol{D}_i}(\boldsymbol{x}) \neq h_{\boldsymbol{D}_j}(\boldsymbol{x})] = \frac{2B_0 B_1}{B(B-1)}. \tag{10}$$

By the definition of of the indicator function $\mathbf{1}$,

$$1 = \frac{1}{B(B-1)} \sum_{i \neq j} \left[\overbrace{\mathbf{1}[h_{\boldsymbol{D}_i}(\boldsymbol{x}) \neq h_{\boldsymbol{D}_j}(\boldsymbol{x})]}^{\text{From } \mathrm{v\hat{a}r}(\mathcal{A}, \hat{\mathbb{D}}, (\boldsymbol{x}, \boldsymbol{g}))_{\text{0-1}}} + \overbrace{\mathbf{1}[h_{\boldsymbol{D}_i}(\boldsymbol{x}) = h_{\boldsymbol{D}_j}(\boldsymbol{x})]}^{\text{From } \mathrm{s\hat{c}}(\mathcal{A}, \{\hat{\boldsymbol{D}}_b\}_{b=1}^B, (\boldsymbol{x}, \boldsymbol{g}))} \right]$$

$$= \overbrace{\frac{2B_0 B_1}{B(B-1)}}^{\text{(10)}} + \frac{1}{B(B-1)} \sum_{i \neq j} \mathbf{1}[h_{\boldsymbol{D}_i}(\boldsymbol{x}) = h_{\boldsymbol{D}_j}(\boldsymbol{x})]. \text{Therefore, rearranging}$$

$$\hat{\text{SC}}\big(\mathcal{A}, \hat{\mathbb{D}}, (\boldsymbol{x}, \boldsymbol{g})\big) = \frac{1}{B(B-1)} \sum_{i \neq j} \mathbf{1}[h_{\boldsymbol{D}_i}(\boldsymbol{x}) = h_{\boldsymbol{D}_j}(\boldsymbol{x})] = 1 - \frac{2B_0 B_1}{B(B-1)}.$$

$\square$

We note that $\hat{\text{SC}}$ (2) is independent of specific costs C_{01} and C_{10}. Nevertheless, the choice of decision threshold τ will of course impact the values of B_0 and B_1 in practice. In turn, this will impact the degree of self-consistency that a learning process exhibits empirically. In short, the measured degree of self-consistency in practice will depend on the choice of ℓ. Further, following an analysis similar to what we can show that $\hat{\text{SC}}$ will be a value in $[0.5 + \epsilon, 1]$, for small positive ϵ. This reality is reflected in the results that we report for our experiments, for which $B = 101$ yields minimal $\hat{\text{SC}} \approx 0.495$.

Cost-independence of self-consistency Intuitively, *self*-consistency of a learning process is a relative metric; it is a quantity that is measured relative to the learning process. We therefore conceive of it as a metric that is normalized with respect to the learning process (Definition 2.1). Such a process can be maximally 100% self-consistent, but it does not make sense for it to be more than that (reflected by the maximum value of 1).

In contrast, as discussed in Appendix B, variance can measure much greater than 1, depending on the magnitude of the sum of the costs C_{01} and C_{10}, in particular, for $C_{01} + C_{10} > 4$ (9). However, it is not necessarily meaningful to compare the magnitude of variance across classifiers. Recall that the effect of changing costs C_{01} and C_{10} corresponds to a change in the binary classification decision threshold, with $\tau = \frac{C_{01}}{C_{01}+C_{10}}$. It is the *relative* costs that change the decision threshold; not the costs themselves. For example, the classifier with costs $C_{01} = 1$ and $C_{10} = 3$ is equivalent to the classifier with costs $C_{01} = \frac{1}{2}$ and $C_{10} = \frac{3}{2}$ (for both, $\tau = \frac{1}{4}$), but the former would measure a larger magnitude for variance.

It is this observation that grounds our cost-independent definition of self-consistency in Section 3 and Appendix B.3. Given the fact that the magnitude of variance measurements can complicate our comparisons of classifiers, as discussed above, we focus on the part of variance that encodes information about arbitrariness in a learning process: its measure of (dis)agreement between classification decisions that result from changing the training dataset. We could alternatively conceive of self-consistency as the additive inverse of normalized variance, but this is more complicated because it would require a computation that depends on the specific costs, $\hat{\text{var}}\big(\mathcal{A}, \hat{\mathbb{D}}, (\boldsymbol{x}, \boldsymbol{g})\big)_{\text{normalized}} = \frac{\hat{\text{var}}\big(\mathcal{A}, \hat{\mathbb{D}}, (\boldsymbol{x}, \boldsymbol{g})\big)}{\hat{\text{var}}\big(\mathcal{A}, \hat{\mathbb{D}}, (\boldsymbol{x}, \boldsymbol{g})\big)_{\text{max}}}.$

B.3.1. ADDITIONAL DETAILS ON OUR CHOICE OF SELF-CONSISTENCY METRIC

Terminology. In logic, the idea of consistent belief has to do with ensuring that we do not draw conclusions that contradcit each other. This is much like the case that we are modeling with self-consistency — the idea that underlying changes in the dataset can lead to predictions that are directly in contradition (Smullyan, 1986; Hintikka, 1962; Stalnaker, 2006). Ideas of consistency in legal rules have a similar flavor; legal rules should not contradict each other; legal judgments should not contradict each other (this is at least an aspiration for the law, based on common ideas in legal theory (Fuller, 1965; Tamanaha, 2004). For both of these reasons, the term "consistent" has a natural mapping to our usage of it in this paper. This is especially true in the legal theory case, given that inconsistency in the law is often considered arbitrary and a source of discrimination.

We nevertheless realize that the word "consistent" is overloaded with many meanings in statistics and different subfields computer science like distributed computing (Zhang et al., 2020; Abadi, 2012, e.g.,). Nevertheless, due to the clear relationship between our purposes concerning arbitrariness and discrimination, and definitions in logic and the law, we believe that it is the most appropriate term for our work.

Quantifying systematic arbitrariness. We depict *systematic arbitrariness* using the Wasserstein-1 distance (Ramdas et al., 2015). This is the natural distance for us to consider because it has a closed form when being applied to CDFs. For our purposes, it should be interpreted as computing the total disparity in self-consistency by examining all possible self-consistency levels κ at once.

Formally, for two groups $\boldsymbol{g} = 0$ and $\boldsymbol{g} = 1$ with respective SC CDFs F_0 and F_1,

$$\mathcal{W}_1 = \int_{\mathbb{R}} |F_0(\kappa) - F_1(\kappa)| \, d\kappa.$$

For self-consistency, which we have defined on $[0.5, 1]$, this is just

$$\mathcal{W}_1 = \int_{0.5}^{1} |F_0(\kappa) - F_1(\kappa)| \, d\kappa.$$

Empirically, we can approximate this with

$$\hat{\mathcal{W}}_1 := \frac{1}{|\hat{\mathbb{K}}|} \sum_{\hat{\mathbb{K}}} |\hat{F}_0(\hat{\kappa}) - \hat{F}_1(\hat{\kappa})|, \quad \text{where } \hat{\mathbb{K}} = \left\{ 1 - \frac{2B_0 B_1}{B(B-1)} \middle| B_0 \in \{0 \ldots B\} \wedge B_1 \in \{0 \ldots B\} \wedge B_0 + B_1 = B \right\}.$$

We typically set $B = 101$, and thus

$$
\begin{aligned}
\hat{\mathbb{K}} = [&0.49505, 0.49545, 0.49624, 0.49743, 0.49901, 0.50099, 0.50337, 0.50614, 0.50931, 0.51287, 0.51683, 0.52119, \\
&0.52594, 0.53109, 0.53663, 0.54257, 0.54891, 0.55564, 0.56277, 0.57030, 0.57822, 0.58653, 0.59525, 0.60436, \\
&0.61386, 0.62376, 0.63406, 0.64475, 0.65584, 0.66733, 0.67921, 0.69149, 0.70416, 0.71723, 0.73069, 0.74455, \\
&0.75881, 0.77347, 0.78851, 0.80396, 0.81980, 0.83604, 0.85267, 0.86970, 0.88713, 0.90495, 0.92317, 0.94178, \\
&0.96079, 0.9802, 1.0],
\end{aligned}
$$

which we use to produce our CDF plots.

When measuring systematic arbitrariness with abstention, we set the probability mass for $< \kappa$ to 0 it. This makes sense because we are effectively re-defining the $\hat{\text{SC}}$ CDFs to not include instances that exhibit below a minimal amount of $\hat{\text{SC}}$. This also makes comparing systematic arbitrariness across CDFs for different interventions more interpretable. It allows us to keep the number of empirical samples for the empirical CDF measures constant when computing averages, so abstaining would then always have the effect of decreasing systematic arbitrariness. If we did not do this, because the Wasserstein-1 distance is an average, changing the set $\hat{\mathbb{K}}$, of course, would change the amount of Wasserstein-1 distance — possibly leading to a relative *increase* (if there are greater discrepancies between g-condition CDF curves at $\geq \kappa$).

C. Related Work and Alternative Notions of Variance

As noted in Section 6, prior work that discusses variance and fair classification often relies on the definition of variance from Domingos (2000a). We deviate from prior work and provide our own definition for two reasons: 1) variance in Domingos (2000a;b) does not cleanly extend to cost-sensitive loss, and 2) the reference point for measuring variance in Domingos (2000a;b) — the *main prediction* — can be unstable/ brittle in practice. We start by explaining the Domingos (2000a;b) definitions, and then use these definitions to support our rationale.

C.1. Defining variance in relation to a "main prediction"

To begin, we restate the definitions from Domingos (2000a;b) concerning the expected model (called the *main predictor*). We change the notation from Domingos to align with our own, as we believe these changes provide greater clarity concerning meaning, significance, and consequent takeaways. Nevertheless, these definitions for quantifying error are equivalent to those in Domingos (2000b), and they fundamentally depend on human decisions for setting up the learning process.

Domingos (2000a;b) define predictive variance in relation to this single point of reference. This reference point captures the general, expected behavior of models that could be produced by the chosen learning process. We can think of each prediction of this point of reference as the "central tendency" of the predictions made by all possible models in μ for $(\boldsymbol{x}, \boldsymbol{g})$. Formally,

Definition C.1. The **main prediction** $\hat{y}$ is the prediction value $y' \in \mathbb{Y}$ that generates the minimum average loss with respect to all of the predictions $\hat{y} \in \hat{\mathbb{Y}}$ generated by the different possible models in μ. It is defined as the expectation over training sets $\mathbb{D}$ for a loss function ℓ, given an example instance $(\boldsymbol{x}, \boldsymbol{g})$. That is,

$$\overline{y} = \underset{y'}{\arg \min} \, \mathbb{E}_{\mathbf{D}}[\ell(\hat{y}, y') | \mathbf{x} = \boldsymbol{x}, \mathbf{g} = \boldsymbol{g}]. \tag{11}$$

The *main predictor* $\overline{h} : \mathbb{X} \to \mathbb{Y}$ produces the main prediction $\overline{y}$ for each $(\boldsymbol{x}, \boldsymbol{g})$.

What (11) evaluates to in practice of course depends on the loss function ℓ. For squared loss, the main prediction is defined as the mean prediction of all the $h_{\boldsymbol{D}_k}$ (Domingos, 2000a; Kong and Dietterich, 1995). Following Kong and Dietterich (1995), for 0-1 loss Domingos (2000a) defines the main prediction as the mode/majority vote — the most frequent prediction for an example instance $(\boldsymbol{x}, \boldsymbol{g})$. We provide a more formal discussion of why this is the case when we discuss problems with the main prediction for cost-sensitive loss (Appendix C.2). Domingos (2000a;b) then define variance in relation to specific models $h_{\boldsymbol{D}_k}$ and the main predictor $\overline{h}$:

Definition C.2. The *variance*-induced error for fresh example instance $(\boldsymbol{x}, \boldsymbol{g})$ is

$$\text{var}(\mathcal{A}, \mathbb{D}, (\boldsymbol{x}, \boldsymbol{g})) = \mathbb{E}_{\mathbf{D}}[\ell(\overline{y}, \hat{y})|\mathbf{x} = \boldsymbol{x}, \mathbf{g} = \boldsymbol{g}],$$

where $\overline{y} = \overline{h}(\boldsymbol{x})$ is the main prediction and the $\hat{y}$ are the predictions for the different $h_{\boldsymbol{D}_k} \sim \mu$.

That is, for a specific $(\boldsymbol{x}, \boldsymbol{g})$, it is possible to compare the individual predictions $\hat{y} = h_{\boldsymbol{D}_k}(\boldsymbol{x})$ to the main prediction $\overline{y} = \overline{h}(\boldsymbol{x})$. Using the main prediction as a reference point, one can compute the extent of disagreement of individual predictions with the main prediction as a source of error. It is this definition (Definition C.2) that prior work on fair classification tends to reference when discussing variance (Chen et al., 2018; Black et al., 2022a). However, as we discuss in more detail below (Appendix C.2), many of the theoretical results in Chen et al. (2018) follow directly from the definitions in Domingos (2000a), and the empirical work does not actually use those results in practice. Black et al. (2022a), in contrast, presents results that rely heavily on the main prediction in Domingos (2000a).

C.2. Why we choose to avoid computing the main prediction

We now compare our definition of variance (Definition 3.1) to the one in Domingos (2000a;b) (Definition C.2). This comparison makes clear in detail why we deviate from prior work that relies on Domingos (2000a;b).

No decomposition result. Following from above, it is worth noting that by not relying on the main prediction, we lose the applicability of the decomposition result that Domingos (2000a;b) develop. However, we believe that this is fine for our purposes, as we are interested in the impact of empirical variance specifically on fair classification outcomes. We do not need to reason about bias or noise in our results to understand the arbitrariness with which we are concerned (Section 3.1). It is also worth noting that prior work on fair classification that leverages Domingos (2000a) also does not leverage the decomposition, either. Chen et al. (2018) extends the decomposition to subgroups in the context of algorithmic fairness,[9] and then informally translates the takeaways of the Domingos (2000a) result to a notion of a "level of discrimination." Moreoever, unlike our work, these prior studies do not actually measure variance directly in its experiments.

No need to compute a "central tendency". In Domingos (2000a;b), variance is defined in terms of both the loss function ℓ and the main prediction $\overline{y}$. This assumes that the main prediction is well-defined for the loss function, and that it is well-behaved. While there is a simple interpretation of the main prediction for squared loss (the mean) and for 0-1 loss (the mode/majority vote), it is significantly messier for cost-sensitive loss, which is a more general formulation that includes 0-1 loss. Domingos (2000a;b) does not discuss this explicitly, so we derive the main prediction for cost-sensitive loss ourselves below. In summary:

- The behavior of the main prediction for cost-sensitive loss reveals that the decomposition result provided in the extended technical report (Theorem 4, Domingos (2000b)) is in fact very carefully constructed. We believe that this construction is so specific that it is not practically useful (it is, in our opinion, hardly "unified" in a more general sense, as it is so carefully adapted to specific loss functions and their behavioral special cases).

- By decoupling from the need to compute a main prediction as a reference point, our variance definition is ultimately much simpler and more general, with respect to how it accommodates different loss functions.[10]

---

[9]This just involves splitting the conditioning on an example instance of features $\boldsymbol{x}$ into conditioning on an example instance whose features are split into $(\boldsymbol{x}, \boldsymbol{g})$.

[10]This reveals a subtle ambiguity in the definition of the loss ℓ in Domingos (2000a;b). Neither paper explicitly defines the signature of ℓ. For the main prediction (Definition C.1) and variance (Definition C.2), there is a lack of clarity in what constitutes a valid domain for ℓ. Computing the main prediction $\overline{y}$ suggests $\ell : \mathbb{Y} \times \mathbb{Y} \to \mathbb{R}_{\geq 0}$, where $\overline{y} \in \mathbb{Y}$, but, since $\hat{\mathbb{Y}} \subseteq \mathbb{Y}$, it is possible that $\overline{y} \notin \mathbb{Y}$. However, the definition of variance suggests that $\ell : \mathbb{Y} \times \hat{\mathbb{Y}} \to \mathbb{R}_{\geq 0}$. Since $\hat{\mathbb{Y}} \subseteq \mathbb{Y}$, it is not guaranteed that $\hat{\mathbb{Y}} = \mathbb{Y}$. This may be fine in practice, especially for squared loss and 0-1 loss (the losses with which Domingos (2000a) explicitly contends), but it does arguably present a problem formally with respect to generalizing.

Brittleness of the main prediction. For high variance instances, the main prediction can flip-flop from $\hat{y} = 1$ to $\hat{y} = 0$ and back. While the strategy in Black et al. (2022a) is to abstain on the prediction in these cases, we believe that a better alternative is to understand that the main prediction is not very meaningful more generally for high-variance examples. That is, for these examples, the ability (and reliability) of breaking close ties to determine the main (simple majority) prediction is not the right approach. Instead, we should ideally be able to embed more confidence into our process than a simple-majority-vote determination.[11] Put different, in cases for which we can reliably estimate the main prediction, but the vote margin is slim, we believe that the main prediction is still uncertain, based on our understanding of variance, intuited in Figure 1. **The main prediction can be reliable, but it can still, in this view, be arbitrary** (Section 6). With a simple-majority voting scheme, there can be huge differences between predictions that are mostly in agreement, and those that are just over the majority reference point. Freeing ourselves of this reference point via our self-consistency metric, we can define thresholds of self-consistency as our criterion for abstention (where simple-majority voting is one instantiation of that criterion).[12]

C.2.1. THE MAIN PREDICTION AND COST-SENSITIVE LOSS

We show here that, for cost-sensitive loss, the main prediction depends on the majority class being predicted, the asymmetry of the costs, and occasional tie-breaking, such that the main prediction can either be the majority vote or the minority vote. Domingos (2000b) provides an error decomposition in Theorem 4, but does not explain the effects on the main prediction. We do so below, and also call attention to 0-1 loss as a special case of cost-sensitive loss, for which the costs are symmetric (and equal to 1). We first summarize the takeaways of the analysis below:

- **Symmetric loss**: The main prediction is the **majority vote**.

- **Asymmetric loss**: Compute 1) the relative cost difference (i.e., $\frac{C_{01} - C_{10}}{C_{10}}$), 2) the majority class (and, as a result, the minority class) for the $\hat{y} \in \hat{\mathbb{Y}}$, and 3) the relative difference in the number of votes in the majority and minority classes (i.e., what we call the *vote margin*; below, $\frac{(i+2j+1)-i}{i}$)

  - If the **majority class** in $\hat{\mathbb{Y}}$ has the **lower cost** of misclassification, then the main prediction is the **majority vote**.
  - If the **majority class** in $\hat{\mathbb{Y}}$ has the **higher cost** of misclassification, then the main prediction **depends on the asymmetry of the costs and the vote margin**, i.e.,
    * If $\frac{C_{01} - C_{10}}{C_{10}} = \frac{(i+2j+1)-i}{i}$, we can choose the main prediction to be **either class** (but must make this choice consistently).
    * If $\frac{C_{01} - C_{10}}{C_{10}} > \frac{(i+2j+1)-i}{i}$, the **minority vote** is the main prediction.
    * If $\frac{C_{01} - C_{10}}{C_{10}} < \frac{(i+2j+1)-i}{i}$, the **majority vote** is the main prediction.

Proof. Let us consider cost-sensitive loss for binary classification, for which $\ell(0,0) = \ell(1,1) = 0$ and we have potentially-asymmetric loss for misclassifications, i.e. $\ell(1,0) = C_{10}$ and $\ell(0,1) = C_{01}$, with $C_{01}, C_{10} \in \mathbb{R}^+$. 0-1 loss is a special case for this type of loss, for which $C_{01} = C_{10} = 1$.

Let us say that the total number of models trained is k, which we evaluate on an example instance $\boldsymbol{x}$. Let us set $|\hat{\mathbb{Y}}| = k = 2i + 2j + 1$, with $i \geq 0$ and $j \geq 0$. We can think of i as the common number of votes that each class has, and $2j + 1$ as the margin of votes between the two classes. Given this setup, this means that $k \geq 1$, i.e., we always have the predictions of at least 1 model to consider, and k is always odd. This means that there is always a strict majority classification.

Without loss of generality, on $\boldsymbol{x}$, of these k model predictions $\hat{y} \in \hat{\mathbb{Y}}$, there are i class-0 predictions and $i + 2j + 1$ class-1 predictions (i.e., we do our analysis with class 1 as the majority prediction). To compute the main prediction $\overline{y}$, each $\hat{y} \in \hat{\mathbb{Y}}$ will get compared to the values of possible predictions $y' \in \mathbb{Y} = \{0, 1\}$. That is, there are two cases to consider:

- **Case** $y' = 0$: $y' = 0$ will get compared i times to the i $\hat{y} = 0$s in $\hat{\mathbb{Y}}$, for which $\ell(0,0) = 0$; $y' = 0$ will similarly get compared $i + 2j + 1$ times to the 1s in in $\hat{\mathbb{Y}}$, for which (by Definition C.1) the comparison is $\ell(1,0) = C_{10}$. By

---

[11]This is also another aspect of the simplicity of not needing to define and compute a "central tendency" prediction. We do not need to encode a notion of a tie-breaking vote to determine a "central tendency." The main prediction can be unclear in cases for which there is no "main outcome" (e.g., Individual 2 in Figure 1), as the vote is split exactly down the middle. By avoiding the need to vote on a main reference point, we also avoid having to ever choose that reference point arbitrarily.

[12]This problem is worse for cost-sensitive loss, where the main prediction is not always the majority vote (Appendix C.2.1).

definition of expectation, the expected loss is

$$\frac{i \times 0 + (i + 2j + 1) \times C_{10}}{2i + 2j + 1} = \frac{C_{10}(i + 2j + 1)}{2i + 2j + 1}. \tag{12}$$

- **Case** $y' = 1$: Similarly, the label 1 will also get compared i times to the 0s in $\hat{\mathbb{Y}}$, for which the comparison is $\ell(0, 1) = C_{01}$; $y' = 1$ will also be compared $i + 2j + 1$ times to the 1s in $\hat{\mathbb{Y}}$, for which $\ell(1, 1) = 0$. The expected loss is

$$\frac{i \times C_{01} + (i + 2j + 1) \times 0}{2i + 2j + 1} = \frac{C_{01}i}{2i + 2j + 1}. \tag{13}$$

We need to compare these two cases for different possible values of C_{10} and C_{01} to understand which expected loss is minimal, which will determine the main prediction $\overline{y}$ that satisfies Equation (11). The three different possible relationships for values of C_{10} and C_{01} are $C_{10} = C_{01}$ (symmetric loss), and $C_{10} > C_{01}$ and $C_{10} < C_{01}$ (asymmetric loss). Since the results of the two cases above share the same denominator, we just need to compare their numerators, $C_{10}(i + 2j + 1)$ (12) and $C_{01}i$ (13).

Symmetric Loss (0-1 Loss). When $C_{10} = C_{01} = 1$, the numerators in (12) and (13) yield expected losses $i + 2j + 1$ and i, respectively. We can rewrite the numerator for (13) as

$$i + \overbrace{2j + 1}^{\geq 1, \text{ given } j \geq 0} \geq i + 1,$$

which makes the comparison of numerators $i < i + 1$, i.e., we are in the case (13) < (12). This means that the case of $y' = 1$ (13) is the minimal one; the expected loss for class 1, the most frequent class, is the minimum, and thus the most frequent/ majority vote class is the main prediction. An analogous result holds if we instead set the most frequent class to be 0. More generally, this holds for all symmetric losses, for which $C_{10} = C_{01}$.

▶ For **symmetric losses**, the main prediction $\overline{y}$ is **majority vote** of the predictions in $\hat{\mathbb{Y}}$.

Asymmetric Loss. For asymmetric/ cost-sensitive loss, we need to examine two sub-cases: $C_{10} > C_{01}$ and $C_{10} < C_{01}$.

- **Case $C_{10} > C_{01}$**: $C_{01}i < C_{10}(i + \overbrace{2j + 1}^{\geq 1})$, given that $j \geq 0$. Therefore, since $C_{01}i$ is minimal and associated with class 1 (the most frequent class in our setup), the majority vote is the main prediction. We can achieve an analogous result if we instead set 0 as the majority class.

  ▶ For **asymmetric losses**, the main prediction $\overline{y}$ is the **majority vote** of the predictions in $\hat{\mathbb{Y}}$, **if the majority class has a cheaper cost associated with misclassification** (i.e., if the majority class is 1 and $C_{10} < C_{01}$, or if the majority class is 0 and $C_{01} < C_{10}$).

- **Case $C_{10} < C_{01}$**: If $C_{10} < C_{10}$, it depends on how asymmetric the costs are and how large the vote margin (i.e., $2j + 1$) between class votes is. There are 3 sub-cases:

  - **Case $C_{01}i = C_{10}(i + 2j + 1)$, i.e. cost equality**: We can look at the relative asymmetric cost difference of the minority class cost (above C_{01}, without loss of generality) and the majority class cost (above C_{10}, without loss of generality), (above $\frac{C_{01} - C_{10}}{C_{10}}$, without loss of generality). If that relative cost difference is equal to the relative difference of the votes between the majority and minority classes (i.e., $\frac{(i + 2j + 1) - i}{i}$), then the costs of predicting

either 1 or 0 are equal. That is, we can rearrange terms as a ratio of costs to votes:

$$C_{01}i = C_{10}(i + \overbrace{2j + 1}^{\geq 1}) \qquad\qquad \text{(The terms in this equality are} > 0)$$

$$\frac{C_{01}}{C_{10}} = \frac{i + 2j + 1}{i} \qquad\qquad \text{(Given the above, } C_{01}i > 0 \text{ so } i > 0)$$

$$= 1 + \frac{2j + 1}{i}$$

$$\frac{C_{01}}{C_{10}} - 1 = \frac{2j + 1}{i}$$

$$\frac{C_{01} - C_{10}}{C_{10}} = \frac{2j + 1}{i} = \frac{(i + 2j + 1) - i}{i} \geq \frac{1}{i} \tag{14}$$

▶ For asymmetric loss **when the majority-class-associated cost is less than the minority-class associated cost and if the expected losses are equal**, then the **main prediction $\overline{y}$ is either 1 or 0**, (and we must make this choice consistently).

− **Case $C_{01}i > C_{10}(i + 2j + 1)$**: We can look at the relative asymmetric cost difference of the minority class cost (above C_{01}, without loss of generality) and the majority class cost (above C_{10}, without loss of generality), (above $\frac{C_{01} - C_{10}}{C_{10}}$, without loss of generality). If that relative cost difference is greater than the relative difference of the votes between the majority and minority classes (i.e., $\frac{(i+2j+1)-i}{i}$), then the *minority vote* yields the minimum cost and is the main prediction $\overline{y}$ (above $\overline{y} = 0$, without loss of generality; an analogous result holds if we had set the majority vote to be 0 and the minority vote to be 1). Following (14) above, this is the same as

$$\frac{C_{01} - C_{10}}{C_{10}} > \frac{(i + 2j + 1) - i}{i}$$

▶ For asymmetric loss **when the majority-class-associated cost is less than the minority-class associated cost**, it is possible for the **minority class** to have a greater associated loss. In this case, the *minority vote* **is the main prediction $\overline{y}$**.

− **Case $C_{01}i < C_{10}(i + 2j + 1)$**: We can look at the relative asymmetric cost difference of the minority class cost (above C_{01}, without loss of generality) and the majority class cost (above C_{10}, without loss of generality), (above $\frac{C_{01} - C_{10}}{C_{10}}$, without loss of generality). If that relative cost difference s less than the relative difference of the votes between the majority and minority classes (i.e., $\frac{(i+2j+1)-i}{i}$), then the majority vote yields to minimum cost and is the main prediction $\overline{y}$ (above $\overline{y} = 1$, without loss of generality; an analogous result holds if we had set the majority vote to be 0 and the minority vote to be 1). Following (14) above, this is the same as

$$\frac{C_{01} - C_{10}}{C_{10}} < \frac{(i + 2j + 1) - i}{i}$$

▶ For asymmetric loss **when the majority-class-associated cost is less than the minority-class associated cost**, it is possible for the **majority class** to have a greater associated loss. In this case, the *majority vote* **is the main prediction $\overline{y}$**.

□

D. Additional Details on Our Algorithmic Framework

A natural question is to see if we can improve self-consistency, with the hope that doing so would reduce arbitrariness in the learning process, improve accuracy, and, for the cases in which there is different self-consistency across subgroups, also perhaps improve fairness. To do so, we consider ways of reducing variance, as, based on our definitions (Definition 3.1 and 3.2), doing so should improve self-consistency.

We consider the classic *bootstrap agg*regation — or, *bagging* — algorithm (Breiman, 1996) as a starting point. It has been well-known since Breiman (1996) that *bagging* can improve the performance of unstable predictors. That is, for models produced by a learning process that is sensitive to the underlying training data, it is (theoretically-grounded) good practice to train an ensemble of models using bootstrapping (Appendix A.4; Efron (1979); Efron and Tibshirani (1993)). When classifying an example instance, we then leverage the whole ensemble by aggregating the predictions produced by its members. This aggregation process identifies the most common prediction in the ensemble, and returns that label as the classification. Put differently, we have combined the information of a lot of unstable classifiers, and averaged over their behavior in order to generate more stable classifications.

Given the the relationship between variance (Definition 3.1) and self-consistency (Definition 3.2), reducing variance will improve self-consistency. However, rather than relying on a simple-majority-vote to decide the aggregated prediction, we also will instill a notion of confidence in our predictions by requiring a minimum level of self-consistency, which is described in Algorithm 1.

D.1. Self-consistent ensembling with abstention

We present a framework that alters the semantics of classification outputs to 0, 1, and `Abstain`, and employ ensembling to determine the $\hat{\text{SC}}$-level that guides the output process. Algorithm 1 suggests many possible ways that this can be achieved. For example, we can change the aggregation rule in regular bagging to use a self-consistency level κ rather than majority vote. Instead of relying on votes, we can bag the underlying prediction probabilities and then apply κ a filter. We could take the top-n most consistent predictions and let a super-ensemble of underlying bagged classifiers decide whether to abstain or predict.

In the experiments in the paper, we provide two examples: Changing the underlying bagging vote aggregation rule (simple ensembling), and applying a round of regular bagging to do variance reduction and then bagging the bagged outputs (super ensembling) to apply a self-consistency threshold. Our ensemble model will not produce predictions for examples for which the lack of self-consistency is too high. We describe our procedure more formally in Algorithm 1.

Simple proof that abstention improves self-consistency (by construction). We briefly show the simple proof that any method that meets the semantics of Algorithm 1 will be more self-consistent than its counterpart that cannot `Abstain`.

We define abstentions to be in agreement with both 0 and 1 predictions. This makes sense intuitively: Algorithm 1 abstains to avoid making predictions that lack self-consistency, so abstaining should not increase disagreement between predictions.

It follows that we can continue to use Definition 3.2 and associated empirical approximations $\hat{\text{SC}}$ (3), but with one small adjustment. Instead of the total number of predictions $B = B_0 + B_1$, with B_0 and B_1 corresponding to 0 and 1 predictions, respectively, we now allow for $B \geq B_0 + B_1$, in order to account for possibly some non-zero number of abstentions.

In more detail, let us denote $\hat{\mathbb{Y}}$ to be the multiset of predictions for models $h_{\boldsymbol{D}_1}, h_{\boldsymbol{D}_2}, \ldots, h_{\boldsymbol{D}_B}$ on $(\boldsymbol{x}, \boldsymbol{g})$, with $|\hat{\mathbb{Y}}| = B = B_0 + B_1 + B_{\texttt{Abstain}}$. This is where we depart from our typical definition of self-consistency, for which $B = B_0 + B_1$ (Section 3, Appendix B.3). We continue to let B_0 and B_1 represent the counts of 0 and 1 predictions, respectively, and now include $B_{\texttt{Abstain}}$ to denote the (possibly nonzero) number of abstentions. This leads to the following adjustment of (3):

$$\hat{\text{SC}}\big(\mathcal{A}, \{\hat{\boldsymbol{D}}_b\}_{b=1}^B, (\boldsymbol{x}, \boldsymbol{g})\big) = 1 - \frac{2(B_0 B_1 + B_0 B_{\texttt{Abstain}} + B_1 B_{\texttt{Abstain}})}{B(B-1)}. \tag{15}$$

Equation (15) follows from a similar analysis of comparing 0s, 1s, and abstentions for Definition 3.2, which lead us to derive (3) in Appendix B.3. However, since the costs of 0-to-`Abstain` comparisons and 1-to-`Abstain` comparisons are both 0, the $B_0 B_{\texttt{Abstain}}$ and $B_1 B_{\texttt{Abstain}}$ terms in (15) reduce to 0. As a result, we yield our original definition for self-consistency (3), with the possibility that $B = B_0 + B_1 + B_{\texttt{Abstain}} > B_0 + B_1$, if there is a nonzero number of abstentions $B_{\texttt{Abstain}}$.

Since $B > 1$ and $B_0, B_1, B_{\texttt{Abstain}} \geq 0$, it is always the case that option to `Abstain` is at least as self-consistent as not

having the option to do so. This follows from the fact that $B_0 + B_1 + B_{\texttt{Abstain}} = B \geq B_0 + B_1$, which would make the denominator in (15) greater than or equal to the corresponding method that cannot $\texttt{Abstain}$; when subtracted from 1, this would produce a $\hat{\text{SC}}$ that is no smaller than the value for the corresponding method without that cannot $\texttt{Abstain}$.

Now, it follows that, given the choice between $\texttt{Abstain}$ and predicting a label that is in disagreement with an existing prediction label, choosing to $\texttt{Abstain}$ will always lead to higher self-consistency. This is because the cost to $\texttt{Abstain}$ is less than disagreeing, so it will always be the minimal choice that maximizes $\hat{\text{SC}}$.

Error and the abstention set. It is very straightforward to see that the *abstention set* will generally exhibit higher than the *prediction set*. When we ensemble and measure $\hat{\text{SC}}$, the exmaples that exhibit low $\hat{\text{SC}}$ contain higher variance-induced error. Let us call the size of the abstention set U (which incurs error u), the size of the prediction set V (which incurs error v), and the size of the test set T (which incurs error t). We can relate the total number of misclassified examples as $T * t = U * u + V * v$, with $T = U + V$. If we assume the bias and noise are equally distributed across the test and abstention sets (this is a reasonable assumption, on average, in our setup), then splitting off the high variance instances from the low variance (high $\hat{\text{SC}}$ instances) requires that $u > v$. The error on the abstention set necessarily has to be larger than the error on the prediction set, in order to retain the above relationship.

E. Additional Empirical Results and Details for Reproducibility

The code for the examples in Sections 1, 3 and 5 can be found in `https://github.com/pasta41/variance`. This repository also contains necessary and sufficient information concerning reproducibility. At the time of writing, we use `Conda` to produce environments with associated package-versioning information, so that our results can be exactly replicated and independently verified. We also use the `Scikit-Learn` (Pedregosa et al., 2011) toolkit for modeling and optimization. More details on our choice of models and hyperparameter optimization can be found in our code repository, cited above. In brief, we consulted prior related work (e.g., Chen et al. (2018)) and performed our own validation for reasonable hyperparameters per model type. We keep these settings fixed to reduce impact on our results, in order to observe in isolation how different training data subsets impact our results.

During these early runs, we collected information on train accuracy, not just test accuracy; while models ultimately have similar test accuracy in most cases for the same task, they can vary significantly in terms of train accuracy (e.g., for logistic regression, `COMPAS` is in the low .70s; for random forests, it is in the mid .90s). We do not include these results for the sake of space.

This section is organized as follows. We first present information on our datasets, models and code, including our `HDMA` toolkit (Appendix E.1). We then provide details on our setup for running experiments on our cluster (Appendix E.2). Appendix E.3 contains more detailed information concerning the experiments performed to produce Figures 1 and 2 in the main paper. In Appendix E.4, we provide more details on the results presented in Section 5, as well as additional experiments. Lastly, in ppendix E.5, we discuss implications of these results for common fairness Abenchmarks like `South German Credit`. We conclude that in many cases, without adequate attention to error estimation, it is likely that training and post-processing a single model for fairness on these models likely is a brittle approach to achieve generalizable fairness (and accuracy) performance. Based on our empirical results, it seems like high variance can be a significant confounding factor when using a small set of models to draw conclusions about performance — whether fairness or accuracy. There is an urgent need for future work concerning reproducibility. More specifically, our results indicate that it would be useful to revisit key algorithmic strategies in fair classification to see how they perform in context with more reliable expected error estimation and variance reduction.

Note on CDF figures. We show our results in terms of the $\hat{\text{SC}}$ of the underlying bagged models because doing so conveys how Algorithm 1 makes decisions to predict or abstain.[13] For both types of ensembling, Algorithm 1 predicts for all examples captured by the area to the right of the κ reference line, and abstains for all examples on the left.

It is also worth noting (though hopefully obvious) that our CDF plots of $\hat{\text{SC}}$ are not continuous, yet we choose to plot them as interpolated curves. This are discrete because we train a concrete number of models (individual models or bags) — typically 101 of them — that we treat as our approximation for B when computing $\hat{\text{SC}}$. This means that there are a finite number of κ-values for $\hat{\text{SC}}$, for which we plot a corresponding concrete number of heights y corresponding to the cumulative proportion of the test set. In this respect, it would perhaps be more precise to plot our curves using a step function, exemplified below (see Appendix B.3.1 for the values in $\hat{\mathbb{K}}$):

---

[13]The $\hat{\text{SC}}$ CDF of Algorithm 1, computed via a *third* round of bootstrapping, has nearly all mass at $\hat{\text{SC}} = 1$; it is difficult to visualize.
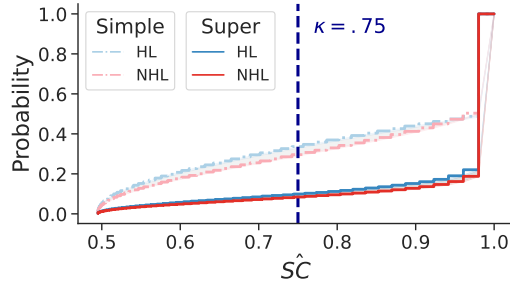
Figure 6. Plotting $\hat{SC}$ with an emphasis on discrete levels κ.

We opted not to do this for two reasons. First, plotting steps for some of our figures, in our opinion, can make the figures more difficult to understand. Second, in experiments for which we increase the number of models used to estimate $\hat{SC}$ (e.g., Appendix E.5), we found that the curves for 101 models were a reasonable approximation of the overall CDF. We therefore concluded that plotting the figures without steps was worth the clarity of presentation, with a sacrifice in correctness for the overall takeaways that we intend with these figures.

A remark on cost. It can be considerably more computationally intensive to train an ensemble of models to compute $\hat{SC}$ than to train a handful of models and perform cross-validation, as is the standard practice in fair classification. However, as our empirical analysis demonstrates, this cost comes with a huge benefit: It enables us to improve self-consistency and to root out the arbitrariness of producing predictions that are effectively close-to-random, which is especially important in high-stakes fairness settings (Cooper et al., 2021a). Moreover, for common fair classification datasets, the increased cost on modern hardware is relatively small; (super-) ensembling with confidence takes under an hour to execute (Appendix E.4).

E.1. Hypothesis classes, datasets, and code

Models. According to a comprehensive recent survey study (Fabris et al., 2022), as well as related work like Chen et al. (2018), we conclude that some of the most common models used in fair classification are logistic regression, decision tree classifiers, random forest classifiers, SVMs, and MLPs. We opted to include comprehensive results for the first three, since they capture different complexities, and therefore encode different degrees of statistical bias, that we expected to have an impact on the underlying sources of error. We provide some results for SVMs and MLPs, which we include in this Appendix. Since we choose not to use stochastic optimizers to reduce the sources of randomness, for our results, training MLPs is slower than it could be. We consistently use a decision threshold of 0.5 (i.e., 0-1 loss) for our experiments, though our results can easily be extended to other thresholds, as discussed in Section 3. Depending on the dataset, we reserve between 20% and 30% of the available data for the test set. This is consistent with standard fair classification training settings, which we validated during our initial experiments to explore the space (for which we also did preliminary hyperparameter optimization, before fixing the hyperparameters for our presented results). Please refer to `https://github.com/pasta41/variance` for more details.

Datasets. Also according to Fabris et al. (2022), the most common tasks in fair classification are `Old Adult` (Kohavi, 1996), `COMPAS` (Larson et al., 2016), and `South German Credit` (Grömping, 2019).[14] These three datasets arguably serve as a *de facto* benchmark in the community, so we felt the need to include them in the present work. In recognition of the fact that these three datasets, however standard, have problems, we also run experiments on 3 tasks in the `New Adult` dataset, introduced by Ding et al. (2021) to replace `Old Adult`. We subset to the `CA` (California) subset of the dataset, and run on `Income`, `Employment`, and `Public Coverage`, and consider `sex` and `race` as protected attributes, which we binarize into {Male, Female} and {White, Non-white}. These are all large-scale tasks, at least in the domain of algorithmic fairness — on the order of hundreds of thousands of example instances. However, the 3 tasks do share example instances and some features. In summary, concerning common tasks in fair classification:

- `COMPAS` (Larson et al., 2016). We run on the commonly-used version of this dataset from Friedler et al. (2019), which has 6167 example instances with 404 features. The target is to predict recidivism within 2 years (1 corresponding to Yes, and 0 to No). The protected attribute is `race`, binarized into "Non-white" (0) and "White" (1) subgroups.

- `Old Adult` (Kohavi, 1996). We run on the commonly-used version of this dataset from Friedler et al. (2019), which

---

[14]Technically, Grömping (2019) is an updated and corrected version of the dataset from 2019.

has 30,162 examples with 97 features. This version of the dataset removes instances with missing values from the original dataset, and changes the encoding of some of the features (Kohavi (1996) has 48842 example instnaces with 88 features). The target is to predict $< \$50,000$ income (0) $>= \$50,000$ income (1). The protected attribute is `sex`, binarized into "Female" (0) and "Male" (1) subgroups.

- `South German Credit` (Grömping, 2019). We download the dataset from UCI and process the data ourselves. We use the provided `codetable.txt` to "translate" the features from German to English. We say "translate" because the authors took some liberties, e.g., the column converted to "credit_history" is labeled "moral" in the German, which is not a translation. There are four categories in the protected attribute "personal_status_sex" column, one of which (2) is used for both "Male (single)" and "Female (non-single)." We therefore remove rows with this value, and binarize the remaining three categories into "Female" (0) and "Male" (1). What results is a dataset with 690 example instances (of the original 1000) with 19 features. The target is "good" credit (1) and "bad" credit (0).

- `Taiwan Credit` (Yeh and Lien, 2009). This task is to predict default on credit card payments (1) or not (0). There are 30,000 example instances and 24 features. The protected attribute is binary `sex`. We download this dataset from UCI.

- `New Adult` (Ding et al., 2021). This dataset contains millions of example instances from US Census data, which can be used for several different targets/tasks. We select three of them (listed below). These tasks share some features, and therefore are not completely independent. Further, given the size of the whole dataset, we subset to `CA` (California), the most populous state in the US. There are two protected attribute columns that we use: `sex`, which is binarized "Female" (0) and "Male" (1) subgroups, and `race`, which we binarize into "Non-white" (0) and "White" (1). In future work, we would like to explore extending our results beyond binary subgroups.

  - `Income`. This task is designed to be analogous to `Old Adult` (Kohavi, 1996). As a result, the target is to predict $< \$50,000$ income (0) $>= \$50,000$ income (1). In the `CA` subset, there are 195,665 example instances with 8 features.
  - `Employment`. This task is to predict whether an individual is employed (1) or not (0). In the `CA` subset, there are 378,817 example instances with 14 features.
  - `Public Coverage`. This task is to predict whether an individual is on public health insurance (1) or not (0). In the `CA` subset, there are 138,554 example instances with 17 features.

E.1.1. THE STANDALONE `HMDA` TOOKIT

In addition to the above standard tasks, we include experiments that use the `NY` and `TX` 2017 subsets of the the Home Mortgage Data Disclosure Act (`HMDA`) 2007-2017 dataset (Federal Financial Institutions Examination Council, 2017). These two datasets have 244,107 and 576,978 examples, respectively, with 18 features. The `HMDA` datasets together contain over 140 million examples of US home mortgage loans from 2007-2017 (newer data exists, but in a different format). We developed a toolkit, described below, to make this dataset easy to use for classification experiments. Similar to `New Adult`, we enable subsetting by US state. For the experiments in this paper, we run on the `NY` (New York) and `TX` (Texas) 2017 subset, in order to add some geographic diversity to complement our `New Adult` experiments. We additionally chose `NY` and `TX` because they are two of the most populous states in the US, alongside `CA`.[15]

The target variable, `action_taken`, concerning loan origination has 8 values, 2 of which we cannot meaningful conclude approval or denial decisions. They are: Action Taken: 1 – Loan originated, 2 – Application approved but not accepted, 3 – Application denied by financial institution, 4 – Application withdrawn by applicant, 5 – File closed for incompleteness, 6 – Loan purchased by the institution, 7 – Preapproval request denied by financial institution, and 8 – Preapproval request approved but not accepted (optional reporting). We filter out 4 and 6, and binarize into `grant`=$\{1, 2, 8\}$ = 1 and `reject`=$\{3, 5, 7\}$ = 0. There are three protected attributes that we consider: `sex`, `race`, and `ethnicity`:

- `sex` has 5 possible values, 2 of which correspond to categories/non-missing values: Male – 1 and Female – 2. We binarize `sex` into F = 0 and M = 1.

- `race` has 8 possible values, 5 of which correspond to categories/ non-missing information: 1 – American Indian or Alaska Native, 2 – Asian, 3 – Black or African American, 4 – Native Hawaiian or Other Pacific Islander, and 5 –

---

[15]Per the 2020 Census, the top-4-most-populous states are `CA`, `TX`, `FL`, and `NY` (Mackun et al., 2021).

White. There are 5 fields for applicant race, which model an applicant belonging to more than one racial group. For our experiments, we only look at the first field. When we binarize `race`, NW = 0 and W = 1.

- `ethnicity` has 5 possible values, 2 of which correspond to categories/ non-missing information: 1 – Hispanic or Latino and 2 – Not Hispanic or Latino. We binarize ethnicity to be HL = 0 and NHL = 1.

After subsetting to only include examples that have values that do not correspond to missing information, `HMDA` has 18 features. The `NY` dataset has 244,107 examples; the `TX` dataset has 576,978 examples, making it the largest dataset in our experiments. As with our experiments using `New Adult`, we would like to extend our results beyond binary subgroups and binary classification in future work.

Releasing a standalone toolkit. These datasets are less-commonly used in current algorithmic fairness literature (Fabris et al., 2022). We believe this is likely due to the fact that the over-100-million data examples are only available in bulk files, which are on the order of 10s of gigabytes and therefore not easily downloadable or explorable on most personal computers. Following the example of Ding et al. (2021), one of our contributions is to pre-process all of these datasets — all locations and years — and release them with a software toolkit. The software engineering effort to produce this toolkit was substantial. Our hope is that wider access to this dataset will further reduce the community's dependency on small (and dated) datasets. Please refer to https://github.com/pasta41/hmda for the latest information on this standalone software package. Our release aligns with the terms of service for this dataset.

E.2. Cluster environment details

While most of the experiments run in this paper can be easily reproduced on a modern laptop, for efficiency, we ran all of our experiments (except the one to produce Figure 1) in a cluster environment. This enabled us to easily execute train/test splits n in parallel on different CPUs, serialize our results, and then reconstitute and combine them to produce plots locally. Our cluster environment runs Ubuntu 20.04 and uses Slurm v20.11.8 to manage jobs. We ran all experiments using `Anaconda3`, which is why we used `Conda` to reproduce environments for easy replicability.

The experiments using `New Adult` and `HMDA` rely on datasets that are (in some cases) orders of magnitude larger than the traditional algorithmic fairness tasks. This is one of the reasons why we recommend running on a cluster, and therefore do not include Jupyter notebooks in our repository for these tasks. We also limit our modeling choices to logistic regression, decision tree classifiers, and random forest classifiers for these results due to the expense of training on the order of thousands of models for each experiment.

E.3. Details on motivating examples in the main paper

This appendix provides extended results for the experiments associated in Sections 1 and 3, which give an intuition for individual- and subgroup-level consistency. The experimental results in the main paper are for logistic regression. We expand the set of models we examine, and associated discussion of how to interpret comparisons between these results.

Reproducing Figure 1. The experiment to produce this figure in Section 1 (also shown in Appendix B.3) trains $B = 10$ logistic regression models on the `COMPAS` dataset (Appendix E.1) using 0-1 loss. We use the bootstrap method to produce each model, which we evaluate on the same test set. We then search for a maximally consistent and minimally consistent individual in the test set, i.e., an individual with 10 predictions that agree and an individual with 5 predictions in each class, which we plot in the bar graph. Please refer to the README in https://github.com/pasta41/variance regarding which `Jupyter` notebook to run to produce the underlying results and figure. The experiments to reproduce this figure can be easily replicated on a laptop.

Reproducing Figure 2. These figures were produced by executing $S = 10$ runs of $B = 101$ bootstrap training replicates to train random forest classifiers for `Old Adult` and `COMPAS`. We reproduce these figures below, so that they can be examined and treated in relation to our additional results for decision tree classifiers and logistic regression. For each s run, we take train/test split, bootstrap the train split $B = 101$ times, and evaluate the resulting model classification decisions on the test set. $\hat{SC}$ can be estimated from the results across those 101 models. We Run this process $S = 10$ times to produce confidence intervals, shown in the figures below. The intervals are not always clearly visible; there is not a lot of variance at the level of comparing whole runs to each other. Please refer to the README in https://github.com/pasta41/variance regarding which `Jupyter` notebook to run to produce the underlying results and figure. There are also scripted version of these experiments, which enable them to be run in parallel in a cluster environment.

Self-consistency of incorrectly-classified instances. Last, we include figures that underscore how self-consistency is independent from correctness that is measured in terms of observed label alignment. That is, it is possible for an instance $(\boldsymbol{x}, \boldsymbol{g})$ to be self-consistent and classified incorrectly, with respect to its observed label o.

We show this using stacked bar plots. For the above experiments, we find the test examples that have the majority of their classifications incorrect ($\hat{y} \neq o$, for $B = 101$, we find the instances with ≥ 51 incorrect classifications) and the majority of their classification correct (similarly), and we examine how self-consistent they are. We bucket self-consistency into different levels, and then plot the relative proportion of majority-incorrectly and majority-correctly classified examples according to subgroup. Subgroups in COMPAS exhibit a similar trend, while subgroups in Adult Old exhibit differences, with the heights of the bars corresponding to the trends we plot in our CDF plots. As we note briefly in Section 3, it may be interesting to examine patterns in examples about which learning processes are confident (i.e., highly self-consistent) but wrong in terms of label alignment. If such issues correlate with subgroup, it may be worth testing the counterfactual that such labels are indicative of label bias. We leave such thoughts to future work.

(a) COMPAS

(b) Adult Old
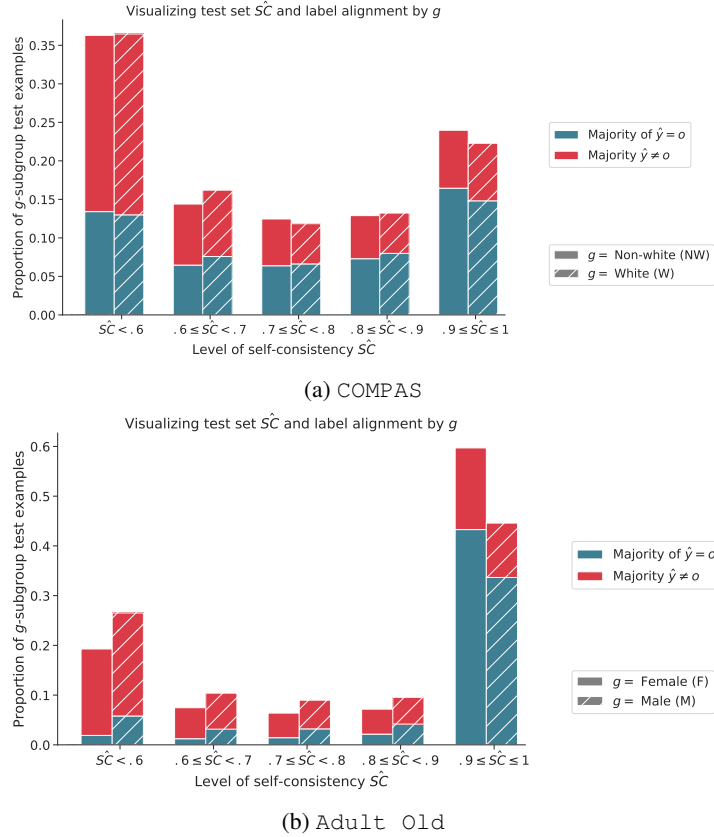
Figure 7. $\hat{SC}$ broken down by $\boldsymbol{g}$ and label alignment with the observed label o. For each train/test split, and for each $\hat{SC}$ range (x-axis), we find the examples that are incorrectly classified the majority of time (≥ 5 splits, we find that $\hat{y} \neq o$), and the examples that are correctly classified the majority of the time (> 5, we find that $\hat{y} = o$). We compute the average the proportion over (over splits) in each $\hat{SC}$ range (y-axis). We plot these proportions with respect to subgroup $\boldsymbol{g}$ (where the sums of the heights of bars for by each $\boldsymbol{g}$ is equal to 1).

E.4. Validating our algorithm in practice

E.4.1. COMPAS

$\hat{\text{SC}}$ CDFs for COMPAS ($g = \text{race}$) and associated error metrics on the prediction set. **Baseline** metrics computed with $B = 101$ models. For **simple**, $B = 101$ models; for **super**, $B = 101$ ensemble models, each composed of 51 underlying models. We repeat for 10 test/train splits. We also report abstention rate $\hat{\text{AR}}$.
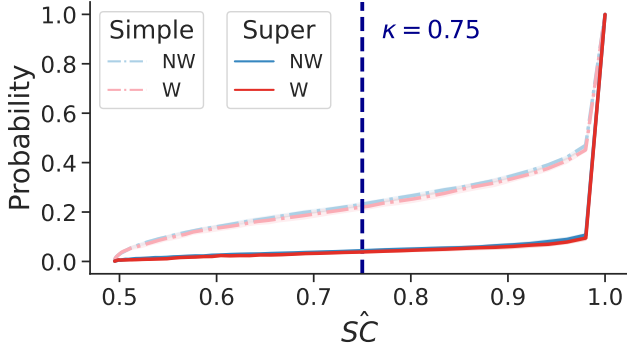
Abstention set metrics		
	Simple	**Super**
$\Delta\hat{\textbf{AR}}$	$1.1 \pm 0.9\%$	$0.5 \pm 0.0\%$
$\hat{\textbf{AR}}_{\textbf{NW}}$	$23.2 \pm 1.3\%$	$4.3 \pm 0.5\%$
$\hat{\textbf{AR}}_{\textbf{W}}$	$22.1 \pm 2.2\%$	$3.8 \pm 0.5\%$

Logistic regression prediction set metrics			
	Baseline	**Simple**	**Super**
$\Delta\hat{\textbf{PR}}$	$14.5 \pm 0.3\%$	$18.7 \pm 0.5\%$	$15.6 \pm 0.1\%$
$\hat{\textbf{PR}}_{\textbf{NW}}$	$45.3 \pm 1.2\%$	$43.8 \pm 1.1\%$	$44.2 \pm 0.7\%$
$\hat{\textbf{PR}}_{\textbf{W}}$	$30.8 \pm 1.5\%$	$25.1 \pm 1.6\%$	$28.6 \pm 0.6\%$
$\Delta\hat{\textbf{Err}}$	$0.2 \pm 0.2\%$	$1.1 \pm 1.5\%$	$0.9 \pm 1.1\%$
$\hat{\textbf{Err}}_{\textbf{NW}}$	$33.0 \pm 1.3\%$	$27.9 \pm 0.9\%$	$31.0 \pm 1.0\%$
$\hat{\textbf{Err}}_{\textbf{W}}$	$33.2 \pm 1.1\%$	$29.0 \pm 2.4\%$	$31.9 \pm 2.1\%$
$\Delta\hat{\textbf{FPR}}$	$2.1 \pm 0.0\%$	$3.0 \pm 0.0\%$	$1.8 \pm 0.2\%$
$\hat{\textbf{FPR}}_{\textbf{NW}}$	$14.7 \pm 1.3\%$	$11.4 \pm 1.0\%$	$12.9 \pm 0.8\%$
$\hat{\textbf{FPR}}_{\textbf{W}}$	$12.6 \pm 1.3\%$	$8.4 \pm 1.0\%$	$11.1 \pm 0.6\%$
$\Delta\hat{\textbf{FNR}}$	$2.4 \pm 0.0\%$	$4.0 \pm 1.1\%$	$2.8 \pm 0.8\%$
$\hat{\textbf{FNR}}_{\textbf{NW}}$	$18.3 \pm 1.1\%$	$16.5 \pm 1.9\%$	$18.0 \pm 1.3\%$
$\hat{\textbf{FNR}}_{\textbf{W}}$	$20.7 \pm 1.1\%$	$20.5 \pm 3.0\%$	$20.8 \pm 2.1\%$

Figure 8. **Logistic regression** on COMPAS
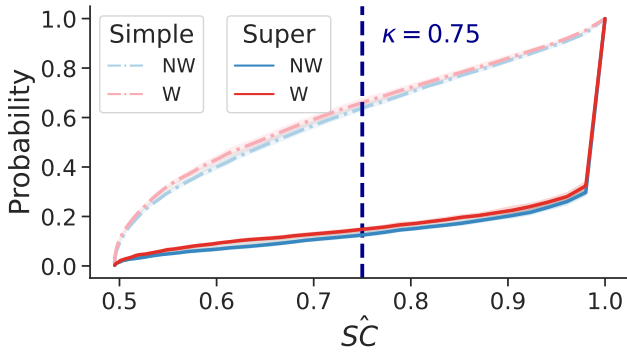
Abstention set metrics		
	Simple	**Super**
$\Delta\hat{\textbf{AR}}$	$1.9 \pm 1.0\%$	$2.3 \pm 0.1\%$
$\hat{\textbf{AR}}_{\textbf{NW}}$	$62.3 \pm 1.8\%$	$12.3 \pm 0.8\%$
$\hat{\textbf{AR}}_{\textbf{W}}$	$64.2 \pm 2.8\%$	$14.6 \pm 0.9\%$

Decision tree prediction set metrics			
	Baseline	**Simple**	**Super**
$\Delta\hat{\textbf{PR}}$	$10.1 \pm 0.6\%$	$22.9 \pm 1.7\%$	$15.8 \pm 0.5\%$
$\hat{\textbf{PR}}_{\textbf{NW}}$	$47.9 \pm 0.7\%$	$43.4 \pm 3.1\%$	$48.5 \pm 1.2\%$
$\hat{\textbf{PR}}_{\textbf{W}}$	$37.8 \pm 1.3\%$	$20.5 \pm 1.4\%$	$32.7 \pm 1.7\%$
$\Delta\hat{\textbf{Err}}$	$0.6 \pm 0.9\%$	$1.7 \pm 0.7\%$	$1.2 \pm 0.8\%$
$\hat{\textbf{Err}}_{\textbf{NW}}$	$38.8 \pm 0.5\%$	$24.0 \pm 0.9\%$	$32.8 \pm 0.4\%$
$\hat{\textbf{Err}}_{\textbf{W}}$	$38.2 \pm 1.4\%$	$22.3 \pm 1.6\%$	$31.6 \pm 1.2\%$
$\Delta\hat{\textbf{FPR}}$	$0.2 \pm 0.4\%$	$4.0 \pm 0.4\%$	$2.5 \pm 0.9\%$
$\hat{\textbf{FPR}}_{\textbf{NW}}$	$18.8 \pm 0.8\%$	$10.4 \pm 1.8\%$	$16.1 \pm 0.9\%$
$\hat{\textbf{FPR}}_{\textbf{W}}$	$18.6 \pm 1.2\%$	$6.4 \pm 1.4\%$	$13.6 \pm 1.8\%$
$\Delta\hat{\textbf{FNR}}$	$0.3 \pm 0.3\%$	$2.3 \pm 1.3\%$	$1.4 \pm 0.1\%$
$\hat{\textbf{FNR}}_{\textbf{NW}}$	$19.9 \pm 0.7\%$	$13.6 \pm 1.0\%$	$16.6 \pm 1.3\%$
$\hat{\textbf{FNR}}_{\textbf{W}}$	$19.6 \pm 1.0\%$	$15.9 \pm 2.3\%$	$18.0 \pm 1.2\%$

Figure 9. **Decision trees** on COMPAS
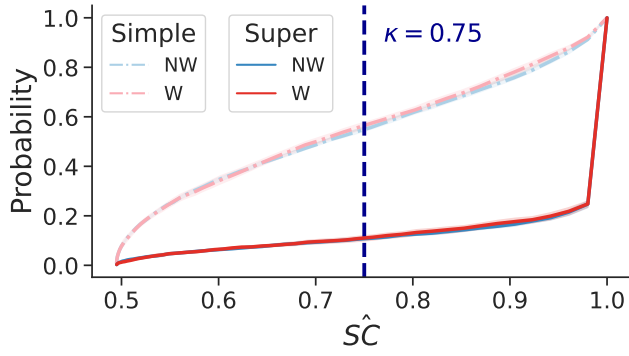
Abstention set metrics		
	Simple	**Super**
$\Delta\hat{\mathbf{AR}}$	$0.3 \pm 0.6\%$	$0.2 \pm 0.7\%$
$\hat{\mathbf{AR}}_{\mathbf{NW}}$	$53.9 \pm 1.6\%$	$10.6 \pm 0.5\%$
$\hat{\mathbf{AR}}_{\mathbf{W}}$	$53.6 \pm 2.2\%$	$10.8 \pm 1.2\%$

Random forest prediction set metrics			
	Baseline	**Simple**	**Super**
$\Delta\hat{\mathbf{PR}}$	$13.0 \pm 0.7\%$	$24.3 \pm 0.4\%$	$18.6 \pm 0.5\%$
$\hat{\mathbf{PR}}_{\mathbf{NW}}$	$48.0 \pm 0.6\%$	$45.6 \pm 1.7\%$	$47.8 \pm 0.9\%$
$\hat{\mathbf{PR}}_{\mathbf{W}}$	$35.0 \pm 1.3\%$	$21.3 \pm 1.3\%$	$29.2 \pm 1.4\%$
$\Delta\hat{\mathbf{Err}}$	$1.0 \pm 0.8\%$	$0.6 \pm 0.8\%$	$2.1 \pm 1.0\%$
$\hat{\mathbf{Err}}_{\mathbf{NW}}$	$36.9 \pm 0.5\%$	$23.3 \pm 0.8\%$	$32.3 \pm 0.4\%$
$\hat{\mathbf{Err}}_{\mathbf{W}}$	$35.9 \pm 1.3\%$	$23.9 \pm 1.6\%$	$30.2 \pm 1.4\%$
$\Delta\hat{\mathbf{FPR}}$	$2.0 \pm 0.4\%$	$3.2 \pm 0.0\%$	$4.5 \pm 0.4\%$
$\hat{\mathbf{FPR}}_{\mathbf{NW}}$	$18.0 \pm 0.8\%$	$10.0 \pm 1.3\%$	$15.3 \pm 1.2\%$
$\hat{\mathbf{FPR}}_{\mathbf{W}}$	$16.0 \pm 1.2\%$	$6.8 \pm 1.3\%$	$10.8 \pm 0.8\%$
$\Delta\hat{\mathbf{FNR}}$	$0.9 \pm 0.4\%$	$3.7 \pm 1.2\%$	$2.4 \pm 0.8\%$
$\hat{\mathbf{FNR}}_{\mathbf{NW}}$	$19.0 \pm 0.7\%$	$13.4 \pm 1.2\%$	$16.9 \pm 1.2\%$
$\hat{\mathbf{FNR}}_{\mathbf{W}}$	$19.9 \pm 1.1\%$	$17.1 \pm 2.4\%$	$19.3 \pm 2.0\%$

Figure 10. **Random forests** on `COMPAS`

E.4.2. OLD ADULT

$\hat{SC}$ CDFs for `Old Adult` ($g = $ `sex`) and associated error metrics on the prediction set. **Baseline** metrics computed with $B = 101$ models. For **simple**, $B = 101$ models; for **super**, $B = 101$ ensemble models, each composed of 51 underlying models. We repeat for 10 test/train splits. We also report abstention rate $\hat{AR}$.
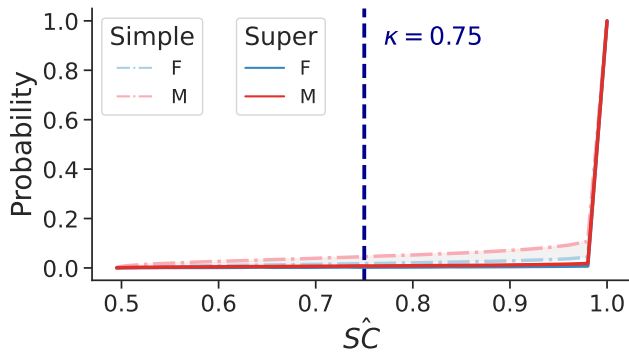
Abstention set metrics		
	Simple	**Super**
$\Delta\hat{\mathbf{AR}}$	$2.6 \pm 0.0\%$	$0.5 \pm 0.0\%$
$\hat{\mathbf{AR}}_{\mathbf{F}}$	$1.8 \pm 0.2\%$	$0.3 \pm 0.1\%$
$\hat{\mathbf{AR}}_{\mathbf{M}}$	$4.4 \pm 0.2\%$	$0.8 \pm 0.1\%$

Logistic regression prediction set metrics			
	Baseline	**Simple**	**Super**
$\Delta\hat{\mathbf{PR}}$	$18.3 \pm 0.2\%$	$17.8 \pm 0.1\%$	$18.1 \pm 0.1\%$
$\hat{\mathbf{PR}}_{\mathbf{F}}$	$8.2 \pm 0.3\%$	$7.1 \pm 0.4\%$	$7.6 \pm 0.4\%$
$\hat{\mathbf{PR}}_{\mathbf{M}}$	$26.5 \pm 0.5\%$	$24.9 \pm 0.5\%$	$25.7 \pm 0.5\%$
$\Delta\hat{\mathbf{Err}}$	$11.3 \pm 0.1\%$	$10.8 \pm 0.1\%$	$11.4 \pm 0.2\%$
$\hat{\mathbf{Err}}_{\mathbf{F}}$	$7.8 \pm 0.4\%$	$7.0 \pm 0.3\%$	$7.5 \pm 0.2\%$
$\hat{\mathbf{Err}}_{\mathbf{M}}$	$19.1 \pm 0.3\%$	$17.8 \pm 0.4\%$	$18.9 \pm 0.4\%$
$\Delta\hat{\mathbf{FPR}}$	$4.7 \pm 0.0\%$	$4.4 \pm 0.2\%$	$4.8 \pm 0.2\%$
$\hat{\mathbf{FPR}}_{\mathbf{F}}$	$2.3 \pm 0.3\%$	$1.6 \pm 0.1\%$	$1.8 \pm 0.1\%$
$\hat{\mathbf{FPR}}_{\mathbf{M}}$	$7.0 \pm 0.3\%$	$6.0 \pm 0.3\%$	$6.6 \pm 0.3\%$
$\Delta\hat{\mathbf{FNR}}$	$6.7 \pm 0.1\%$	$6.5 \pm 0.1\%$	$6.6 \pm 0.1\%$
$\hat{\mathbf{FNR}}_{\mathbf{F}}$	$5.5 \pm 0.3\%$	$5.4 \pm 0.2\%$	$5.7 \pm 0.2\%$
$\hat{\mathbf{FNR}}_{\mathbf{M}}$	$12.2 \pm 0.2\%$	$11.9 \pm 0.1\%$	$12.3 \pm 0.1\%$

Figure 11. **Logistic regression** on `Old Adult`

Decision tree prediction set metrics			
	Baseline	**Simple**	**Super**
$\Delta\hat{PR}$	$20.3 \pm 0.1\%$	$18.2 \pm 0.1\%$	$19.9 \pm 0.1\%$
$\hat{PR}_F$	$12.1 \pm 0.4\%$	$4.5 \pm 0.3\%$	$7.8 \pm 0.5\%$
$\hat{PR}_M$	$32.4 \pm 0.5\%$	$22.7 \pm 0.2\%$	$27.7 \pm 0.4\%$
$\Delta\hat{Err}$	$12.3 \pm 0.0\%$	$6.0 \pm 0.1\%$	$10.9 \pm 0.2\%$
$\hat{Err}_F$	$10.8 \pm 0.3\%$	$3.0 \pm 0.2\%$	$6.6 \pm 0.4\%$
$\hat{Err}_M$	$23.1 \pm 0.3\%$	$9.0 \pm 0.3\%$	$17.5 \pm 0.2\%$
$\Delta\hat{FPR}$	$6.2 \pm 0.1\%$	$2.5 \pm 0.1\%$	$5.4 \pm 0.2\%$
$\hat{FPR}_F$	$5.7 \pm 0.2\%$	$0.4 \pm 0.0\%$	$1.9 \pm 0.3\%$
$\hat{FPR}_M$	$11.9 \pm 0.3\%$	$2.9 \pm 0.1\%$	$7.3 \pm 0.1\%$
$\Delta\hat{FNR}$	$6.1 \pm 0.1\%$	$3.4 \pm 0.0\%$	$5.5 \pm 0.1\%$
$\hat{FNR}_F$	$5.1 \pm 0.3\%$	$2.7 \pm 0.2\%$	$4.7 \pm 0.1\%$
$\hat{FNR}_M$	$11.2 \pm 0.2\%$	$6.1 \pm 0.2\%$	$10.2 \pm 0.2\%$

Abstention set metrics		
	Simple	**Super**
$\Delta\hat{AR}$	$19.5 \pm 0.2\%$	$3.5 \pm 0.1\%$
$\hat{AR}_F$	$18.2 \pm 0.4\%$	$3.4 \pm 0.2\%$
$\hat{AR}_M$	$37.7 \pm 0.6\%$	$6.9 \pm 0.3\%$

Figure 12. **Decision trees** on `Old Adult`

Random forest prediction set metrics			
	Baseline	**Simple**	**Super**
$\Delta\hat{PR}$	$20.0 \pm 0.2\%$	$17.1 \pm 0.3\%$	$19.0 \pm 0.2\%$
$\hat{PR}_F$	$9.8 \pm 0.2\%$	$4.8 \pm 0.2\%$	$7.7 \pm 0.3\%$
$\hat{PR}_M$	$29.8 \pm 0.4\%$	$21.9 \pm 0.5\%$	$26.7 \pm 0.5\%$
$\Delta\hat{Err}$	$12.2 \pm 0.0\%$	$6.5 \pm 0.1\%$	$10.7 \pm 0.0\%$
$\hat{Err}_F$	$9.0 \pm 0.3\%$	$4.2 \pm 0.2\%$	$6.6 \pm 0.2\%$
$\hat{Err}_M$	$21.2 \pm 0.3\%$	$10.7 \pm 0.3\%$	$17.3 \pm 0.2\%$
$\Delta\hat{FPR}$	$6.0 \pm 0.2\%$	$2.5 \pm 0.1\%$	$5.0 \pm 0.1\%$
$\hat{FPR}_F$	$3.7 \pm 0.1\%$	$0.7 \pm 0.1\%$	$1.7 \pm 0.2\%$
$\hat{FPR}_M$	$9.7 \pm 0.3\%$	$3.2 \pm 0.2\%$	$6.7 \pm 0.3\%$
$\Delta\hat{FNR}$	$6.3 \pm 0.2\%$	$4.1 \pm 0.2\%$	$5.8 \pm 0.1\%$
$\hat{FNR}_F$	$5.3 \pm 0.3\%$	$3.5 \pm 0.1\%$	$4.9 \pm 0.2\%$
$\hat{FNR}_M$	$11.6 \pm 0.1\%$	$7.6 \pm 0.3\%$	$10.7 \pm 0.3\%$
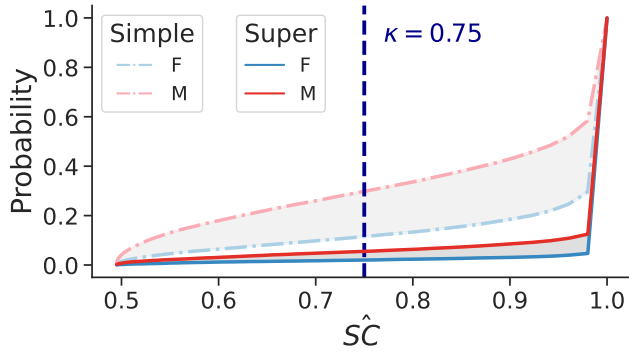
Abstention set metrics		
	Simple	**Super**
$\Delta\hat{AR}$	$17.2 \pm 0.4\%$	$3.4 \pm 0.1\%$
$\hat{AR}_F$	$11.2 \pm 0.3\%$	$2.0 \pm 0.3\%$
$\hat{AR}_M$	$28.4 \pm 0.7\%$	$5.4 \pm 0.2\%$

Figure 13. **Random forests** on `Old Adult`

E.4.3. SOUTH GERMAN CREDIT

$\hat{\text{SC}}$ CDFs for `German Credit` (g = `sex`) and associated error metrics on the prediction set. **Baseline** metrics computed with $B = 101$ models. For **simple**, $B = 101$ models; for **super**, $B = 101$ ensemble models, each composed of 51 underlying models. We repeat for 10 test/train splits. We also report abstention rate $\hat{\text{AR}}$.
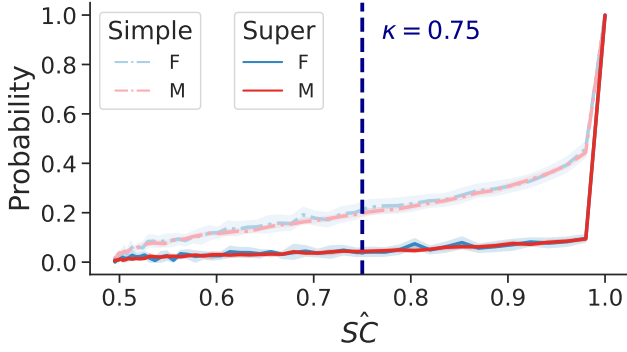
Abstention set metrics		
	Simple	**Super**
$\Delta\hat{\textbf{AR}}$	$0.4 \pm 3.9\%$	$0.1 \pm 1.8\%$
$\hat{\textbf{AR}}_{\textbf{F}}$	$20.8 \pm 6.6\%$	$4.1 \pm 3.1\%$
$\hat{\textbf{AR}}_{\textbf{M}}$	$21.2 \pm 2.7\%$	$4.0 \pm 1.3\%$

Logistic regression prediction set metrics			
	Baseline	**Simple**	**Super**
$\Delta\hat{\textbf{PR}}$	$8.8 \pm 1.4\%$	$9.1 \pm 1.2\%$	$9.7 \pm 1.7\%$
$\hat{\textbf{PR}}_{\textbf{F}}$	$88.8 \pm 4.7\%$	$96.0 \pm 4.1\%$	$91.7 \pm 5.0\%$
$\hat{\textbf{PR}}_{\textbf{M}}$	$80.0 \pm 3.3\%$	$86.9 \pm 2.9\%$	$82.0 \pm 3.3\%$
$\Delta\hat{\textbf{Err}}$	$0.9 \pm 4.0\%$	$4.6 \pm 6.1\%$	$3.4 \pm 4.7\%$
$\hat{\textbf{Err}}_{\textbf{F}}$	$23.3 \pm 6.9\%$	$22.8 \pm 8.7\%$	$25.5 \pm 7.4\%$
$\hat{\textbf{Err}}_{\textbf{M}}$	$24.2 \pm 2.9\%$	$18.2 \pm 2.6\%$	$22.1 \pm 2.7\%$
$\Delta\hat{\textbf{FPR}}$	$0.7 \pm 3.9\%$	$5.4 \pm 5.8\%$	$5.5 \pm 5.4\%$
$\hat{\textbf{FPR}}_{\textbf{F}}$	$16.2 \pm 6.2\%$	$19.6 \pm 8.4\%$	$21.1 \pm 7.8\%$
$\hat{\textbf{FPR}}_{\textbf{M}}$	$15.5 \pm 2.3\%$	$14.2 \pm 2.6\%$	$15.6 \pm 2.4\%$
$\Delta\hat{\textbf{FNR}}$	$1.6 \pm 1.1\%$	$0.8 \pm 1.9\%$	$2.1 \pm 1.8\%$
$\hat{\textbf{FNR}}_{\textbf{F}}$	$7.1 \pm 3.7\%$	$3.2 \pm 3.5\%$	$4.4 \pm 3.8\%$
$\hat{\textbf{FNR}}_{\textbf{M}}$	$8.7 \pm 2.6\%$	$4.0 \pm 1.6\%$	$6.5 \pm 2.0\%$

Figure 14. **Logistic regression** on `German Credit`
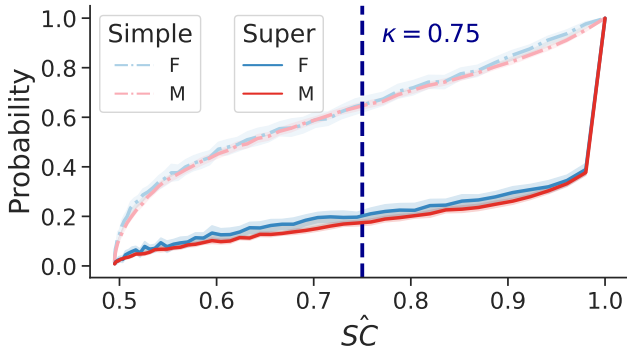
Abstention set metrics		
	Simple	**Super**
$\Delta\hat{\textbf{AR}}$	$0.0 \pm 2.6\%$	$2.8 \pm 2.6\%$
$\hat{\textbf{AR}}_{\textbf{F}}$	$65.2 \pm 6.0\%$	$19.9 \pm 5.9\%$
$\hat{\textbf{AR}}_{\textbf{M}}$	$65.2 \pm 3.4\%$	$17.1 \pm 3.3\%$

Decision tree prediction set metrics			
	Baseline	**Simple**	**Super**
$\Delta\hat{\textbf{PR}}$	$0.3 \pm 2.5\%$	$1.9 \pm 0.4\%$	$3.7 \pm 3.7\%$
$\hat{\textbf{PR}}_{\textbf{F}}$	$71.2 \pm 4.6\%$	$99.6 \pm 0.8\%$	$87.9 \pm 6.0\%$
$\hat{\textbf{PR}}_{\textbf{M}}$	$70.9 \pm 2.1\%$	$97.7 \pm 1.2\%$	$84.2 \pm 2.3\%$
$\Delta\hat{\textbf{Err}}$	$1.1 \pm 2.9\%$	$0.3 \pm 5.2\%$	$0.1 \pm 4.9\%$
$\hat{\textbf{Err}}_{\textbf{F}}$	$33.0 \pm 4.8\%$	$9.8 \pm 8.0\%$	$20.3 \pm 8.2\%$
$\hat{\textbf{Err}}_{\textbf{M}}$	$31.9 \pm 1.9\%$	$9.5 \pm 2.8\%$	$20.2 \pm 3.3\%$
$\Delta\hat{\textbf{FPR}}$	$0.7 \pm 3.5\%$	$0.8 \pm 5.1\%$	$0.4 \pm 3.9\%$
$\hat{\textbf{FPR}}_{\textbf{F}}$	$15.6 \pm 5.9\%$	$9.6 \pm 7.9\%$	$15.2 \pm 7.0\%$
$\hat{\textbf{FPR}}_{\textbf{M}}$	$14.9 \pm 2.4\%$	$8.8 \pm 2.8\%$	$14.8 \pm 3.1\%$
$\Delta\hat{\textbf{FNR}}$	$0.5 \pm 2.3\%$	$0.4 \pm 0.0\%$	$0.2 \pm 3.6\%$
$\hat{\textbf{FNR}}_{\textbf{F}}$	$17.4 \pm 4.4\%$	$0.2 \pm 0.7\%$	$5.2 \pm 5.1\%$
$\hat{\textbf{FNR}}_{\textbf{M}}$	$16.9 \pm 2.1\%$	$0.6 \pm 0.7\%$	$5.4 \pm 1.5\%$

Figure 15. **Decision trees** on `German Credit`

Random forest prediction set metrics			
	Baseline	**Simple**	**Super**
$\Delta\hat{\mathbf{PR}}$	$3.9 \pm 1.6\%$	$1.9 \pm 0.6\%$	$4.9 \pm 2.1\%$
$\hat{\mathbf{PR}}_{\mathbf{F}}$	$81.7 \pm 3.3\%$	$99.9 \pm 0.4\%$	$94.5 \pm 4.0\%$
$\hat{\mathbf{PR}}_{\mathbf{M}}$	$77.8 \pm 1.7\%$	$98.0 \pm 1.0\%$	$89.6 \pm 1.9\%$
$\Delta\hat{\mathbf{Err}}$	$2.3 \pm 3.1\%$	$0.0 \pm 4.7\%$	$2.4 \pm 6.7\%$
$\hat{\mathbf{Err}}_{\mathbf{F}}$	$25.8 \pm 5.1\%$	$11.9 \pm 7.8\%$	$23.5 \pm 9.6\%$
$\hat{\mathbf{Err}}_{\mathbf{M}}$	$28.1 \pm 2.0\%$	$11.9 \pm 3.1\%$	$21.1 \pm 2.9\%$
$\Delta\hat{\mathbf{FPR}}$	$2.5 \pm 3.0\%$	$0.3 \pm 4.7\%$	$2.2 \pm 6.2\%$
$\hat{\mathbf{FPR}}_{\mathbf{F}}$	$13.8 \pm 4.8\%$	$11.8 \pm 7.8\%$	$20.5 \pm 9.1\%$
$\hat{\mathbf{FPR}}_{\mathbf{M}}$	$16.3 \pm 1.8\%$	$11.5 \pm 3.1\%$	$18.3 \pm 2.9\%$
$\Delta\hat{\mathbf{FNR}}$	$0.2 \pm 1.7\%$	$0.3 \pm 0.1\%$	$0.1 \pm 2.3\%$
$\hat{\mathbf{FNR}}_{\mathbf{F}}$	$11.9 \pm 3.1\%$	$0.1 \pm 0.3\%$	$3.0 \pm 3.5\%$
$\hat{\mathbf{FNR}}_{\mathbf{M}}$	$11.7 \pm 1.4\%$	$0.4 \pm 0.4\%$	$2.9 \pm 1.2\%$
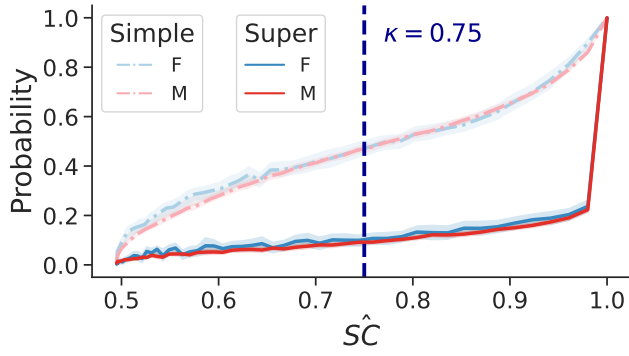
Abstention set metrics		
	Simple	**Super**
$\Delta\hat{\mathbf{AR}}$	$0.1 \pm 2.8\%$	$1.2 \pm 3.3\%$
$\hat{\mathbf{AR}}_{\mathbf{F}}$	$47.5 \pm 6.7\%$	$9.9 \pm 5.7\%$
$\hat{\mathbf{AR}}_{\mathbf{M}}$	$47.4 \pm 3.9\%$	$8.7 \pm 2.4\%$

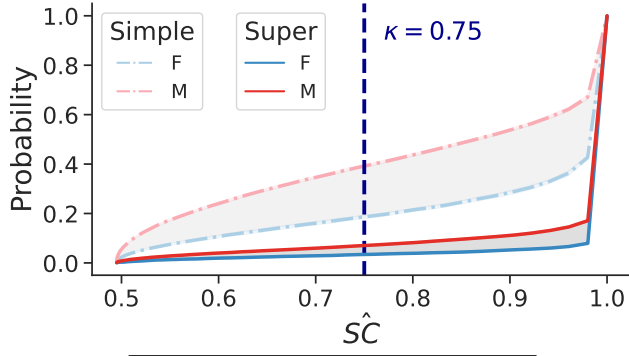*Figure 16.* **Random forests** on `German Credit`

E.4.4. TAIWAN CREDIT

$\hat{SC}$ CDFs for `Taiwan Credit` ($g = $ sex) and associated error metrics on the prediction set. **Baseline** metrics computed with $B = 101$ models. For **simple**, $B = 101$ models; for **super**, $B = 101$ ensemble models, each composed of 41 underlying models. We repeat for 10 test/train splits. We also report abstention rate $\hat{AR}$.

Logistic regression prediction set metrics			
	Baseline	**Simple**	**Super**
$\Delta\hat{\mathbf{PR}}$	$1.5 \pm 0.1\%$	$1.0 \pm 0.1\%$	$1.0 \pm 0.1\%$
$\hat{\mathbf{PR}}_{\mathbf{F}}$	$6.7 \pm 0.3\%$	$6.2 \pm 0.1\%$	$6.9 \pm 0.1\%$
$\hat{\mathbf{PR}}_{\mathbf{M}}$	$8.2 \pm 0.4\%$	$7.2 \pm 0.2\%$	$7.9 \pm 0.2\%$
$\Delta\hat{\mathbf{Err}}$	$3.1 \pm 0.1\%$	$3.1 \pm 0.3\%$	$3.2 \pm 0.3\%$
$\hat{\mathbf{Err}}_{\mathbf{F}}$	$17.8 \pm 0.5\%$	$17.0 \pm 0.2\%$	$17.5 \pm 0.3\%$
$\hat{\mathbf{Err}}_{\mathbf{M}}$	$20.9 \pm 0.4\%$	$20.1 \pm 0.5\%$	$20.7 \pm 0.6\%$
$\Delta\hat{\mathbf{FPR}}$	$0.7 \pm 0.2\%$	$0.3 \pm 0.0\%$	$0.3 \pm 0.1\%$
$\hat{\mathbf{FPR}}_{\mathbf{F}}$	$1.8 \pm 0.1\%$	$1.7 \pm 0.1\%$	$2.0 \pm 0.1\%$
$\hat{\mathbf{FPR}}_{\mathbf{M}}$	$2.5 \pm 0.3\%$	$2.0 \pm 0.1\%$	$2.3 \pm 0.2\%$
$\Delta\hat{\mathbf{FNR}}$	$2.4 \pm 0.2\%$	$2.7 \pm 0.4\%$	$2.8 \pm 0.3\%$
$\hat{\mathbf{FNR}}_{\mathbf{F}}$	$16.0 \pm 0.6\%$	$15.3 \pm 0.2\%$	$15.6 \pm 0.3\%$
$\hat{\mathbf{FNR}}_{\mathbf{M}}$	$18.4 \pm 0.4\%$	$18.0 \pm 0.6\%$	$18.4 \pm 0.6\%$
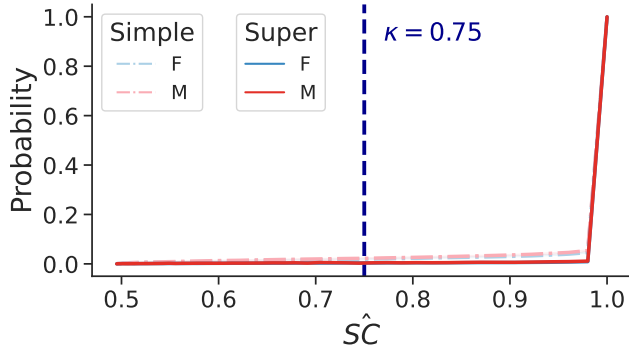
Abstention set metrics		
	Simple	**Super**
$\Delta\hat{\mathbf{AR}}$	$0.4 \pm 0.2\%$	$0.0 \pm 0.2\%$
$\hat{\mathbf{AR}}_{\mathbf{F}}$	$2.1 \pm 0.1\%$	$0.4 \pm 0.0\%$
$\hat{\mathbf{AR}}_{\mathbf{M}}$	$2.5 \pm 0.3\%$	$0.4 \pm 0.2\%$

Figure 17. **Logistic regression** on `Taiwan Credit`
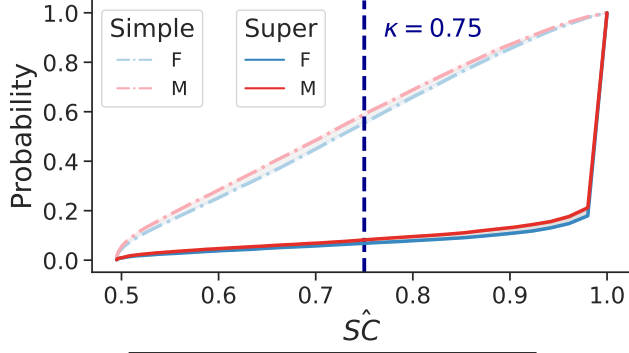
Abstention set metrics		
	Simple	**Super**
$\Delta\hat{\mathbf{AR}}$	$3.2 \pm 0.1\%$	$1.3 \pm 0.1\%$
$\hat{\mathbf{AR}}_{\mathbf{F}}$	$56.7 \pm 0.6\%$	$6.7 \pm 0.2\%$
$\hat{\mathbf{AR}}_{\mathbf{M}}$	$59.9 \pm 0.5\%$	$8.0 \pm 0.1\%$

Decision tree prediction set metrics			
	Baseline	**Simple**	**Super**
$\Delta\hat{\mathbf{PR}}$	$2.1 \pm 0.1\%$	$1.2 \pm 0.0\%$	$2.0 \pm 0.2\%$
$\hat{\mathbf{PR}}_{\mathbf{F}}$	$22.9 \pm 0.2\%$	$3.0 \pm 0.4\%$	$9.9 \pm 0.3\%$
$\hat{\mathbf{PR}}_{\mathbf{M}}$	$25.0 \pm 0.3\%$	$4.2 \pm 0.4\%$	$11.9 \pm 0.5\%$
$\Delta\hat{\mathbf{Err}}$	$2.3 \pm 0.1\%$	$1.6 \pm 0.0\%$	$2.5 \pm 0.1\%$
$\hat{\mathbf{Err}}_{\mathbf{F}}$	$26.8 \pm 0.2\%$	$9.6 \pm 0.4\%$	$15.3 \pm 0.3\%$
$\hat{\mathbf{Err}}_{\mathbf{M}}$	$29.1 \pm 0.3\%$	$11.2 \pm 0.4\%$	$17.8 \pm 0.4\%$
$\Delta\hat{\mathbf{FPR}}$	$0.6 \pm 0.1\%$	$0.2 \pm 0.1\%$	$0.7 \pm 0.2\%$
$\hat{\mathbf{FPR}}_{\mathbf{F}}$	$14.4 \pm 0.2\%$	$0.6 \pm 0.1\%$	$3.0 \pm 0.1\%$
$\hat{\mathbf{FPR}}_{\mathbf{M}}$	$15.0 \pm 0.3\%$	$0.8 \pm 0.2\%$	$3.7 \pm 0.3\%$
$\Delta\hat{\mathbf{FNR}}$	$1.7 \pm 0.2\%$	$1.3 \pm 0.1\%$	$1.9 \pm 0.1\%$
$\hat{\mathbf{FNR}}_{\mathbf{F}}$	$12.4 \pm 0.4\%$	$9.0 \pm 0.4\%$	$12.3 \pm 0.3\%$
$\hat{\mathbf{FNR}}_{\mathbf{M}}$	$14.1 \pm 0.2\%$	$10.3 \pm 0.5\%$	$14.2 \pm 0.4\%$

Figure 18. **Decision trees** on `Taiwan Credit`
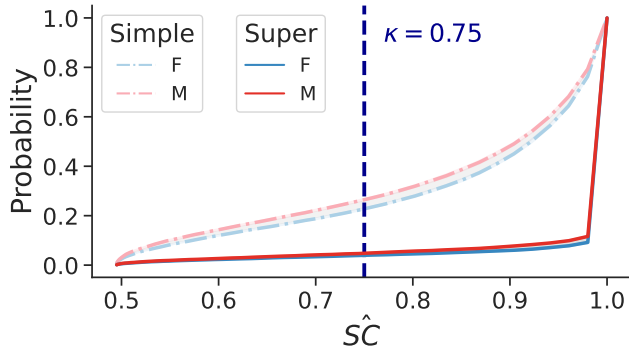
Abstention set metrics		
	Simple	**Super**
$\Delta\hat{\mathbf{AR}}$	$4.1 \pm 0.1\%$	$0.8 \pm 0.0\%$
$\hat{\mathbf{AR}}_{\mathbf{F}}$	$24.0 \pm 0.8\%$	$3.9 \pm 0.3\%$
$\hat{\mathbf{AR}}_{\mathbf{M}}$	$28.1 \pm 0.7\%$	$4.7 \pm 0.3\%$

Random forest prediction set metrics			
	Baseline	**Simple**	**Super**
$\Delta\hat{\mathbf{PR}}$	$2.5 \pm 0.1\%$	$1.0 \pm 0.1\%$	$2.1 \pm 0.2\%$
$\hat{\mathbf{PR}}_{\mathbf{F}}$	$14.9 \pm 0.2\%$	$4.1 \pm 0.3\%$	$10.3 \pm 0.2\%$
$\hat{\mathbf{PR}}_{\mathbf{M}}$	$17.4 \pm 0.3\%$	$5.1 \pm 0.2\%$	$12.4 \pm 0.4\%$
$\Delta\hat{\mathbf{Err}}$	$2.8 \pm 0.1\%$	$1.9 \pm 0.0\%$	$2.5 \pm 0.1\%$
$\hat{\mathbf{Err}}_{\mathbf{F}}$	$20.5 \pm 0.3\%$	$12.0 \pm 0.4\%$	$15.8 \pm 0.4\%$
$\hat{\mathbf{Err}}_{\mathbf{M}}$	$23.3 \pm 0.4\%$	$13.9 \pm 0.4\%$	$18.3 \pm 0.5\%$
$\Delta\hat{\mathbf{FPR}}$	$1.0 \pm 0.1\%$	$0.3 \pm 0.1\%$	$0.6 \pm 0.1\%$
$\hat{\mathbf{FPR}}_{\mathbf{F}}$	$7.2 \pm 0.2\%$	$0.9 \pm 0.1\%$	$3.3 \pm 0.1\%$
$\hat{\mathbf{FPR}}_{\mathbf{M}}$	$8.2 \pm 0.3\%$	$1.2 \pm 0.2\%$	$3.9 \pm 0.2\%$
$\Delta\hat{\mathbf{FNR}}$	$1.7 \pm 0.1\%$	$1.6 \pm 0.1\%$	$1.8 \pm 0.0\%$
$\hat{\mathbf{FNR}}_{\mathbf{F}}$	$13.3 \pm 0.4\%$	$11.0 \pm 0.3\%$	$12.6 \pm 0.4\%$
$\hat{\mathbf{FNR}}_{\mathbf{M}}$	$15.0 \pm 0.3\%$	$12.6 \pm 0.4\%$	$14.4 \pm 0.4\%$

Figure 19. **Random forests** on `Taiwan Credit`

E.4.5. NEW ADULT – CA

$\hat{SC}$ CDFs for three tasks (Income, Employment, Public Coverage) in New Adult – CA, using $g =$ sex and race, and associated error metrics on the prediction set. **Baseline** metrics computed with $B = 101$ models. For **simple**, $B = 101$ models; for **super**, $B = 101$ ensemble models, each composed of 21 underlying models for Income and Public Coverage; 15 for Employment. We repeat for 5 test/train splits. We also report abstention rate $\hat{AR}$.

Income - by sex.

Logistic regression prediction set metrics			
	Baseline	**Simple**	**Super**
$\Delta\hat{PR}$	$2.7 \pm 0.1\%$	$2.9 \pm 0.1\%$	$2.8 \pm 0.1\%$
$\hat{PR}_F$	$38.4 \pm 0.2\%$	$38.1 \pm 0.2\%$	$38.2 \pm 0.2\%$
$\hat{PR}_M$	$41.1 \pm 0.3\%$	$41.0 \pm 0.1\%$	$41.0 \pm 0.1\%$
$\Delta\hat{Err}$	$0.9 \pm 0.0\%$	$1.0 \pm 0.2\%$	$1.0 \pm 0.2\%$
$\hat{Err}_F$	$21.5 \pm 0.2\%$	$21.1 \pm 0.3\%$	$21.3 \pm 0.3\%$
$\hat{Err}_M$	$22.4 \pm 0.2\%$	$22.1 \pm 0.1\%$	$22.3 \pm 0.1\%$
$\Delta\hat{FPR}$	$4.0 \pm 0.1\%$	$3.9 \pm 0.0\%$	$3.9 \pm 0.0\%$
$\hat{FPR}_F$	$12.5 \pm 0.2\%$	$12.2 \pm 0.1\%$	$12.3 \pm 0.1\%$
$\hat{FPR}_M$	$8.5 \pm 0.1\%$	$8.3 \pm 0.1\%$	$8.4 \pm 0.1\%$
$\Delta\hat{FNR}$	$4.9 \pm 0.0\%$	$4.9 \pm 0.1\%$	$4.8 \pm 0.1\%$
$\hat{FNR}_F$	$9.0 \pm 0.1\%$	$8.9 \pm 0.2\%$	$9.1 \pm 0.2\%$
$\hat{FNR}_M$	$13.9 \pm 0.1\%$	$13.8 \pm 0.1\%$	$13.9 \pm 0.1\%$
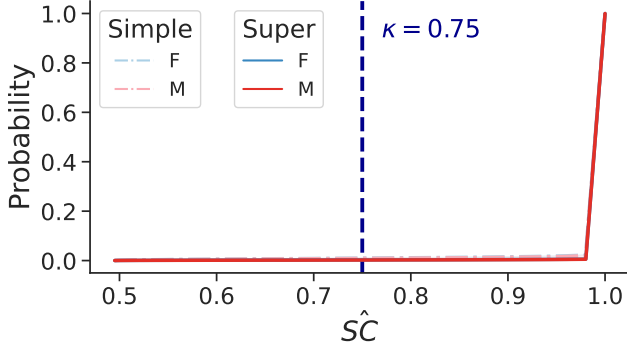
Abstention set metrics		
	Simple	**Super**
$\Delta\hat{AR}$	$0.1 \pm 0.0\%$	$0.1 \pm 0.0\%$
$\hat{AR}_F$	$1.0 \pm 0.0\%$	$0.3 \pm 0.0\%$
$\hat{AR}_M$	$0.9 \pm 0.0\%$	$0.2 \pm 0.0\%$

Figure 20. **Logistic regression** on New Adult – CA – Income, by sex

Decision tree prediction set metrics			
	Baseline	**Simple**	**Super**
$\Delta\hat{PR}$	$7.5 \pm 0.1\%$	$12.5 \pm 0.1\%$	$9.7 \pm 0.1\%$
$\hat{PR}_F$	$37.4 \pm 0.2\%$	$26.8 \pm 0.4\%$	$34.1 \pm 0.3\%$
$\hat{PR}_M$	$44.9 \pm 0.1\%$	$39.3 \pm 0.3\%$	$43.8 \pm 0.2\%$
$\Delta\hat{Err}$	$1.4 \pm 0.0\%$	$1.0 \pm 0.0\%$	$1.4 \pm 0.0\%$
$\hat{Err}_F$	$24.4 \pm 0.1\%$	$6.9 \pm 0.1\%$	$14.5 \pm 0.2\%$
$\hat{Err}_M$	$25.8 \pm 0.1\%$	$7.9 \pm 0.1\%$	$15.9 \pm 0.2\%$
$\Delta\hat{FPR}$	$1.4 \pm 0.0\%$	$0.1 \pm 0.1\%$	$0.5 \pm 0.1\%$
$\hat{FPR}_F$	$13.5 \pm 0.1\%$	$3.6 \pm 0.1\%$	$7.6 \pm 0.1\%$
$\hat{FPR}_M$	$12.1 \pm 0.1\%$	$3.5 \pm 0.2\%$	$7.1 \pm 0.2\%$
$\Delta\hat{FNR}$	$2.9 \pm 0.0\%$	$1.1 \pm 0.0\%$	$1.9 \pm 0.1\%$
$\hat{FNR}_F$	$10.9 \pm 0.1\%$	$3.3 \pm 0.1\%$	$6.9 \pm 0.1\%$
$\hat{FNR}_M$	$13.8 \pm 0.1\%$	$4.4 \pm 0.1\%$	$8.8 \pm 0.2\%$
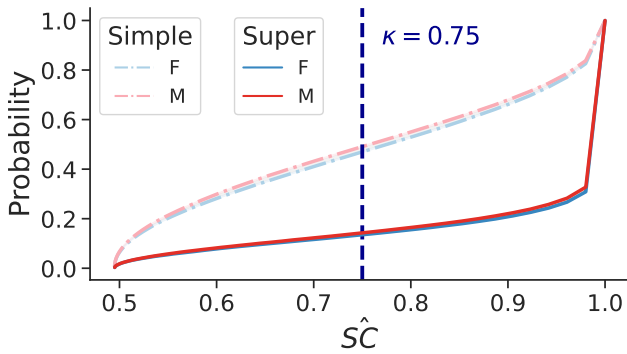
Abstention set metrics		
	Simple	**Super**
$\Delta\hat{AR}$	$2.1 \pm 0.2\%$	$0.8 \pm 0.0\%$
$\hat{AR}_F$	$49.2 \pm 0.3\%$	$13.3 \pm 0.2\%$
$\hat{AR}_M$	$51.3 \pm 0.1\%$	$14.1 \pm 0.2\%$

Figure 21. **Decision trees** on New Adult – CA – Income, by sex

Random forest prediction set metrics			
	Baseline	**Simple**	**Super**
$\Delta\hat{\text{PR}}$	$7.4 \pm 0.1\%$	$10.3 \pm 0.2\%$	$8.6 \pm 0.1\%$
$\hat{\text{PR}}_F$	$36.7 \pm 0.2\%$	$30.6 \pm 0.4\%$	$34.9 \pm 0.3\%$
$\hat{\text{PR}}_M$	$44.1 \pm 0.1\%$	$40.9 \pm 0.2\%$	$43.5 \pm 0.2\%$
$\Delta\hat{\text{Err}}$	$1.4 \pm 0.0\%$	$1.2 \pm 0.0\%$	$1.4 \pm 0.1\%$
$\hat{\text{Err}}_F$	$21.0 \pm 0.1\%$	$9.3 \pm 0.2\%$	$15.3 \pm 0.1\%$
$\hat{\text{Err}}_M$	$22.4 \pm 0.1\%$	$10.5 \pm 0.2\%$	$16.7 \pm 0.2\%$
$\Delta\hat{\text{FPR}}$	$1.4 \pm 0.0\%$	$0.5 \pm 0.0\%$	$0.9 \pm 0.1\%$
$\hat{\text{FPR}}_F$	$11.4 \pm 0.1\%$	$4.9 \pm 0.1\%$	$8.1 \pm 0.1\%$
$\hat{\text{FPR}}_M$	$10.0 \pm 0.1\%$	$4.4 \pm 0.1\%$	$7.2 \pm 0.2\%$
$\Delta\hat{\text{FNR}}$	$2.8 \pm 0.0\%$	$1.8 \pm 0.0\%$	$2.3 \pm 0.0\%$
$\hat{\text{FNR}}_F$	$9.6 \pm 0.1\%$	$4.4 \pm 0.1\%$	$7.2 \pm 0.1\%$
$\hat{\text{FNR}}_M$	$12.4 \pm 0.1\%$	$6.2 \pm 0.1\%$	$9.5 \pm 0.1\%$
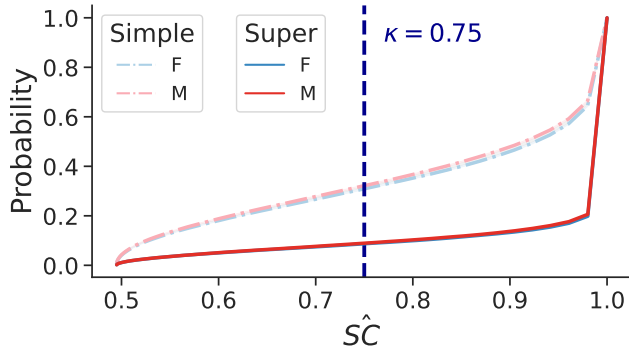
Abstention set metrics		
	Simple	**Super**
$\Delta\hat{\text{AR}}$	$1.4 \pm 0.1\%$	$0.2 \pm 0.0\%$
$\hat{\text{AR}}_F$	$32.8 \pm 0.2\%$	$8.6 \pm 0.1\%$
$\hat{\text{AR}}_M$	$34.2 \pm 0.1\%$	$8.8 \pm 0.1\%$

Figure 22. **Random forests** on `New Adult - CA - Income`, by `sex`

Income - by race.

Logistic regression prediction set metrics			
	Baseline	**Simple**	**Super**
$\Delta\hat{\text{PR}}$	$9.2 \pm 0.1\%$	$9.2 \pm 0.3\%$	$9.2 \pm 0.2\%$
$\hat{\text{PR}}_{NW}$	$34.1 \pm 0.3\%$	$33.9 \pm 0.3\%$	$34.0 \pm 0.3\%$
$\hat{\text{PR}}_W$	$43.3 \pm 0.2\%$	$43.1 \pm 0.0\%$	$43.2 \pm 0.1\%$
$\Delta\hat{\text{Err}}$	$0.6 \pm 0.1\%$	$0.4 \pm 0.1\%$	$0.4 \pm 0.1\%$
$\hat{\text{Err}}_{NW}$	$21.6 \pm 0.2\%$	$21.4 \pm 0.2\%$	$21.6 \pm 0.2\%$
$\hat{\text{Err}}_W$	$22.2 \pm 0.1\%$	$21.8 \pm 0.1\%$	$22.0 \pm 0.1\%$
$\Delta\hat{\text{FPR}}$	$0.6 \pm 0.1\%$	$0.5 \pm 0.0\%$	$0.5 \pm 0.0\%$
$\hat{\text{FPR}}_{NW}$	$10.0 \pm 0.2\%$	$9.8 \pm 0.1\%$	$9.9 \pm 0.1\%$
$\hat{\text{FPR}}_W$	$10.6 \pm 0.1\%$	$10.3 \pm 0.1\%$	$10.4 \pm 0.1\%$
$\Delta\hat{\text{FNR}}$	$0.0 \pm 0.1\%$	$0.1 \pm 0.2\%$	$0.1 \pm 0.3\%$
$\hat{\text{FNR}}_{NW}$	$11.6 \pm 0.2\%$	$11.6 \pm 0.3\%$	$11.7 \pm 0.3\%$
$\hat{\text{FNR}}_W$	$11.6 \pm 0.1\%$	$11.5 \pm 0.1\%$	$11.6 \pm 0.0\%$
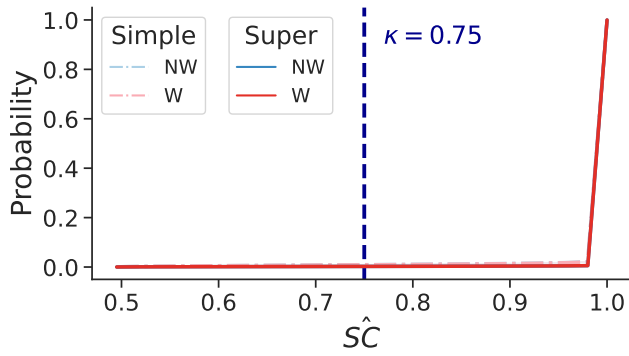
Abstention set metrics		
	Simple	**Super**
$\Delta\hat{\text{AR}}$	$0.0 \pm 0.0\%$	$0.0 \pm 0.0\%$
$\hat{\text{AR}}_{NW}$	$1.0 \pm 0.0\%$	$0.2 \pm 0.0\%$
$\hat{\text{AR}}_W$	$1.0 \pm 0.0\%$	$0.2 \pm 0.0\%$

Figure 23. **Logistic regression** on `New Adult - CA - Income`, by `race`

Decision tree prediction set metrics			
	Baseline	**Simple**	**Super**
$\Delta\hat{\mathbf{PR}}$	$7.0 \pm 0.0\%$	$10.3 \pm 0.0\%$	$9.9 \pm 0.1\%$
$\hat{\mathbf{PR}}_{NW}$	$37.0 \pm 0.2\%$	$27.0 \pm 0.2\%$	$33.1 \pm 0.3\%$
$\hat{\mathbf{PR}}_{W}$	$44.0 \pm 0.2\%$	$37.3 \pm 0.2\%$	$43.0 \pm 0.2\%$
$\Delta\hat{\mathbf{Err}}$	$1.1 \pm 0.0\%$	$0.2 \pm 0.1\%$	$0.6 \pm 0.0\%$
$\hat{\mathbf{Err}}_{NW}$	$24.5 \pm 0.1\%$	$7.3 \pm 0.0\%$	$14.8 \pm 0.2\%$
$\hat{\mathbf{Err}}_{W}$	$25.6 \pm 0.1\%$	$7.5 \pm 0.1\%$	$15.4 \pm 0.2\%$
$\Delta\hat{\mathbf{FPR}}$	$0.2 \pm 0.0\%$	$0.5 \pm 0.0\%$	$0.7 \pm 0.0\%$
$\hat{\mathbf{FPR}}_{NW}$	$12.9 \pm 0.1\%$	$3.2 \pm 0.1\%$	$6.9 \pm 0.1\%$
$\hat{\mathbf{FPR}}_{W}$	$12.7 \pm 0.1\%$	$3.7 \pm 0.1\%$	$7.6 \pm 0.1\%$
$\Delta\hat{\mathbf{FNR}}$	$1.4 \pm 0.0\%$	$0.3 \pm 0.0\%$	$0.2 \pm 0.0\%$
$\hat{\mathbf{FNR}}_{NW}$	$11.6 \pm 0.1\%$	$4.1 \pm 0.1\%$	$8.0 \pm 0.1\%$
$\hat{\mathbf{FNR}}_{W}$	$13.0 \pm 0.1\%$	$3.8 \pm 0.1\%$	$7.8 \pm 0.1\%$
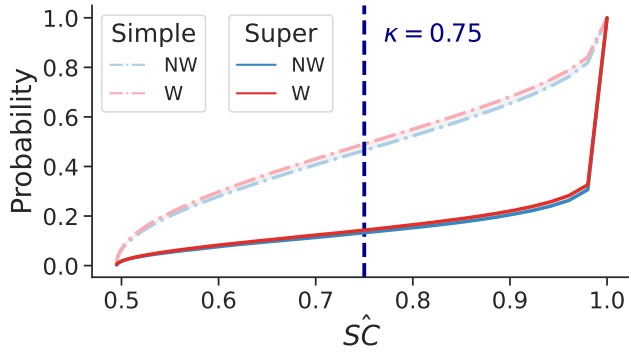
Abstention set metrics		
	Simple	**Super**
$\Delta\hat{\mathbf{AR}}$	$2.5 \pm 0.1\%$	$1.0 \pm 0.1\%$
$\hat{\mathbf{AR}}_{NW}$	$48.8 \pm 0.3\%$	$13.1 \pm 0.2\%$
$\hat{\mathbf{AR}}_{W}$	$51.3 \pm 0.2\%$	$14.1 \pm 0.1\%$

Figure 24. **Decision trees** on `New Adult - CA - Income,` by `race`

Random forest prediction set metrics			
	Baseline	**Simple**	**Super**
$\Delta\hat{\mathbf{PR}}$	$8.4 \pm 0.0\%$	$11.0 \pm 0.0\%$	$10.1 \pm 0.2\%$
$\hat{\mathbf{PR}}_{NW}$	$35.4 \pm 0.2\%$	$29.3 \pm 0.2\%$	$33.2 \pm 0.3\%$
$\hat{\mathbf{PR}}_{W}$	$43.8 \pm 0.2\%$	$40.3 \pm 0.2\%$	$43.3 \pm 0.1\%$
$\Delta\hat{\mathbf{Err}}$	$1.2 \pm 0.0\%$	$0.2 \pm 0.1\%$	$0.5 \pm 0.0\%$
$\hat{\mathbf{Err}}_{NW}$	$21.0 \pm 0.1\%$	$9.8 \pm 0.1\%$	$15.7 \pm 0.2\%$
$\hat{\mathbf{Err}}_{W}$	$22.2 \pm 0.1\%$	$10.0 \pm 0.2\%$	$16.2 \pm 0.2\%$
$\Delta\hat{\mathbf{FPR}}$	$0.6 \pm 0.0\%$	$0.7 \pm 0.0\%$	$0.9 \pm 0.1\%$
$\hat{\mathbf{FPR}}_{NW}$	$10.3 \pm 0.1\%$	$4.2 \pm 0.1\%$	$7.1 \pm 0.2\%$
$\hat{\mathbf{FPR}}_{W}$	$10.9 \pm 0.1\%$	$4.9 \pm 0.1\%$	$8.0 \pm 0.1\%$
$\Delta\hat{\mathbf{FNR}}$	$0.7 \pm 0.0\%$	$0.5 \pm 0.0\%$	$0.3 \pm 0.1\%$
$\hat{\mathbf{FNR}}_{NW}$	$10.7 \pm 0.1\%$	$5.6 \pm 0.1\%$	$8.6 \pm 0.2\%$
$\hat{\mathbf{FNR}}_{W}$	$11.4 \pm 0.1\%$	$5.1 \pm 0.1\%$	$8.3 \pm 0.1\%$
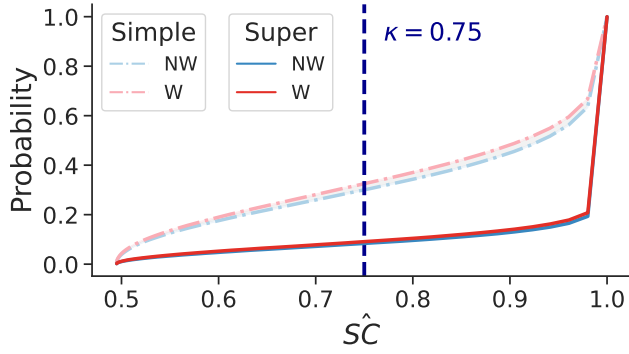
Abstention set metrics		
	Simple	**Super**
$\Delta\hat{\mathbf{AR}}$	$2.4 \pm 0.2\%$	$0.8 \pm 0.0\%$
$\hat{\mathbf{AR}}_{NW}$	$32.1 \pm 0.3\%$	$8.2 \pm 0.1\%$
$\hat{\mathbf{AR}}_{W}$	$34.5 \pm 0.1\%$	$9.0 \pm 0.1\%$

Figure 25. **Random forests** on `New Adult - CA - Income,` by `race`

Employment - by sex.

Logistic regression prediction set metrics			
	Baseline	**Simple**	**Super**
$\Delta\hat{\mathbf{PR}}$	$4.3 \pm 0.1\%$	$4.5 \pm 0.0\%$	$4.4 \pm 0.0\%$
$\hat{\mathbf{PR}}_F$	$56.6 \pm 0.1\%$	$56.8 \pm 0.1\%$	$56.7 \pm 0.1\%$
$\hat{\mathbf{PR}}_M$	$52.3 \pm 0.2\%$	$52.3 \pm 0.1\%$	$52.3 \pm 0.1\%$
$\Delta\hat{\mathbf{Err}}$	$5.0 \pm 0.0\%$	$4.9 \pm 0.0\%$	$4.9 \pm 0.0\%$
$\hat{\mathbf{Err}}_F$	$25.8 \pm 0.1\%$	$25.5 \pm 0.1\%$	$25.6 \pm 0.1\%$
$\hat{\mathbf{Err}}_M$	$20.8 \pm 0.1\%$	$20.6 \pm 0.1\%$	$20.7 \pm 0.1\%$
$\Delta\hat{\mathbf{FPR}}$	$8.1 \pm 0.0\%$	$8.1 \pm 0.0\%$	$8.1 \pm 0.0\%$
$\hat{\mathbf{FPR}}_F$	$20.1 \pm 0.1\%$	$20.1 \pm 0.0\%$	$20.1 \pm 0.0\%$
$\hat{\mathbf{FPR}}_M$	$12.0 \pm 0.1\%$	$12.0 \pm 0.0\%$	$12.0 \pm 0.0\%$
$\Delta\hat{\mathbf{FNR}}$	$3.1 \pm 0.0\%$	$3.2 \pm 0.1\%$	$3.2 \pm 0.1\%$
$\hat{\mathbf{FNR}}_F$	$5.7 \pm 0.1\%$	$5.4 \pm 0.0\%$	$5.5 \pm 0.0\%$
$\hat{\mathbf{FNR}}_M$	$8.8 \pm 0.1\%$	$8.6 \pm 0.1\%$	$8.7 \pm 0.1\%$
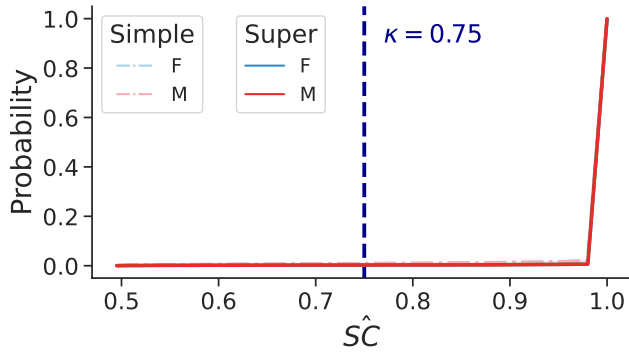
Abstention set metrics		
	Simple	**Super**
$\Delta\hat{\mathbf{AR}}$	$0.1 \pm 0.0\%$	$0.0 \pm 0.0\%$
$\hat{\mathbf{AR}}_F$	$0.8 \pm 0.0\%$	$0.3 \pm 0.0\%$
$\hat{\mathbf{AR}}_M$	$0.7 \pm 0.0\%$	$0.3 \pm 0.0\%$

Figure 26. **Logistic regression** on `New Adult – CA – Employment`, by `sex`

Decision tree prediction set metrics			
	Baseline	**Simple**	**Super**
$\Delta\hat{\mathbf{PR}}$	$0.5 \pm 0.0\%$	$0.7 \pm 0.2\%$	$0.6 \pm 0.1\%$
$\hat{\mathbf{PR}}_F$	$50.3 \pm 0.2\%$	$50.0 \pm 0.3\%$	$51.2 \pm 0.2\%$
$\hat{\mathbf{PR}}_M$	$49.8 \pm 0.2\%$	$49.3 \pm 0.1\%$	$50.6 \pm 0.1\%$
$\Delta\hat{\mathbf{Err}}$	$4.8 \pm 0.0\%$	$5.8 \pm 0.1\%$	$5.8 \pm 0.1\%$
$\hat{\mathbf{Err}}_F$	$24.8 \pm 0.1\%$	$17.8 \pm 0.1\%$	$20.9 \pm 0.2\%$
$\hat{\mathbf{Err}}_M$	$20.0 \pm 0.1\%$	$12.0 \pm 0.0\%$	$15.1 \pm 0.1\%$
$\Delta\hat{\mathbf{FPR}}$	$6.1 \pm 0.0\%$	$6.2 \pm 0.1\%$	$6.4 \pm 0.1\%$
$\hat{\mathbf{FPR}}_F$	$16.5 \pm 0.1\%$	$13.8 \pm 0.1\%$	$15.2 \pm 0.1\%$
$\hat{\mathbf{FPR}}_M$	$10.4 \pm 0.1\%$	$7.6 \pm 0.0\%$	$8.8 \pm 0.0\%$
$\Delta\hat{\mathbf{FNR}}$	$1.3 \pm 0.0\%$	$0.3 \pm 0.0\%$	$0.7 \pm 0.1\%$
$\hat{\mathbf{FNR}}_F$	$8.3 \pm 0.1\%$	$4.0 \pm 0.1\%$	$5.6 \pm 0.2\%$
$\hat{\mathbf{FNR}}_M$	$9.6 \pm 0.1\%$	$4.3 \pm 0.1\%$	$6.3 \pm 0.1\%$
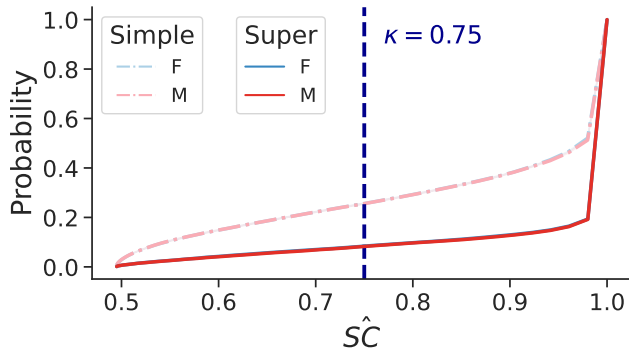
Abstention set metrics		
	Simple	**Super**
$\Delta\hat{\mathbf{AR}}$	$0.2 \pm 0.0\%$	$0.1 \pm 0.1\%$
$\hat{\mathbf{AR}}_F$	$22.5 \pm 0.1\%$	$8.2 \pm 0.2\%$
$\hat{\mathbf{AR}}_M$	$22.3 \pm 0.1\%$	$8.1 \pm 0.3\%$

Figure 27. **Decision trees** on `New Adult – CA – Employment`, by `sex`

Random forest prediction set metrics			
	Baseline	**Simple**	**Super**
$\Delta\hat{\mathbf{PR}}$	$0.3 \pm 0.1\%$	$0.5 \pm 0.1\%$	$0.6 \pm 0.1\%$
$\hat{\mathbf{PR}}_{\mathbf{F}}$	$49.2 \pm 0.1\%$	$48.7 \pm 0.2\%$	$50.3 \pm 0.2\%$
$\hat{\mathbf{PR}}_{\mathbf{M}}$	$48.9 \pm 0.2\%$	$48.2 \pm 0.1\%$	$49.7 \pm 0.1\%$
$\Delta\hat{\mathbf{Err}}$	$4.8 \pm 0.0\%$	$5.4 \pm 0.0\%$	$5.5 \pm 0.2\%$
$\hat{\mathbf{Err}}_{\mathbf{F}}$	$24.0 \pm 0.1\%$	$17.5 \pm 0.1\%$	$20.5 \pm 0.2\%$
$\hat{\mathbf{Err}}_{\mathbf{M}}$	$19.2 \pm 0.1\%$	$12.1 \pm 0.1\%$	$15.0 \pm 0.0\%$
$\Delta\hat{\mathbf{FPR}}$	$6.0 \pm 0.0\%$	$5.9 \pm 0.1\%$	$6.2 \pm 0.1\%$
$\hat{\mathbf{FPR}}_{\mathbf{F}}$	$15.5 \pm 0.1\%$	$13.2 \pm 0.1\%$	$14.7 \pm 0.1\%$
$\hat{\mathbf{FPR}}_{\mathbf{M}}$	$9.5 \pm 0.1\%$	$7.3 \pm 0.0\%$	$8.5 \pm 0.0\%$
$\Delta\hat{\mathbf{FNR}}$	$1.2 \pm 0.0\%$	$0.4 \pm 0.0\%$	$0.8 \pm 0.0\%$
$\hat{\mathbf{FNR}}_{\mathbf{F}}$	$8.5 \pm 0.1\%$	$4.3 \pm 0.1\%$	$5.8 \pm 0.1\%$
$\hat{\mathbf{FNR}}_{\mathbf{M}}$	$9.7 \pm 0.1\%$	$4.7 \pm 0.1\%$	$6.6 \pm 0.1\%$
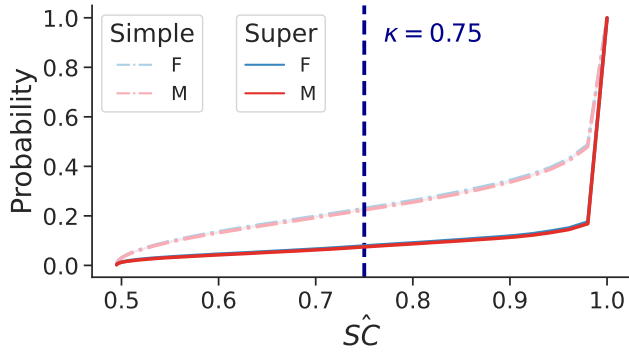
Abstention set metrics		
	Simple	**Super**
$\Delta\hat{\mathbf{AR}}$	$0.7 \pm 0.0\%$	$0.4 \pm 0.0\%$
$\hat{\mathbf{AR}}_{\mathbf{F}}$	$20.3 \pm 0.2\%$	$7.7 \pm 0.2\%$
$\hat{\mathbf{AR}}_{\mathbf{M}}$	$19.6 \pm 0.2\%$	$7.3 \pm 0.2\%$

Figure 28. **Random forests** on `New Adult - CA - Employment`, by `sex`

Employment - by race.

Logistic regression prediction set metrics			
	Baseline	**Simple**	**Super**
$\Delta\hat{\mathbf{PR}}$	$1.1 \pm 0.2\%$	$1.2 \pm 0.1\%$	$1.2 \pm 0.1\%$
$\hat{\mathbf{PR}}_{\mathbf{NW}}$	$55.2 \pm 0.3\%$	$55.3 \pm 0.2\%$	$55.3 \pm 0.2\%$
$\hat{\mathbf{PR}}_{\mathbf{W}}$	$54.1 \pm 0.1\%$	$54.1 \pm 0.1\%$	$54.1 \pm 0.1\%$
$\Delta\hat{\mathbf{Err}}$	$0.1 \pm 0.0\%$	$0.2 \pm 0.1\%$	$0.2 \pm 0.0\%$
$\hat{\mathbf{Err}}_{\mathbf{NW}}$	$23.3 \pm 0.1\%$	$23.0 \pm 0.0\%$	$23.1 \pm 0.1\%$
$\hat{\mathbf{Err}}_{\mathbf{W}}$	$23.4 \pm 0.1\%$	$23.2 \pm 0.1\%$	$23.3 \pm 0.1\%$
$\Delta\hat{\mathbf{FPR}}$	$0.8 \pm 0.0\%$	$0.7 \pm 0.1\%$	$0.7 \pm 0.1\%$
$\hat{\mathbf{FPR}}_{\mathbf{NW}}$	$16.6 \pm 0.1\%$	$16.5 \pm 0.0\%$	$16.6 \pm 0.0\%$
$\hat{\mathbf{FPR}}_{\mathbf{W}}$	$15.8 \pm 0.1\%$	$15.8 \pm 0.1\%$	$15.9 \pm 0.1\%$
$\Delta\hat{\mathbf{FNR}}$	$1.0 \pm 0.0\%$	$1.0 \pm 0.0\%$	$0.9 \pm 0.0\%$
$\hat{\mathbf{FNR}}_{\mathbf{NW}}$	$6.6 \pm 0.1\%$	$6.4 \pm 0.1\%$	$6.5 \pm 0.1\%$
$\hat{\mathbf{FNR}}_{\mathbf{W}}$	$7.6 \pm 0.1\%$	$7.4 \pm 0.1\%$	$7.4 \pm 0.1\%$
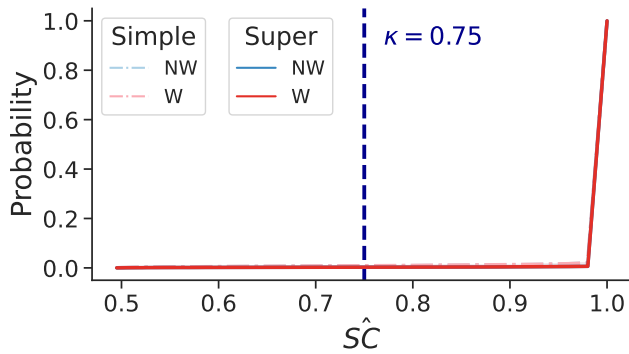
Abstention set metrics		
	Simple	**Super**
$\Delta\hat{\mathbf{AR}}$	$0.1 \pm 0.0\%$	$0.0 \pm 0.0\%$
$\hat{\mathbf{AR}}_{\mathbf{NW}}$	$0.7 \pm 0.0\%$	$0.3 \pm 0.0\%$
$\hat{\mathbf{AR}}_{\mathbf{W}}$	$0.8 \pm 0.0\%$	$0.3 \pm 0.0\%$

Figure 29. **Logistic regression** on `New Adult - CA - Employment`, by `race`

Decision tree prediction set metrics			
	Baseline	**Simple**	**Super**
$\Delta\hat{\text{PR}}$	$0.2 \pm 0.2\%$	$0.6 \pm 0.0\%$	$1.2 \pm 0.1\%$
$\hat{\text{PR}}_{\text{NW}}$	$50.2 \pm 0.3\%$	$50.0 \pm 0.1\%$	$51.6 \pm 0.0\%$
$\hat{\text{PR}}_{\text{W}}$	$50.0 \pm 0.1\%$	$49.4 \pm 0.1\%$	$50.4 \pm 0.1\%$
$\Delta\hat{\text{Err}}$	$0.6 \pm 0.0\%$	$0.7 \pm 0.0\%$	$0.6 \pm 0.0\%$
$\hat{\text{Err}}_{\text{NW}}$	$22.1 \pm 0.1\%$	$14.5 \pm 0.1\%$	$17.7 \pm 0.1\%$
$\hat{\text{Err}}_{\text{W}}$	$22.7 \pm 0.1\%$	$15.2 \pm 0.1\%$	$18.3 \pm 0.1\%$
$\Delta\hat{\text{FPR}}$	$0.1 \pm 0.0\%$	$0.6 \pm 0.1\%$	$0.7 \pm 0.1\%$
$\hat{\text{FPR}}_{\text{NW}}$	$13.5 \pm 0.1\%$	$11.1 \pm 0.0\%$	$12.5 \pm 0.0\%$
$\hat{\text{FPR}}_{\text{W}}$	$13.4 \pm 0.1\%$	$10.5 \pm 0.1\%$	$11.8 \pm 0.1\%$
$\Delta\hat{\text{FNR}}$	$0.6 \pm 0.0\%$	$1.3 \pm 0.0\%$	$1.3 \pm 0.0\%$
$\hat{\text{FNR}}_{\text{NW}}$	$8.6 \pm 0.1\%$	$3.4 \pm 0.1\%$	$5.2 \pm 0.1\%$
$\hat{\text{FNR}}_{\text{W}}$	$9.2 \pm 0.1\%$	$4.7 \pm 0.1\%$	$6.5 \pm 0.1\%$
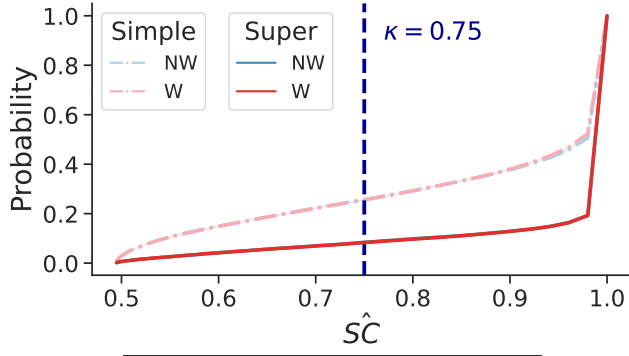
Abstention set metrics		
	Simple	**Super**
$\Delta\hat{\text{AR}}$	$0.0 \pm 0.1\%$	$0.1 \pm 0.1\%$
$\hat{\text{AR}}_{\text{NW}}$	$22.4 \pm 0.2\%$	$8.2 \pm 0.2\%$
$\hat{\text{AR}}_{\text{W}}$	$22.4 \pm 0.1\%$	$8.1 \pm 0.3\%$

Figure 30. **Decision trees** on New Adult - CA - Employment, by race

Random forest prediction set metrics			
	Baseline	**Simple**	**Super**
$\Delta\hat{\text{PR}}$	$0.6 \pm 0.2\%$	$1.0 \pm 0.0\%$	$1.4 \pm 0.1\%$
$\hat{\text{PR}}_{\text{NW}}$	$49.4 \pm 0.3\%$	$49.1 \pm 0.1\%$	$50.9 \pm 0.0\%$
$\hat{\text{PR}}_{\text{W}}$	$48.8 \pm 0.1\%$	$48.1 \pm 0.1\%$	$49.5 \pm 0.1\%$
$\Delta\hat{\text{Err}}$	$0.5 \pm 0.0\%$	$0.7 \pm 0.0\%$	$0.5 \pm 0.0\%$
$\hat{\text{Err}}_{\text{NW}}$	$21.3 \pm 0.1\%$	$14.4 \pm 0.1\%$	$17.5 \pm 0.1\%$
$\hat{\text{Err}}_{\text{W}}$	$21.8 \pm 0.1\%$	$15.1 \pm 0.1\%$	$18.0 \pm 0.1\%$
$\Delta\hat{\text{FPR}}$	$0.3 \pm 0.1\%$	$0.6 \pm 0.1\%$	$0.8 \pm 0.1\%$
$\hat{\text{FPR}}_{\text{NW}}$	$12.7 \pm 0.2\%$	$10.7 \pm 0.0\%$	$12.1 \pm 0.0\%$
$\hat{\text{FPR}}_{\text{W}}$	$12.4 \pm 0.1\%$	$10.1 \pm 0.1\%$	$11.3 \pm 0.1\%$
$\Delta\hat{\text{FNR}}$	$0.8 \pm 0.0\%$	$1.3 \pm 0.0\%$	$1.3 \pm 0.0\%$
$\hat{\text{FNR}}_{\text{NW}}$	$8.6 \pm 0.1\%$	$3.7 \pm 0.1\%$	$5.4 \pm 0.1\%$
$\hat{\text{FNR}}_{\text{W}}$	$9.4 \pm 0.1\%$	$5.0 \pm 0.1\%$	$6.7 \pm 0.1\%$
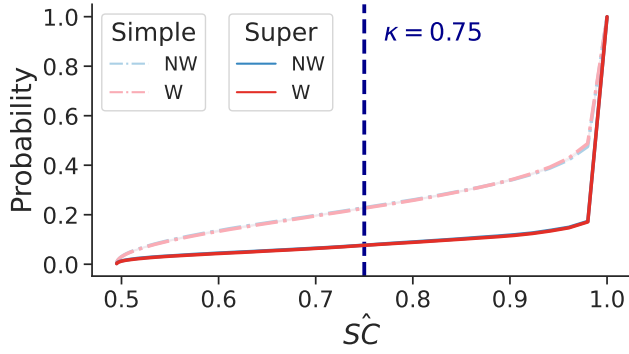
Abstention set metrics		
	Simple	**Super**
$\Delta\hat{\text{AR}}$	$0.3 \pm 0.1\%$	$0.1 \pm 0.1\%$
$\hat{\text{AR}}_{\text{NW}}$	$20.1 \pm 0.1\%$	$7.6 \pm 0.1\%$
$\hat{\text{AR}}_{\text{W}}$	$19.8 \pm 0.2\%$	$7.5 \pm 0.2\%$

Figure 31. **Random forests** on New Adult - CA - Employment, by race

Public Coverage - by sex.
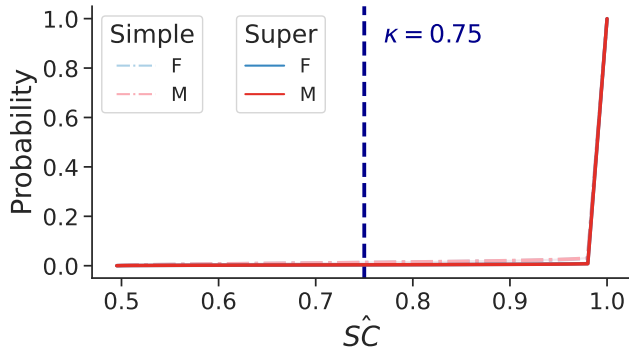
Abstention set metrics		
	Simple	**Super**
$\triangle\hat{\mathbf{AR}}$	$0.0 \pm 0.1\%$	$0.0 \pm 0.0\%$
$\hat{\mathbf{AR}}_{\mathbf{F}}$	$1.4 \pm 0.1\%$	$0.4 \pm 0.0\%$
$\hat{\mathbf{AR}}_{\mathbf{M}}$	$1.4 \pm 0.0\%$	$0.4 \pm 0.0\%$

Logistic regression prediction set metrics			
	Baseline	**Simple**	**Super**
$\triangle\hat{\mathbf{PR}}$	$2.6 \pm 0.1\%$	$2.6 \pm 0.1\%$	$2.6 \pm 0.1\%$
$\hat{\mathbf{PR}}_{\mathbf{F}}$	$15.1 \pm 0.2\%$	$14.7 \pm 0.2\%$	$15.1 \pm 0.2\%$
$\hat{\mathbf{PR}}_{\mathbf{M}}$	$17.7 \pm 0.3\%$	$17.3 \pm 0.1\%$	$17.7 \pm 0.1\%$
$\triangle\hat{\mathbf{Err}}$	$0.9 \pm 0.1\%$	$0.5 \pm 0.0\%$	$0.5 \pm 0.0\%$
$\hat{\mathbf{Err}}_{\mathbf{F}}$	$31.2 \pm 0.3\%$	$30.8 \pm 0.2\%$	$31.0 \pm 0.2\%$
$\hat{\mathbf{Err}}_{\mathbf{M}}$	$32.1 \pm 0.2\%$	$31.3 \pm 0.2\%$	$31.5 \pm 0.2\%$
$\triangle\hat{\mathbf{FPR}}$	$0.0 \pm 0.0\%$	$0.0 \pm 0.0\%$	$0.0 \pm 0.1\%$
$\hat{\mathbf{FPR}}_{\mathbf{F}}$	$5.5 \pm 0.1\%$	$5.1 \pm 0.1\%$	$5.3 \pm 0.1\%$
$\hat{\mathbf{FPR}}_{\mathbf{M}}$	$5.5 \pm 0.1\%$	$5.1 \pm 0.1\%$	$5.3 \pm 0.2\%$
$\triangle\hat{\mathbf{FNR}}$	$0.9 \pm 0.1\%$	$0.5 \pm 0.0\%$	$0.5 \pm 0.0\%$
$\hat{\mathbf{FNR}}_{\mathbf{F}}$	$25.7 \pm 0.3\%$	$25.7 \pm 0.2\%$	$25.7 \pm 0.2\%$
$\hat{\mathbf{FNR}}_{\mathbf{M}}$	$26.6 \pm 0.2\%$	$26.2 \pm 0.2\%$	$26.2 \pm 0.2\%$

Figure 32. **Logistic regression** on `New Adult - CA - Public Coverage`, by `sex`
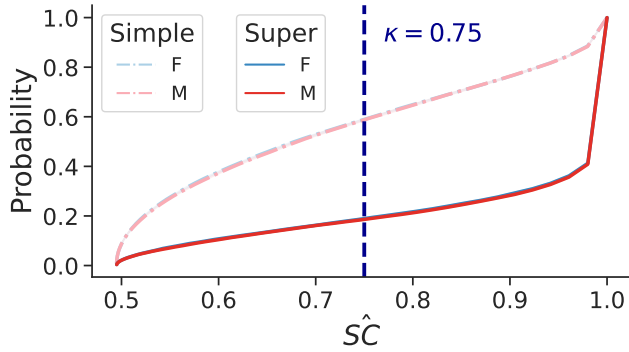
Abstention set metrics		
	Simple	**Super**
$\triangle\hat{\mathbf{AR}}$	$0.1 \pm 0.0\%$	$0.3 \pm 0.1\%$
$\hat{\mathbf{AR}}_{\mathbf{F}}$	$60.7 \pm 0.4\%$	$18.6 \pm 0.4\%$
$\hat{\mathbf{AR}}_{\mathbf{M}}$	$60.8 \pm 0.4\%$	$18.3 \pm 0.3\%$

Decision tree prediction set metrics			
	Baseline	**Simple**	**Super**
$\triangle\hat{\mathbf{PR}}$	$2.6 \pm 0.2\%$	$6.8 \pm 0.1\%$	$3.0 \pm 0.0\%$
$\hat{\mathbf{PR}}_{\mathbf{F}}$	$35.5 \pm 0.2\%$	$20.5 \pm 0.4\%$	$27.5 \pm 0.3\%$
$\hat{\mathbf{PR}}_{\mathbf{M}}$	$38.1 \pm 0.4\%$	$27.3 \pm 0.3\%$	$30.5 \pm 0.3\%$
$\triangle\hat{\mathbf{Err}}$	$0.1 \pm 0.1\%$	$0.5 \pm 0.1\%$	$0.2 \pm 0.1\%$
$\hat{\mathbf{Err}}_{\mathbf{F}}$	$35.1 \pm 0.2\%$	$18.8 \pm 0.3\%$	$26.7 \pm 0.3\%$
$\hat{\mathbf{Err}}_{\mathbf{M}}$	$35.2 \pm 0.1\%$	$19.3 \pm 0.2\%$	$26.9 \pm 0.4\%$
$\triangle\hat{\mathbf{FPR}}$	$0.4 \pm 0.1\%$	$0.1 \pm 0.0\%$	$0.3 \pm 0.1\%$
$\hat{\mathbf{FPR}}_{\mathbf{F}}$	$17.6 \pm 0.1\%$	$4.6 \pm 0.2\%$	$10.0 \pm 0.3\%$
$\hat{\mathbf{FPR}}_{\mathbf{M}}$	$17.2 \pm 0.2\%$	$4.5 \pm 0.2\%$	$9.7 \pm 0.2\%$
$\triangle\hat{\mathbf{FNR}}$	$0.6 \pm 0.0\%$	$0.6 \pm 0.2\%$	$0.5 \pm 0.3\%$
$\hat{\mathbf{FNR}}_{\mathbf{F}}$	$17.4 \pm 0.2\%$	$14.2 \pm 0.3\%$	$16.7 \pm 0.1\%$
$\hat{\mathbf{FNR}}_{\mathbf{M}}$	$18.0 \pm 0.2\%$	$14.8 \pm 0.1\%$	$17.2 \pm 0.4\%$

Figure 33. **Decision trees** on `New Adult - CA - Public Coverage`, by `sex`

Random forest prediction set metrics			
	Baseline	**Simple**	**Super**
$\Delta\hat{PR}$	$2.5 \pm 0.1\%$	$5.3 \pm 0.0\%$	$2.6 \pm 0.1\%$
$\hat{PR}_F$	$31.9 \pm 0.3\%$	$19.5 \pm 0.3\%$	$25.7 \pm 0.4\%$
$\hat{PR}_M$	$34.4 \pm 0.4\%$	$24.8 \pm 0.3\%$	$28.3 \pm 0.3\%$
$\Delta\hat{Err}$	$0.4 \pm 0.1\%$	$1.0 \pm 0.2\%$	$0.3 \pm 0.1\%$
$\hat{Err}_F$	$32.3 \pm 0.2\%$	$19.3 \pm 0.3\%$	$26.3 \pm 0.2\%$
$\hat{Err}_M$	$32.7 \pm 0.1\%$	$20.3 \pm 0.1\%$	$26.6 \pm 0.3\%$
$\Delta\hat{FPR}$	$0.2 \pm 0.1\%$	$0.2 \pm 0.1\%$	$0.3 \pm 0.1\%$
$\hat{FPR}_F$	$14.4 \pm 0.1\%$	$4.1 \pm 0.2\%$	$8.7 \pm 0.3\%$
$\hat{FPR}_M$	$14.2 \pm 0.2\%$	$4.3 \pm 0.1\%$	$8.4 \pm 0.2\%$
$\Delta\hat{FNR}$	$0.7 \pm 0.1\%$	$0.8 \pm 0.0\%$	$0.7 \pm 0.1\%$
$\hat{FNR}_F$	$17.9 \pm 0.2\%$	$15.2 \pm 0.2\%$	$17.5 \pm 0.2\%$
$\hat{FNR}_M$	$18.6 \pm 0.3\%$	$16.0 \pm 0.2\%$	$18.2 \pm 0.3\%$
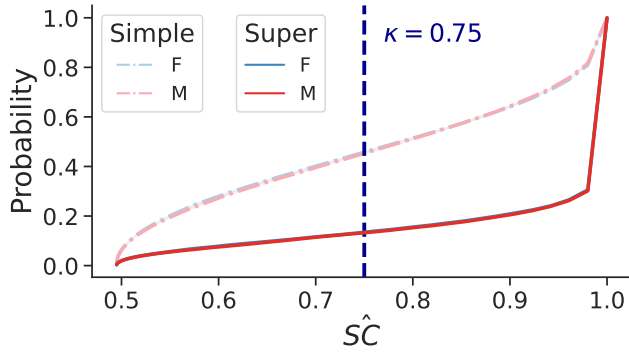
Abstention set metrics		
	Simple	**Super**
$\Delta\hat{AR}$	$0.2 \pm 0.0\%$	$0.2 \pm 0.2\%$
$\hat{AR}_F$	$48.1 \pm 0.3\%$	$13.2 \pm 0.1\%$
$\hat{AR}_M$	$47.9 \pm 0.3\%$	$13.0 \pm 0.3\%$

Figure 34. **Random forests** on `New Adult - CA - Public Coverage`, by `sex`

`Public Coverage`‑by `race`.

Logistic regression prediction set metrics			
	Baseline	**Simple**	**Super**
$\Delta\hat{PR}$	$0.1 \pm 0.1\%$	$0.1 \pm 0.1\%$	$0.1 \pm 0.1\%$
$\hat{PR}_{NW}$	$16.3 \pm 0.3\%$	$15.8 \pm 0.3\%$	$16.2 \pm 0.3\%$
$\hat{PR}_W$	$16.2 \pm 0.2\%$	$15.9 \pm 0.2\%$	$16.3 \pm 0.2\%$
$\Delta\hat{Err}$	$3.2 \pm 0.0\%$	$2.7 \pm 0.1\%$	$2.8 \pm 0.1\%$
$\hat{Err}_{NW}$	$33.4 \pm 0.3\%$	$32.6 \pm 0.3\%$	$32.8 \pm 0.3\%$
$\hat{Err}_W$	$30.2 \pm 0.3\%$	$29.9 \pm 0.2\%$	$30.0 \pm 0.2\%$
$\Delta\hat{FPR}$	$0.2 \pm 0.0\%$	$0.1 \pm 0.0\%$	$0.1 \pm 0.1\%$
$\hat{FPR}_{NW}$	$5.6 \pm 0.1\%$	$5.2 \pm 0.1\%$	$5.4 \pm 0.2\%$
$\hat{FPR}_W$	$5.4 \pm 0.1\%$	$5.1 \pm 0.1\%$	$5.3 \pm 0.1\%$
$\Delta\hat{FNR}$	$3.0 \pm 0.1\%$	$2.6 \pm 0.1\%$	$2.7 \pm 0.1\%$
$\hat{FNR}_{NW}$	$27.8 \pm 0.3\%$	$27.4 \pm 0.3\%$	$27.4 \pm 0.3\%$
$\hat{FNR}_W$	$24.8 \pm 0.2\%$	$24.8 \pm 0.2\%$	$24.7 \pm 0.2\%$
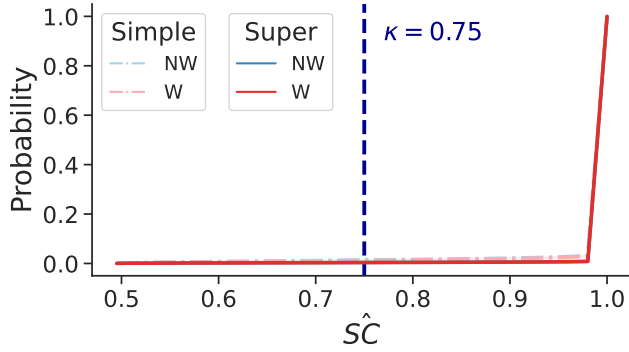
Abstention set metrics		
	Simple	**Super**
$\Delta\hat{AR}$	$0.1 \pm 0.0\%$	$0.1 \pm 0.0\%$
$\hat{AR}_{NW}$	$1.5 \pm 0.1\%$	$0.4 \pm 0.0\%$
$\hat{AR}_W$	$1.4 \pm 0.1\%$	$0.3 \pm 0.0\%$

Figure 35. **Logistic regression** on `New Adult - CA - Public Coverage`, by `race`

Decision tree prediction set metrics			
	Baseline	**Simple**	**Super**
$\Delta\hat{\mathbf{PR}}$	$1.5 \pm 0.0\%$	$2.2 \pm 0.2\%$	$1.4 \pm 0.1\%$
$\hat{\mathbf{PR}}_{\mathbf{NW}}$	$37.5 \pm 0.2\%$	$24.8 \pm 0.3\%$	$29.6 \pm 0.4\%$
$\hat{\mathbf{PR}}_{\mathbf{W}}$	$36.0 \pm 0.2\%$	$22.6 \pm 0.5\%$	$28.2 \pm 0.3\%$
$\Delta\hat{\mathbf{Err}}$	$2.2 \pm 0.0\%$	$3.0 \pm 0.1\%$	$2.7 \pm 0.1\%$
$\hat{\mathbf{Err}}_{\mathbf{NW}}$	$36.4 \pm 0.1\%$	$20.8 \pm 0.3\%$	$28.3 \pm 0.2\%$
$\hat{\mathbf{Err}}_{\mathbf{W}}$	$34.2 \pm 0.1\%$	$17.8 \pm 0.4\%$	$25.6 \pm 0.3\%$
$\Delta\hat{\mathbf{FPR}}$	$0.4 \pm 0.1\%$	$0.7 \pm 0.0\%$	$0.6 \pm 0.0\%$
$\hat{\mathbf{FPR}}_{\mathbf{NW}}$	$17.7 \pm 0.1\%$	$5.0 \pm 0.2\%$	$10.2 \pm 0.2\%$
$\hat{\mathbf{FPR}}_{\mathbf{W}}$	$17.3 \pm 0.2\%$	$4.3 \pm 0.2\%$	$9.6 \pm 0.2\%$
$\Delta\hat{\mathbf{FNR}}$	$1.8 \pm 0.0\%$	$2.3 \pm 0.1\%$	$2.1 \pm 0.1\%$
$\hat{\mathbf{FNR}}_{\mathbf{NW}}$	$18.7 \pm 0.2\%$	$15.8 \pm 0.2\%$	$18.1 \pm 0.3\%$
$\hat{\mathbf{FNR}}_{\mathbf{W}}$	$16.9 \pm 0.2\%$	$13.5 \pm 0.3\%$	$16.0 \pm 0.2\%$
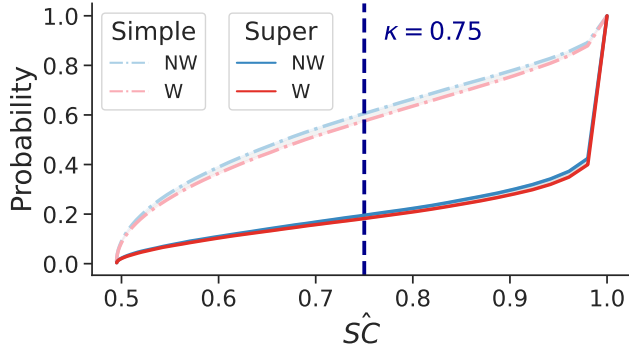
Abstention set metrics		
	Simple	**Super**
$\Delta\hat{\mathbf{AR}}$	$2.8 \pm 0.0\%$	$1.3 \pm 0.1\%$
$\hat{\mathbf{AR}}_{\mathbf{NW}}$	$62.3 \pm 0.3\%$	$19.2 \pm 0.3\%$
$\hat{\mathbf{AR}}_{\mathbf{W}}$	$59.5 \pm 0.3\%$	$17.9 \pm 0.4\%$

Figure 36. **Decision trees** on `New Adult - CA - Public Coverage,` by `race`

Random forest prediction set metrics			
	Baseline	**Simple**	**Super**
$\Delta\hat{\mathbf{PR}}$	$1.2 \pm 0.0\%$	$1.0 \pm 0.1\%$	$0.7 \pm 0.0\%$
$\hat{\mathbf{PR}}_{\mathbf{NW}}$	$33.7 \pm 0.3\%$	$22.4 \pm 0.4\%$	$27.3 \pm 0.4\%$
$\hat{\mathbf{PR}}_{\mathbf{W}}$	$32.5 \pm 0.3\%$	$21.4 \pm 0.3\%$	$26.6 \pm 0.4\%$
$\Delta\hat{\mathbf{Err}}$	$2.7 \pm 0.1\%$	$2.9 \pm 0.0\%$	$2.6 \pm 0.1\%$
$\hat{\mathbf{Err}}_{\mathbf{NW}}$	$34.0 \pm 0.2\%$	$21.4 \pm 0.3\%$	$27.9 \pm 0.3\%$
$\hat{\mathbf{Err}}_{\mathbf{W}}$	$31.3 \pm 0.1\%$	$18.5 \pm 0.3\%$	$25.3 \pm 0.2\%$
$\Delta\hat{\mathbf{FPR}}$	$0.5 \pm 0.0\%$	$0.4 \pm 0.0\%$	$0.5 \pm 0.0\%$
$\hat{\mathbf{FPR}}_{\mathbf{NW}}$	$14.6 \pm 0.2\%$	$4.4 \pm 0.2\%$	$8.9 \pm 0.2\%$
$\hat{\mathbf{FPR}}_{\mathbf{W}}$	$14.1 \pm 0.2\%$	$4.0 \pm 0.2\%$	$8.4 \pm 0.2\%$
$\Delta\hat{\mathbf{FNR}}$	$2.2 \pm 0.0\%$	$2.5 \pm 0.0\%$	$2.3 \pm 0.1\%$
$\hat{\mathbf{FNR}}_{\mathbf{NW}}$	$19.4 \pm 0.2\%$	$17.0 \pm 0.2\%$	$19.1 \pm 0.3\%$
$\hat{\mathbf{FNR}}_{\mathbf{W}}$	$17.2 \pm 0.2\%$	$14.5 \pm 0.2\%$	$16.8 \pm 0.2\%$
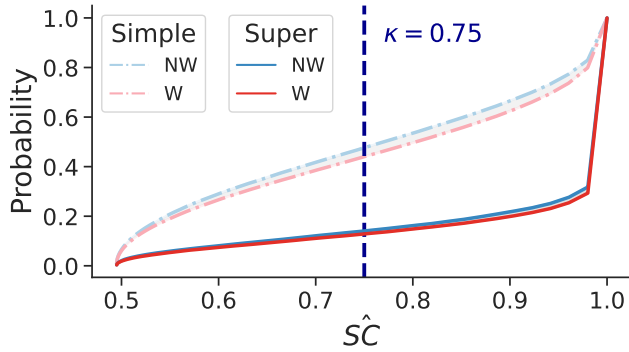
Abstention set metrics		
	Simple	**Super**
$\Delta\hat{\mathbf{AR}}$	$3.6 \pm 0.1\%$	$1.2 \pm 0.0\%$
$\hat{\mathbf{AR}}_{\mathbf{NW}}$	$50.0 \pm 0.2\%$	$13.8 \pm 0.2\%$
$\hat{\mathbf{AR}}_{\mathbf{W}}$	$46.4 \pm 0.3\%$	$12.6 \pm 0.2\%$

Figure 37. **Random forests** on `New Adult - CA - Public Coverage,` by `race`

E.4.6. `HMDA`

$\hat{SC}$ CDFs for two states (`NY`, `TX`) in `HMDA -2017`, using $\boldsymbol{g} = $ `ethnicity`, `race`, and `sex`, and associated error metrics on the prediction set. **Baseline** metrics computed with $B = 101$ models. For **simple**, $B = 101$ models; for **super**, $B = 101$ ensemble models, each composed of 21 underlying models for `NY`; 15 for `TX`. We repeat for 5 test/train splits. We also report abstention rate $\hat{AR}$.

`NY - 2017` - by `ethnicity`.

Logistic regression prediction set metrics			
	Baseline	**Simple**	**Super**
$\Delta\hat{\mathbf{PR}}$	$9.7 \pm 0.2\%$	$10.5 \pm 0.5\%$	$10.4 \pm 0.5\%$
$\hat{\mathbf{PR}}_{\mathrm{HL}}$	$73.5 \pm 0.3\%$	$73.5 \pm 0.6\%$	$73.1 \pm 0.6\%$
$\hat{\mathbf{PR}}_{\mathrm{NHL}}$	$83.2 \pm 0.1\%$	$84.0 \pm 0.1\%$	$83.5 \pm 0.1\%$
$\Delta\hat{\mathbf{Err}}$	$1.3 \pm 0.2\%$	$1.7 \pm 0.3\%$	$1.8 \pm 0.4\%$
$\hat{\mathbf{Err}}_{\mathrm{HL}}$	$18.7 \pm 0.3\%$	$18.4 \pm 0.4\%$	$18.9 \pm 0.5\%$
$\hat{\mathbf{Err}}_{\mathrm{NHL}}$	$17.4 \pm 0.1\%$	$16.7 \pm 0.1\%$	$17.1 \pm 0.1\%$
$\Delta\hat{\mathbf{FPR}}$	$0.8 \pm 0.1\%$	$1.0 \pm 0.1\%$	$0.9 \pm 0.1\%$
$\hat{\mathbf{FPR}}_{\mathrm{HL}}$	$10.7 \pm 0.2\%$	$10.2 \pm 0.2\%$	$10.5 \pm 0.2\%$
$\hat{\mathbf{FPR}}_{\mathrm{NHL}}$	$11.5 \pm 0.1\%$	$11.2 \pm 0.1\%$	$11.4 \pm 0.1\%$
$\Delta\hat{\mathbf{FNR}}$	$2.2 \pm 0.2\%$	$2.7 \pm 0.3\%$	$2.7 \pm 0.4\%$
$\hat{\mathbf{FNR}}_{\mathrm{HL}}$	$8.0 \pm 0.3\%$	$8.2 \pm 0.4\%$	$8.4 \pm 0.5\%$
$\hat{\mathbf{FNR}}_{\mathrm{NHL}}$	$5.8 \pm 0.1\%$	$5.5 \pm 0.1\%$	$5.7 \pm 0.1\%$
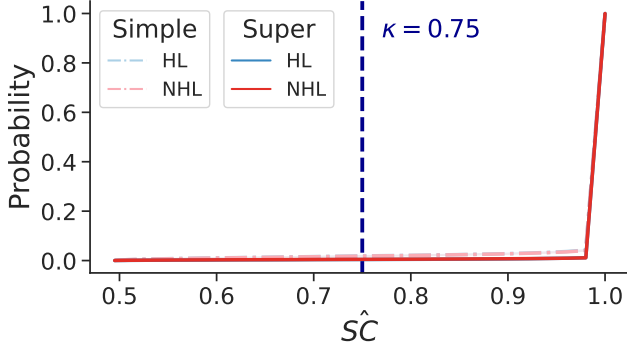
Abstention set metrics		
	Simple	**Super**
$\Delta\hat{\mathbf{AR}}$	$0.0 \pm 0.1\%$	$0.1 \pm 0.1\%$
$\hat{\mathbf{AR}}_{\mathrm{HL}}$	$1.9 \pm 0.2\%$	$0.4 \pm 0.1\%$
$\hat{\mathbf{AR}}_{\mathrm{NHL}}$	$1.9 \pm 0.1\%$	$0.5 \pm 0.0\%$

Figure 38. **Logistic regression** on `HMDA - 2017 - NY`, by `ethnicity`

Decision tree prediction set metrics			
	Baseline	**Simple**	**Super**
$\Delta\hat{\mathbf{PR}}$	$4.0 \pm 0.2\%$	$1.5 \pm 0.2\%$	$8.2 \pm 0.5\%$
$\hat{\mathbf{PR}}_{\mathrm{HL}}$	$73.1 \pm 0.3\%$	$94.0 \pm 0.3\%$	$74.1 \pm 0.6\%$
$\hat{\mathbf{PR}}_{\mathrm{NHL}}$	$77.1 \pm 0.1\%$	$95.5 \pm 0.1\%$	$82.3 \pm 0.1\%$
$\Delta\hat{\mathbf{Err}}$	$0.4 \pm 0.1\%$	$0.7 \pm 0.1\%$	$0.6 \pm 0.2\%$
$\hat{\mathbf{Err}}_{\mathrm{HL}}$	$20.6 \pm 0.2\%$	$2.7 \pm 0.2\%$	$12.2 \pm 0.4\%$
$\hat{\mathbf{Err}}_{\mathrm{NHL}}$	$20.2 \pm 0.1\%$	$3.4 \pm 0.1\%$	$11.6 \pm 0.2\%$
$\Delta\hat{\mathbf{FPR}}$	$1.6 \pm 0.1\%$	$1.1 \pm 0.0\%$	$1.1 \pm 0.1\%$
$\hat{\mathbf{FPR}}_{\mathrm{HL}}$	$11.5 \pm 0.2\%$	$1.4 \pm 0.1\%$	$5.1 \pm 0.2\%$
$\hat{\mathbf{FPR}}_{\mathrm{NHL}}$	$9.9 \pm 0.1\%$	$2.5 \pm 0.1\%$	$6.2 \pm 0.1\%$
$\Delta\hat{\mathbf{FNR}}$	$1.2 \pm 0.1\%$	$0.4 \pm 0.1\%$	$1.6 \pm 0.3\%$
$\hat{\mathbf{FNR}}_{\mathrm{HL}}$	$9.1 \pm 0.2\%$	$1.3 \pm 0.2\%$	$7.0 \pm 0.4\%$
$\hat{\mathbf{FNR}}_{\mathrm{NHL}}$	$10.3 \pm 0.1\%$	$0.9 \pm 0.1\%$	$5.4 \pm 0.1\%$
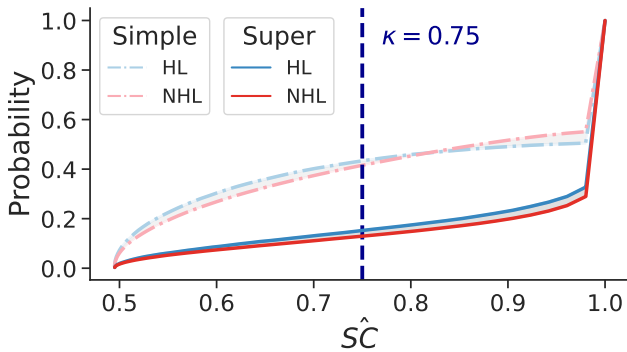
Abstention set metrics		
	Simple	**Super**
$\Delta\hat{\mathbf{AR}}$	$1.2 \pm 0.2\%$	$2.3 \pm 0.1\%$
$\hat{\mathbf{AR}}_{\mathrm{HL}}$	$43.3 \pm 0.4\%$	$15.0 \pm 0.3\%$
$\hat{\mathbf{AR}}_{\mathrm{NHL}}$	$42.1 \pm 0.2\%$	$12.7 \pm 0.2\%$

Figure 39. **Decision trees** on `HMDA - 2017 - NY`, by `ethnicity`

Random forest prediction set metrics			
	Baseline	**Simple**	**Super**
$\Delta\hat{\mathbf{PR}}$	$6.3 \pm 0.3\%$	$5.1 \pm 0.6\%$	$9.4 \pm 0.4\%$
$\hat{\mathbf{PR}}_{\mathbf{HL}}$	$71.9 \pm 0.4\%$	$85.8 \pm 0.7\%$	$71.6 \pm 0.5\%$
$\hat{\mathbf{PR}}_{\mathbf{NHL}}$	$78.2 \pm 0.1\%$	$90.9 \pm 0.1\%$	$81.0 \pm 0.1\%$
$\Delta\hat{\mathbf{Err}}$	$1.1 \pm 0.1\%$	$0.5 \pm 0.3\%$	$1.0 \pm 0.2\%$
$\hat{\mathbf{Err}}_{\mathbf{HL}}$	$19.1 \pm 0.2\%$	$5.8 \pm 0.4\%$	$13.7 \pm 0.4\%$
$\hat{\mathbf{Err}}_{\mathbf{NHL}}$	$18.0 \pm 0.1\%$	$6.3 \pm 0.1\%$	$12.7 \pm 0.2\%$
$\Delta\hat{\mathbf{FPR}}$	$0.7 \pm 0.1\%$	$1.5 \pm 0.1\%$	$1.2 \pm 0.1\%$
$\hat{\mathbf{FPR}}_{\mathbf{HL}}$	$10.1 \pm 0.2\%$	$2.6 \pm 0.1\%$	$5.7 \pm 0.2\%$
$\hat{\mathbf{FPR}}_{\mathbf{NHL}}$	$9.4 \pm 0.1\%$	$4.1 \pm 0.0\%$	$6.9 \pm 0.1\%$
$\Delta\hat{\mathbf{FNR}}$	$0.4 \pm 0.2\%$	$1.1 \pm 0.2\%$	$2.2 \pm 0.2\%$
$\hat{\mathbf{FNR}}_{\mathbf{HL}}$	$9.0 \pm 0.3\%$	$3.3 \pm 0.3\%$	$8.0 \pm 0.3\%$
$\hat{\mathbf{FNR}}_{\mathbf{NHL}}$	$8.6 \pm 0.1\%$	$2.2 \pm 0.1\%$	$5.8 \pm 0.1\%$
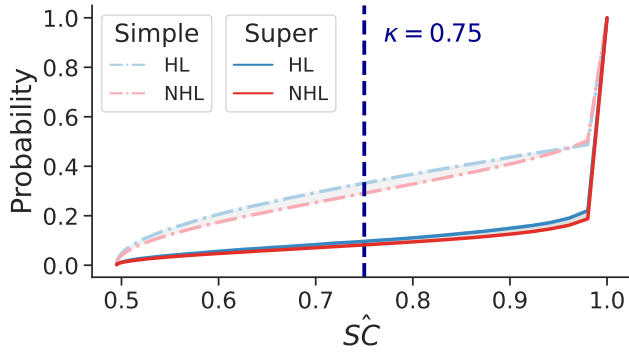
Abstention set metrics		
	Simple	**Super**
$\Delta\hat{\mathbf{AR}}$	$3.6 \pm 0.2\%$	$1.4 \pm 0.3\%$
$\hat{\mathbf{AR}}_{\mathbf{HL}}$	$33.9 \pm 0.3\%$	$9.5 \pm 0.4\%$
$\hat{\mathbf{AR}}_{\mathbf{NHL}}$	$30.3 \pm 0.1\%$	$8.1 \pm 0.1\%$

Figure 40. **Random forests** on `HMDA - 2017 - NY`, by `ethnicity`

NY - 2017 - by race.

Logistic regression prediction set metrics			
	Baseline	**Simple**	**Super**
$\Delta\hat{\mathbf{PR}}$	$11.5 \pm 0.2\%$	$11.6 \pm 0.3\%$	$11.5 \pm 0.3\%$
$\hat{\mathbf{PR}}_{\mathbf{NW}}$	$73.3 \pm 0.3\%$	$73.9 \pm 0.4\%$	$73.5 \pm 0.4\%$
$\hat{\mathbf{PR}}_{\mathbf{W}}$	$84.8 \pm 0.1\%$	$85.5 \pm 0.1\%$	$85.0 \pm 0.1\%$
$\Delta\hat{\mathbf{Err}}$	$2.8 \pm 0.1\%$	$2.9 \pm 0.1\%$	$3.0 \pm 0.0\%$
$\hat{\mathbf{Err}}_{\mathbf{NW}}$	$19.7 \pm 0.2\%$	$19.1 \pm 0.2\%$	$19.6 \pm 0.1\%$
$\hat{\mathbf{Err}}_{\mathbf{W}}$	$16.9 \pm 0.1\%$	$16.2 \pm 0.1\%$	$16.6 \pm 0.1\%$
$\Delta\hat{\mathbf{FPR}}$	$0.2 \pm 0.1\%$	$0.2 \pm 0.1\%$	$0.1 \pm 0.1\%$
$\hat{\mathbf{FPR}}_{\mathbf{NW}}$	$11.3 \pm 0.2\%$	$11.0 \pm 0.2\%$	$11.3 \pm 0.2\%$
$\hat{\mathbf{FPR}}_{\mathbf{W}}$	$11.5 \pm 0.1\%$	$11.2 \pm 0.1\%$	$11.4 \pm 0.1\%$
$\Delta\hat{\mathbf{FNR}}$	$3.0 \pm 0.1\%$	$3.0 \pm 0.1\%$	$3.0 \pm 0.1\%$
$\hat{\mathbf{FNR}}_{\mathbf{NW}}$	$8.4 \pm 0.2\%$	$8.1 \pm 0.2\%$	$8.3 \pm 0.2\%$
$\hat{\mathbf{FNR}}_{\mathbf{W}}$	$5.4 \pm 0.1\%$	$5.1 \pm 0.1\%$	$5.3 \pm 0.1\%$
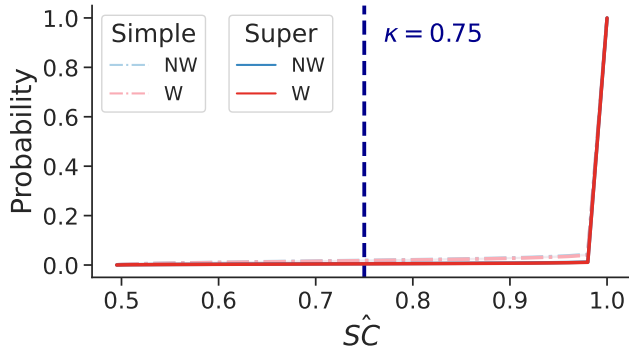
Abstention set metrics		
	Simple	**Super**
$\Delta\hat{\mathbf{AR}}$	$0.1 \pm 0.0\%$	$0.0 \pm 0.1\%$
$\hat{\mathbf{AR}}_{\mathbf{NW}}$	$2.0 \pm 0.1\%$	$0.5 \pm 0.1\%$
$\hat{\mathbf{AR}}_{\mathbf{W}}$	$1.9 \pm 0.1\%$	$0.5 \pm 0.0\%$

Figure 41. **Logistic regression** on `HMDA - 2017 - NY`, by `race`

Decision tree prediction set metrics			
	Baseline	**Simple**	**Super**
$\Delta\hat{\mathbf{PR}}$	$6.2 \pm 0.2\%$	$3.0 \pm 0.1\%$	$9.1 \pm 0.1\%$
$\hat{\mathbf{PR}}_{\mathbf{NW}}$	$71.8 \pm 0.3\%$	$93.0 \pm 0.2\%$	$74.4 \pm 0.3\%$
$\hat{\mathbf{PR}}_{\mathbf{W}}$	$78.0 \pm 0.1\%$	$96.0 \pm 0.1\%$	$83.5 \pm 0.2\%$
$\Delta\hat{\mathbf{Err}}$	$3.1 \pm 0.0\%$	$0.7 \pm 0.1\%$	$2.3 \pm 0.0\%$
$\hat{\mathbf{Err}}_{\mathbf{NW}}$	$22.7 \pm 0.1\%$	$3.9 \pm 0.2\%$	$13.5 \pm 0.2\%$
$\hat{\mathbf{Err}}_{\mathbf{W}}$	$19.6 \pm 0.1\%$	$3.2 \pm 0.1\%$	$11.2 \pm 0.2\%$
$\Delta\hat{\mathbf{FPR}}$	$2.5 \pm 0.1\%$	$0.2 \pm 0.0\%$	$0.6 \pm 0.1\%$
$\hat{\mathbf{FPR}}_{\mathbf{NW}}$	$12.0 \pm 0.2\%$	$2.5 \pm 0.1\%$	$6.6 \pm 0.2\%$
$\hat{\mathbf{FPR}}_{\mathbf{W}}$	$9.5 \pm 0.1\%$	$2.3 \pm 0.1\%$	$6.0 \pm 0.1\%$
$\Delta\hat{\mathbf{FNR}}$	$0.6 \pm 0.0\%$	$0.5 \pm 0.1\%$	$1.8 \pm 0.1\%$
$\hat{\mathbf{FNR}}_{\mathbf{NW}}$	$10.7 \pm 0.1\%$	$1.4 \pm 0.1\%$	$6.9 \pm 0.2\%$
$\hat{\mathbf{FNR}}_{\mathbf{W}}$	$10.1 \pm 0.1\%$	$0.9 \pm 0.0\%$	$5.1 \pm 0.1\%$
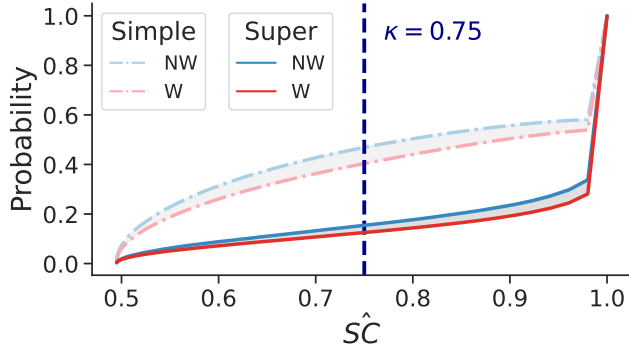
Abstention set metrics		
	Simple	**Super**
$\Delta\hat{\mathbf{AR}}$	$6.2 \pm 0.2\%$	$2.9 \pm 0.0\%$
$\hat{\mathbf{AR}}_{\mathbf{NW}}$	$47.1 \pm 0.4\%$	$15.2 \pm 0.2\%$
$\hat{\mathbf{AR}}_{\mathbf{W}}$	$40.9 \pm 0.2\%$	$12.3 \pm 0.2\%$

Figure 42. **Decision trees** on HMDA - 2017 - NY, by race

Random forest prediction set metrics			
	Baseline	**Simple**	**Super**
$\Delta\hat{\mathbf{PR}}$	$7.9 \pm 0.2\%$	$6.5 \pm 0.1\%$	$9.8 \pm 0.1\%$
$\hat{\mathbf{PR}}_{\mathbf{NW}}$	$71.5 \pm 0.3\%$	$85.3 \pm 0.2\%$	$72.4 \pm 0.3\%$
$\hat{\mathbf{PR}}_{\mathbf{W}}$	$79.4 \pm 0.1\%$	$91.8 \pm 0.1\%$	$82.2 \pm 0.2\%$
$\Delta\hat{\mathbf{Err}}$	$3.3 \pm 0.0\%$	$1.3 \pm 0.0\%$	$2.7 \pm 0.1\%$
$\hat{\mathbf{Err}}_{\mathbf{NW}}$	$20.7 \pm 0.1\%$	$7.3 \pm 0.1\%$	$14.9 \pm 0.1\%$
$\hat{\mathbf{Err}}_{\mathbf{W}}$	$17.4 \pm 0.1\%$	$6.0 \pm 0.1\%$	$12.2 \pm 0.2\%$
$\Delta\hat{\mathbf{FPR}}$	$1.8 \pm 0.1\%$	$0.0 \pm 0.0\%$	$0.5 \pm 0.0\%$
$\hat{\mathbf{FPR}}_{\mathbf{NW}}$	$10.9 \pm 0.2\%$	$4.0 \pm 0.1\%$	$7.2 \pm 0.1\%$
$\hat{\mathbf{FPR}}_{\mathbf{W}}$	$9.1 \pm 0.1\%$	$4.0 \pm 0.1\%$	$6.7 \pm 0.1\%$
$\Delta\hat{\mathbf{FNR}}$	$1.4 \pm 0.0\%$	$1.3 \pm 0.0\%$	$2.1 \pm 0.1\%$
$\hat{\mathbf{FNR}}_{\mathbf{NW}}$	$9.8 \pm 0.1\%$	$3.3 \pm 0.1\%$	$7.7 \pm 0.1\%$
$\hat{\mathbf{FNR}}_{\mathbf{W}}$	$8.4 \pm 0.1\%$	$2.0 \pm 0.1\%$	$5.6 \pm 0.2\%$
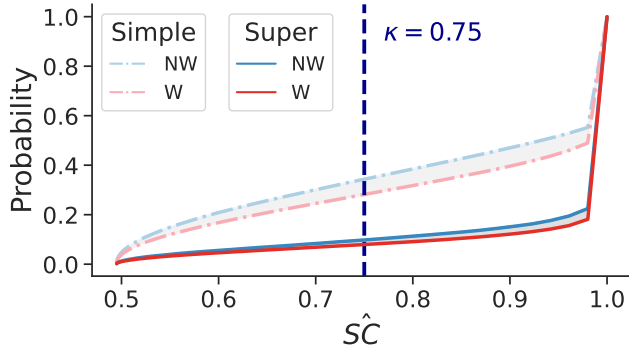
Abstention set metrics		
	Simple	**Super**
$\Delta\hat{\mathbf{AR}}$	$6.3 \pm 0.3\%$	$1.8 \pm 0.2\%$
$\hat{\mathbf{AR}}_{\mathbf{NW}}$	$35.6 \pm 0.4\%$	$9.6 \pm 0.3\%$
$\hat{\mathbf{AR}}_{\mathbf{W}}$	$29.3 \pm 0.1\%$	$7.8 \pm 0.1\%$

Figure 43. **Random forests** on HMDA - 2017 - NY, by race

NY – 2017 - by sex.

Logistic regression prediction set metrics			
	Baseline	**Simple**	**Super**
$\Delta\hat{\textbf{PR}}$	$1.5 \pm 0.0\%$	$1.5 \pm 0.2\%$	$1.5 \pm 0.2\%$
$\hat{\textbf{PR}}_F$	$81.5 \pm 0.1\%$	$82.2 \pm 0.3\%$	$81.7 \pm 0.3\%$
$\hat{\textbf{PR}}_M$	$83.0 \pm 0.1\%$	$83.7 \pm 0.1\%$	$83.2 \pm 0.1\%$
$\Delta\hat{\textbf{Err}}$	$0.5 \pm 0.1\%$	$0.6 \pm 0.1\%$	$0.6 \pm 0.1\%$
$\hat{\textbf{Err}}_F$	$17.8 \pm 0.2\%$	$17.2 \pm 0.2\%$	$17.6 \pm 0.2\%$
$\hat{\textbf{Err}}_M$	$17.3 \pm 0.1\%$	$16.6 \pm 0.1\%$	$17.0 \pm 0.1\%$
$\Delta\hat{\textbf{FPR}}$	$0.2 \pm 0.1\%$	$0.3 \pm 0.1\%$	$0.2 \pm 0.1\%$
$\hat{\textbf{FPR}}_F$	$11.6 \pm 0.2\%$	$11.3 \pm 0.2\%$	$11.5 \pm 0.2\%$
$\hat{\textbf{FPR}}_M$	$11.4 \pm 0.1\%$	$11.0 \pm 0.1\%$	$11.3 \pm 0.1\%$
$\Delta\hat{\textbf{FNR}}$	$0.3 \pm 0.0\%$	$0.3 \pm 0.0\%$	$0.3 \pm 0.1\%$
$\hat{\textbf{FNR}}_F$	$6.2 \pm 0.1\%$	$5.9 \pm 0.1\%$	$6.1 \pm 0.2\%$
$\hat{\textbf{FNR}}_M$	$5.9 \pm 0.1\%$	$5.6 \pm 0.1\%$	$5.8 \pm 0.1\%$
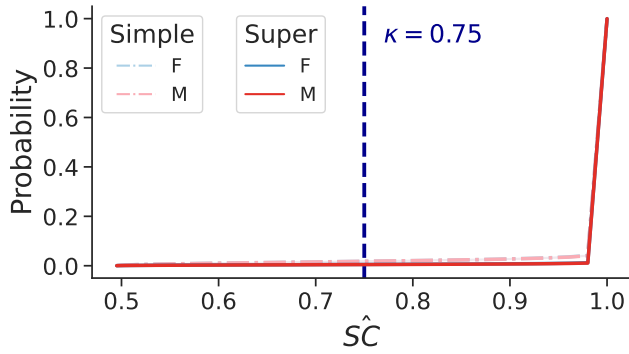
Abstention set metrics		
	Simple	**Super**
$\Delta\hat{\textbf{AR}}$	$0.1 \pm 0.0\%$	$0.0 \pm 0.0\%$
$\hat{\textbf{AR}}_F$	$2.0 \pm 0.1\%$	$0.5 \pm 0.0\%$
$\hat{\textbf{AR}}_M$	$1.9 \pm 0.1\%$	$0.5 \pm 0.0\%$

Figure 44. **Logistic regression** on `HMDA - 2017 - NY`, by `sex`

Decision tree prediction set metrics			
	Baseline	**Simple**	**Super**
$\Delta\hat{\textbf{PR}}$	$1.5 \pm 0.0\%$	$0.9 \pm 0.0\%$	$1.9 \pm 0.1\%$
$\hat{\textbf{PR}}_F$	$75.8 \pm 0.1\%$	$94.8 \pm 0.1\%$	$80.5 \pm 0.2\%$
$\hat{\textbf{PR}}_M$	$77.3 \pm 0.1\%$	$95.7 \pm 0.1\%$	$82.4 \pm 0.1\%$
$\Delta\hat{\textbf{Err}}$	$0.4 \pm 0.1\%$	$0.1 \pm 0.0\%$	$0.3 \pm 0.0\%$
$\hat{\textbf{Err}}_F$	$20.5 \pm 0.2\%$	$3.3 \pm 0.1\%$	$11.8 \pm 0.2\%$
$\hat{\textbf{Err}}_M$	$20.1 \pm 0.1\%$	$3.4 \pm 0.1\%$	$11.5 \pm 0.2\%$
$\Delta\hat{\textbf{FPR}}$	$0.2 \pm 0.1\%$	$0.2 \pm 0.0\%$	$0.1 \pm 0.0\%$
$\hat{\textbf{FPR}}_F$	$10.1 \pm 0.2\%$	$2.2 \pm 0.1\%$	$6.1 \pm 0.1\%$
$\hat{\textbf{FPR}}_M$	$9.9 \pm 0.1\%$	$2.4 \pm 0.1\%$	$6.2 \pm 0.1\%$
$\Delta\hat{\textbf{FNR}}$	$0.3 \pm 0.0\%$	$0.2 \pm 0.0\%$	$0.4 \pm 0.1\%$
$\hat{\textbf{FNR}}_F$	$10.4 \pm 0.1\%$	$1.1 \pm 0.1\%$	$5.7 \pm 0.2\%$
$\hat{\textbf{FNR}}_M$	$10.1 \pm 0.1\%$	$0.9 \pm 0.1\%$	$5.3 \pm 0.1\%$
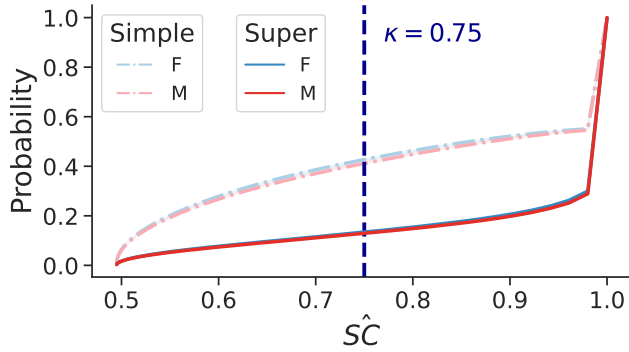
Abstention set metrics		
	Simple	**Super**
$\Delta\hat{\textbf{AR}}$	$1.4 \pm 0.0\%$	$0.5 \pm 0.0\%$
$\hat{\textbf{AR}}_F$	$43.1 \pm 0.2\%$	$13.2 \pm 0.2\%$
$\hat{\textbf{AR}}_M$	$41.7 \pm 0.2\%$	$12.7 \pm 0.2\%$

Figure 45. **Decision trees** on `HMDA - 2017 - NY`, by `sex`

Random forest prediction set metrics			
	Baseline	**Simple**	**Super**
$\Delta\hat{\mathbf{PR}}$	$1.8 \pm 0.0\%$	$1.7 \pm 0.1\%$	$2.0 \pm 0.1\%$
$\hat{\mathbf{PR}}_{\mathbf{F}}$	$76.6 \pm 0.1\%$	$89.4 \pm 0.2\%$	$79.0 \pm 0.2\%$
$\hat{\mathbf{PR}}_{\mathbf{M}}$	$78.4 \pm 0.1\%$	$91.1 \pm 0.1\%$	$81.0 \pm 0.1\%$
$\Delta\hat{\mathbf{Err}}$	$0.4 \pm 0.1\%$	$0.1 \pm 0.1\%$	$0.4 \pm 0.0\%$
$\hat{\mathbf{Err}}_{\mathbf{F}}$	$18.4 \pm 0.2\%$	$6.3 \pm 0.2\%$	$13.0 \pm 0.2\%$
$\hat{\mathbf{Err}}_{\mathbf{M}}$	$18.0 \pm 0.1\%$	$6.2 \pm 0.1\%$	$12.6 \pm 0.2\%$
$\Delta\hat{\mathbf{FPR}}$	$0.0 \pm 0.1\%$	$0.3 \pm 0.0\%$	$0.1 \pm 0.1\%$
$\hat{\mathbf{FPR}}_{\mathbf{F}}$	$9.4 \pm 0.2\%$	$3.8 \pm 0.1\%$	$6.7 \pm 0.2\%$
$\hat{\mathbf{FPR}}_{\mathbf{M}}$	$9.4 \pm 0.1\%$	$4.1 \pm 0.1\%$	$6.8 \pm 0.1\%$
$\Delta\hat{\mathbf{FNR}}$	$0.4 \pm 0.0\%$	$0.4 \pm 0.1\%$	$0.5 \pm 0.1\%$
$\hat{\mathbf{FNR}}_{\mathbf{F}}$	$8.9 \pm 0.1\%$	$2.5 \pm 0.1\%$	$6.3 \pm 0.2\%$
$\hat{\mathbf{FNR}}_{\mathbf{M}}$	$8.5 \pm 0.1\%$	$2.1 \pm 0.0\%$	$5.8 \pm 0.1\%$
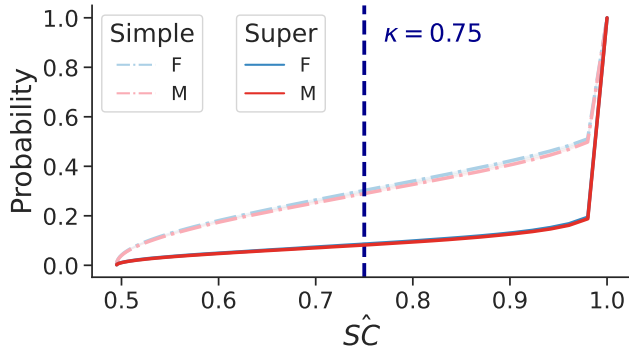
Abstention set metrics		
	Simple	**Super**
$\Delta\hat{\mathbf{AR}}$	$1.2 \pm 0.0\%$	$0.3 \pm 0.0\%$
$\hat{\mathbf{AR}}_{\mathbf{F}}$	$31.4 \pm 0.1\%$	$8.4 \pm 0.1\%$
$\hat{\mathbf{AR}}_{\mathbf{M}}$	$30.2 \pm 0.1\%$	$8.1 \pm 0.1\%$

Figure 46. **Random forests** on `HMDA - 2017 - NY`, by `sex`

`TX - 2017` - by `ethnicity`.

Logistic regression prediction set metrics			
	Baseline	**Simple**	**Super**
$\Delta\hat{\mathbf{PR}}$	$13.3 \pm 0.1\%$	$13.1 \pm 0.0\%$	$13.0 \pm 0.0\%$
$\hat{\mathbf{PR}}_{\mathbf{HL}}$	$64.7 \pm 0.2\%$	$65.1 \pm 0.1\%$	$65.0 \pm 0.1\%$
$\hat{\mathbf{PR}}_{\mathbf{NHL}}$	$78.0 \pm 0.1\%$	$78.2 \pm 0.1\%$	$78.0 \pm 0.1\%$
$\Delta\hat{\mathbf{Err}}$	$3.4 \pm 0.1\%$	$3.3 \pm 0.1\%$	$3.3 \pm 0.2\%$
$\hat{\mathbf{Err}}_{\mathbf{HL}}$	$17.5 \pm 0.2\%$	$17.1 \pm 0.2\%$	$17.3 \pm 0.2\%$
$\hat{\mathbf{Err}}_{\mathbf{NHL}}$	$14.1 \pm 0.1\%$	$13.8 \pm 0.1\%$	$14.0 \pm 0.0\%$
$\Delta\hat{\mathbf{FPR}}$	$0.6 \pm 0.0\%$	$0.6 \pm 0.0\%$	$0.6 \pm 0.0\%$
$\hat{\mathbf{FPR}}_{\mathbf{HL}}$	$5.6 \pm 0.1\%$	$5.4 \pm 0.0\%$	$5.6 \pm 0.0\%$
$\hat{\mathbf{FPR}}_{\mathbf{NHL}}$	$6.2 \pm 0.1\%$	$6.0 \pm 0.0\%$	$6.2 \pm 0.0\%$
$\Delta\hat{\mathbf{FNR}}$	$4.0 \pm 0.1\%$	$3.9 \pm 0.1\%$	$3.9 \pm 0.2\%$
$\hat{\mathbf{FNR}}_{\mathbf{HL}}$	$11.9 \pm 0.2\%$	$11.6 \pm 0.2\%$	$11.7 \pm 0.2\%$
$\hat{\mathbf{FNR}}_{\mathbf{NHL}}$	$7.9 \pm 0.1\%$	$7.7 \pm 0.1\%$	$7.8 \pm 0.0\%$
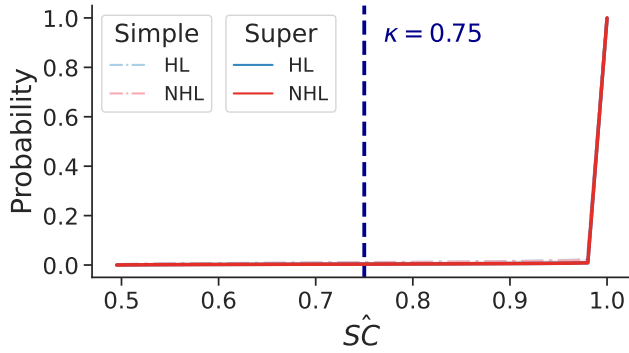
Abstention set metrics		
	Simple	**Super**
$\Delta\hat{\mathbf{AR}}$	$0.1 \pm 0.0\%$	$0.0 \pm 0.0\%$
$\hat{\mathbf{AR}}_{\mathbf{HL}}$	$1.1 \pm 0.0\%$	$0.4 \pm 0.0\%$
$\hat{\mathbf{AR}}_{\mathbf{NHL}}$	$1.0 \pm 0.0\%$	$0.4 \pm 0.0\%$

Figure 47. **Logistic regression** on `HMDA - 2017 - TX`, by `ethnicity`

Decision tree prediction set metrics			
	Baseline	**Simple**	**Super**
$\Delta\hat{\text{PR}}$	$7.4 \pm 0.2\%$	$4.1 \pm 0.1\%$	$8.5 \pm 0.0\%$
$\hat{\text{PR}}_{\text{HL}}$	$72.0 \pm 0.2\%$	$90.2 \pm 0.1\%$	$76.7 \pm 0.1\%$
$\hat{\text{PR}}_{\text{NHL}}$	$79.4 \pm 0.0\%$	$94.3 \pm 0.0\%$	$85.2 \pm 0.1\%$
$\Delta\hat{\text{Err}}$	$3.7 \pm 0.0\%$	$0.7 \pm 0.1\%$	$2.0 \pm 0.1\%$
$\hat{\text{Err}}_{\text{HL}}$	$19.2 \pm 0.1\%$	$3.0 \pm 0.1\%$	$8.4 \pm 0.1\%$
$\hat{\text{Err}}_{\text{NHL}}$	$15.5 \pm 0.1\%$	$2.3 \pm 0.0\%$	$6.4 \pm 0.0\%$
$\Delta\hat{\text{FPR}}$	$2.4 \pm 0.1\%$	$0.1 \pm 0.0\%$	$0.4 \pm 0.0\%$
$\hat{\text{FPR}}_{\text{HL}}$	$10.1 \pm 0.1\%$	$1.2 \pm 0.0\%$	$3.3 \pm 0.1\%$
$\hat{\text{FPR}}_{\text{NHL}}$	$7.7 \pm 0.0\%$	$1.1 \pm 0.0\%$	$2.9 \pm 0.1\%$
$\Delta\hat{\text{FNR}}$	$1.2 \pm 0.1\%$	$0.6 \pm 0.0\%$	$1.5 \pm 0.0\%$
$\hat{\text{FNR}}_{\text{HL}}$	$9.1 \pm 0.1\%$	$1.8 \pm 0.0\%$	$5.0 \pm 0.0\%$
$\hat{\text{FNR}}_{\text{NHL}}$	$7.9 \pm 0.0\%$	$1.2 \pm 0.0\%$	$3.5 \pm 0.0\%$
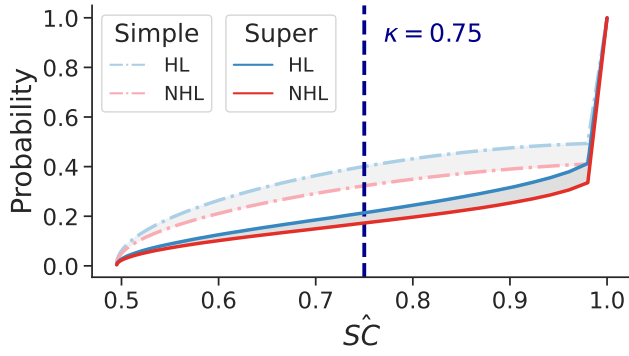
Abstention set metrics		
	Simple	**Super**
$\Delta\hat{\text{AR}}$	$7.3 \pm 0.0\%$	$4.1 \pm 0.1\%$
$\hat{\text{AR}}_{\text{HL}}$	$40.3 \pm 0.1\%$	$21.1 \pm 0.0\%$
$\hat{\text{AR}}_{\text{NHL}}$	$33.0 \pm 0.1\%$	$17.0 \pm 0.1\%$

Figure 48. **Decision trees** on `HMDA - 2017 - TX`, by `ethnicity`

Random forest prediction set metrics			
	Baseline	**Simple**	**Super**
$\Delta\hat{\text{PR}}$	$8.7 \pm 0.1\%$	$6.9 \pm 0.1\%$	$9.6 \pm 0.1\%$
$\hat{\text{PR}}_{\text{HL}}$	$70.6 \pm 0.2\%$	$83.0 \pm 0.2\%$	$72.7 \pm 0.2\%$
$\hat{\text{PR}}_{\text{NHL}}$	$79.3 \pm 0.1\%$	$89.9 \pm 0.1\%$	$82.3 \pm 0.1\%$
$\Delta\hat{\text{Err}}$	$3.3 \pm 0.0\%$	$1.2 \pm 0.1\%$	$2.2 \pm 0.1\%$
$\hat{\text{Err}}_{\text{HL}}$	$17.3 \pm 0.1\%$	$5.3 \pm 0.1\%$	$10.5 \pm 0.1\%$
$\hat{\text{Err}}_{\text{NHL}}$	$14.0 \pm 0.1\%$	$4.1 \pm 0.0\%$	$8.3 \pm 0.0\%$
$\Delta\hat{\text{FPR}}$	$1.7 \pm 0.1\%$	$0.1 \pm 0.0\%$	$0.4 \pm 0.1\%$
$\hat{\text{FPR}}_{\text{HL}}$	$8.5 \pm 0.1\%$	$2.0 \pm 0.0\%$	$4.1 \pm 0.0\%$
$\hat{\text{FPR}}_{\text{NHL}}$	$6.8 \pm 0.0\%$	$1.9 \pm 0.0\%$	$3.7 \pm 0.1\%$
$\Delta\hat{\text{FNR}}$	$1.7 \pm 0.1\%$	$1.1 \pm 0.1\%$	$1.8 \pm 0.1\%$
$\hat{\text{FNR}}_{\text{HL}}$	$8.9 \pm 0.1\%$	$3.3 \pm 0.1\%$	$6.4 \pm 0.1\%$
$\hat{\text{FNR}}_{\text{NHL}}$	$7.2 \pm 0.0\%$	$2.2 \pm 0.0\%$	$4.6 \pm 0.0\%$
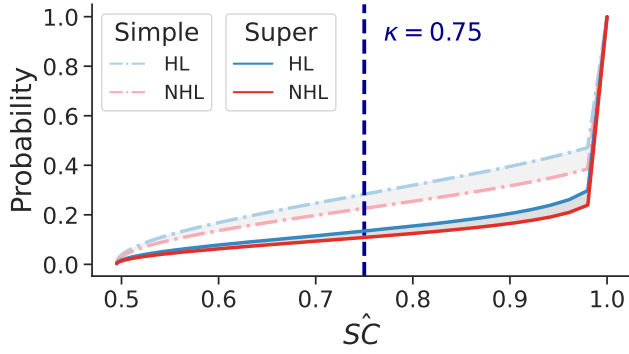
Abstention set metrics		
	Simple	**Super**
$\Delta\hat{\text{AR}}$	$6.2 \pm 0.1\%$	$2.5 \pm 0.1\%$
$\hat{\text{AR}}_{\text{HL}}$	$31.9 \pm 0.0\%$	$13.2 \pm 0.1\%$
$\hat{\text{AR}}_{\text{NHL}}$	$25.7 \pm 0.1\%$	$10.7 \pm 0.0\%$

Figure 49. **Random forests** on `HMDA - 2017 - TX`, by `ethnicity`

`TX – 2017 - by race.`

Logistic regression prediction set metrics			
	Baseline	**Simple**	**Super**
$\Delta\hat{\mathbf{PR}}$	$2.6 \pm 0.1\%$	$2.7 \pm 0.1\%$	$2.7 \pm 0.1\%$
$\hat{\mathbf{PR}}_{\mathbf{NW}}$	$72.6 \pm 0.2\%$	$72.8 \pm 0.0\%$	$72.6 \pm 0.0\%$
$\hat{\mathbf{PR}}_{\mathbf{W}}$	$75.2 \pm 0.1\%$	$75.5 \pm 0.1\%$	$75.3 \pm 0.1\%$
$\Delta\hat{\mathbf{Err}}$	$0.0 \pm 0.1\%$	$0.2 \pm 0.0\%$	$0.1 \pm 0.0\%$
$\hat{\mathbf{Err}}_{\mathbf{NW}}$	$14.9 \pm 0.2\%$	$14.4 \pm 0.1\%$	$14.7 \pm 0.1\%$
$\hat{\mathbf{Err}}_{\mathbf{W}}$	$14.9 \pm 0.1\%$	$14.6 \pm 0.1\%$	$14.8 \pm 0.1\%$
$\Delta\hat{\mathbf{FPR}}$	$1.1 \pm 0.1\%$	$0.9 \pm 0.1\%$	$1.0 \pm 0.1\%$
$\hat{\mathbf{FPR}}_{\mathbf{NW}}$	$7.0 \pm 0.2\%$	$6.6 \pm 0.1\%$	$6.8 \pm 0.1\%$
$\hat{\mathbf{FPR}}_{\mathbf{W}}$	$5.9 \pm 0.1\%$	$5.7 \pm 0.0\%$	$5.8 \pm 0.0\%$
$\Delta\hat{\mathbf{FNR}}$	$1.2 \pm 0.0\%$	$1.1 \pm 0.1\%$	$1.1 \pm 0.0\%$
$\hat{\mathbf{FNR}}_{\mathbf{NW}}$	$7.9 \pm 0.1\%$	$7.8 \pm 0.0\%$	$7.9 \pm 0.1\%$
$\hat{\mathbf{FNR}}_{\mathbf{W}}$	$9.1 \pm 0.1\%$	$8.9 \pm 0.1\%$	$9.0 \pm 0.1\%$
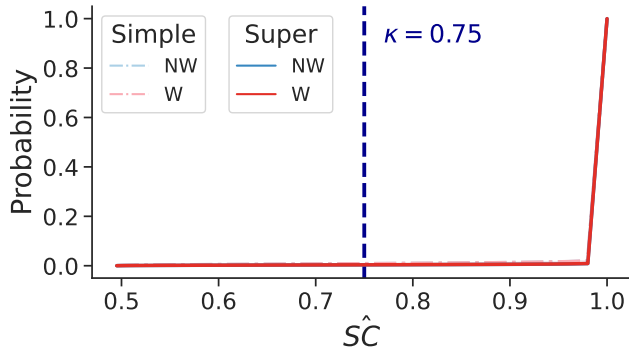
Abstention set metrics		
	Simple	**Super**
$\Delta\hat{\mathbf{AR}}$	$0.0 \pm 0.0\%$	$0.0 \pm 0.0\%$
$\hat{\mathbf{AR}}_{\mathbf{NW}}$	$1.0 \pm 0.0\%$	$0.4 \pm 0.0\%$
$\hat{\mathbf{AR}}_{\mathbf{W}}$	$1.0 \pm 0.0\%$	$0.4 \pm 0.0\%$

Figure 50. **Logistic regression** on `HMDA – 2017 – TX`, by `race`

Decision tree prediction set metrics			
	Baseline	**Simple**	**Super**
$\Delta\hat{\mathbf{PR}}$	$2.5 \pm 0.0\%$	$1.8 \pm 0.0\%$	$4.1 \pm 0.2\%$
$\hat{\mathbf{PR}}_{\mathbf{NW}}$	$75.6 \pm 0.1\%$	$91.9 \pm 0.1\%$	$79.8 \pm 0.3\%$
$\hat{\mathbf{PR}}_{\mathbf{W}}$	$78.1 \pm 0.1\%$	$93.7 \pm 0.1\%$	$83.9 \pm 0.1\%$
$\Delta\hat{\mathbf{Err}}$	$0.0 \pm 0.1\%$	$0.0 \pm 0.0\%$	$0.1 \pm 0.0\%$
$\hat{\mathbf{Err}}_{\mathbf{NW}}$	$16.4 \pm 0.1\%$	$2.4 \pm 0.0\%$	$6.8 \pm 0.1\%$
$\hat{\mathbf{Err}}_{\mathbf{W}}$	$16.4 \pm 0.0\%$	$2.4 \pm 0.0\%$	$6.9 \pm 0.1\%$
$\Delta\hat{\mathbf{FPR}}$	$1.2 \pm 0.0\%$	$0.0 \pm 0.0\%$	$0.1 \pm 0.0\%$
$\hat{\mathbf{FPR}}_{\mathbf{NW}}$	$9.2 \pm 0.1\%$	$1.1 \pm 0.0\%$	$2.9 \pm 0.1\%$
$\hat{\mathbf{FPR}}_{\mathbf{W}}$	$8.0 \pm 0.1\%$	$1.1 \pm 0.0\%$	$3.0 \pm 0.1\%$
$\Delta\hat{\mathbf{FNR}}$	$1.2 \pm 0.0\%$	$0.1 \pm 0.0\%$	$0.1 \pm 0.1\%$
$\hat{\mathbf{FNR}}_{\mathbf{NW}}$	$7.2 \pm 0.1\%$	$1.4 \pm 0.0\%$	$3.9 \pm 0.1\%$
$\hat{\mathbf{FNR}}_{\mathbf{W}}$	$8.4 \pm 0.1\%$	$1.3 \pm 0.0\%$	$3.8 \pm 0.0\%$
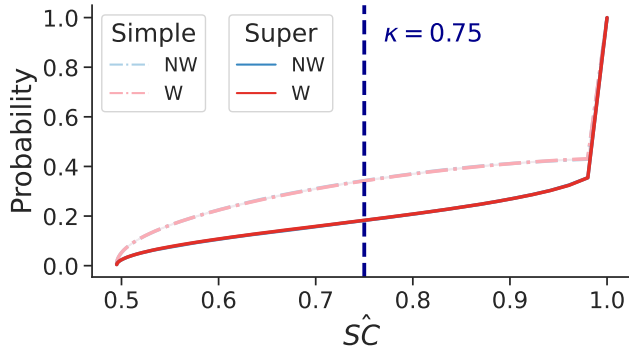
Abstention set metrics		
	Simple	**Super**
$\Delta\hat{\mathbf{AR}}$	$0.1 \pm 0.1\%$	$0.0 \pm 0.3\%$
$\hat{\mathbf{AR}}_{\mathbf{NW}}$	$34.7 \pm 0.2\%$	$18.0 \pm 0.3\%$
$\hat{\mathbf{AR}}_{\mathbf{W}}$	$34.8 \pm 0.1\%$	$18.0 \pm 0.0\%$

Figure 51. **Decision trees** on `HMDA – 2017 – TX`, by `race`
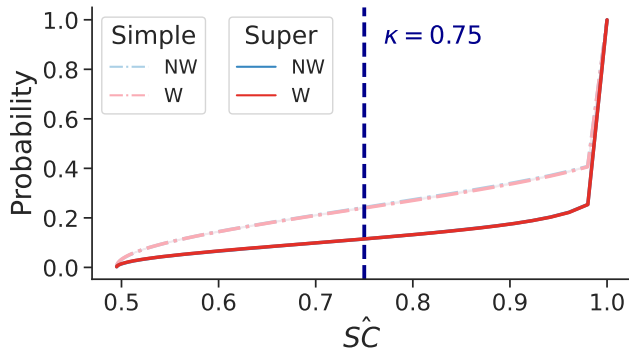
Abstention set metrics		
	Simple	**Super**
$\Delta\hat{\mathbf{AR}}$	$0.1 \pm 0.1\%$	$0.0 \pm 0.0\%$
$\hat{\mathbf{AR}}_{NW}$	$27.3 \pm 0.1\%$	$11.3 \pm 0.1\%$
$\hat{\mathbf{AR}}_{W}$	$27.2 \pm 0.0\%$	$11.3 \pm 0.1\%$

Random forest prediction set metrics			
	Baseline	**Simple**	**Super**
$\Delta\hat{\mathbf{PR}}$	$3.4 \pm 0.1\%$	$3.3 \pm 0.0\%$	$4.5 \pm 0.0\%$
$\hat{\mathbf{PR}}_{NW}$	$74.4 \pm 0.2\%$	$85.6 \pm 0.1\%$	$76.3 \pm 0.1\%$
$\hat{\mathbf{PR}}_{W}$	$77.8 \pm 0.1\%$	$88.9 \pm 0.1\%$	$80.8 \pm 0.1\%$
$\Delta\hat{\mathbf{Err}}$	$0.0 \pm 0.0\%$	$0.1 \pm 0.0\%$	$0.1 \pm 0.0\%$
$\hat{\mathbf{Err}}_{NW}$	$14.8 \pm 0.1\%$	$4.3 \pm 0.0\%$	$8.7 \pm 0.1\%$
$\hat{\mathbf{Err}}_{W}$	$14.8 \pm 0.1\%$	$4.4 \pm 0.0\%$	$8.8 \pm 0.1\%$
$\Delta\hat{\mathbf{FPR}}$	$0.7 \pm 0.0\%$	$0.1 \pm 0.1\%$	$0.1 \pm 0.1\%$
$\hat{\mathbf{FPR}}_{NW}$	$7.8 \pm 0.1\%$	$1.8 \pm 0.1\%$	$3.7 \pm 0.1\%$
$\hat{\mathbf{FPR}}_{W}$	$7.1 \pm 0.1\%$	$1.9 \pm 0.0\%$	$3.8 \pm 0.0\%$
$\Delta\hat{\mathbf{FNR}}$	$0.8 \pm 0.0\%$	$0.2 \pm 0.1\%$	$0.1 \pm 0.0\%$
$\hat{\mathbf{FNR}}_{NW}$	$6.9 \pm 0.1\%$	$2.6 \pm 0.1\%$	$5.1 \pm 0.0\%$
$\hat{\mathbf{FNR}}_{W}$	$7.7 \pm 0.1\%$	$2.4 \pm 0.0\%$	$5.0 \pm 0.0\%$

Figure 52. **Random forests** on `HMDA - 2017 - TX`, by `race`

TX - 2017 - by sex.
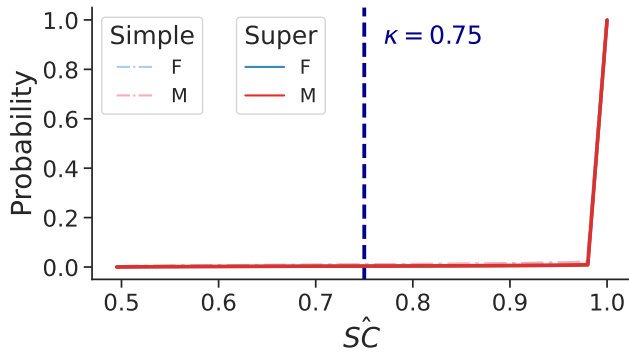
Abstention set metrics		
	Simple	**Super**
$\Delta\hat{\mathbf{AR}}$	$0.0 \pm 0.0\%$	$0.0 \pm 0.0\%$
$\hat{\mathbf{AR}}_{F}$	$1.0 \pm 0.0\%$	$0.4 \pm 0.0\%$
$\hat{\mathbf{AR}}_{M}$	$1.0 \pm 0.0\%$	$0.4 \pm 0.0\%$

Logistic regression prediction set metrics			
	Baseline	**Simple**	**Super**
$\Delta\hat{\mathbf{PR}}$	$5.7 \pm 0.1\%$	$5.6 \pm 0.0\%$	$5.6 \pm 0.0\%$
$\hat{\mathbf{PR}}_{F}$	$70.8 \pm 0.2\%$	$71.1 \pm 0.2\%$	$70.9 \pm 0.2\%$
$\hat{\mathbf{PR}}_{M}$	$76.5 \pm 0.1\%$	$76.7 \pm 0.2\%$	$76.5 \pm 0.2\%$
$\Delta\hat{\mathbf{Err}}$	$1.1 \pm 0.2\%$	$1.0 \pm 0.1\%$	$1.0 \pm 0.1\%$
$\hat{\mathbf{Err}}_{F}$	$15.7 \pm 0.2\%$	$15.3 \pm 0.0\%$	$15.5 \pm 0.0\%$
$\hat{\mathbf{Err}}_{M}$	$14.6 \pm 0.0\%$	$14.3 \pm 0.1\%$	$14.5 \pm 0.1\%$
$\Delta\hat{\mathbf{FPR}}$	$0.4 \pm 0.0\%$	$0.4 \pm 0.1\%$	$0.4 \pm 0.0\%$
$\hat{\mathbf{FPR}}_{F}$	$5.8 \pm 0.1\%$	$5.6 \pm 0.1\%$	$5.7 \pm 0.0\%$
$\hat{\mathbf{FPR}}_{M}$	$6.2 \pm 0.1\%$	$6.0 \pm 0.0\%$	$6.1 \pm 0.0\%$
$\Delta\hat{\mathbf{FNR}}$	$1.4 \pm 0.1\%$	$1.3 \pm 0.1\%$	$1.3 \pm 0.1\%$
$\hat{\mathbf{FNR}}_{F}$	$9.8 \pm 0.2\%$	$9.6 \pm 0.0\%$	$9.7 \pm 0.0\%$
$\hat{\mathbf{FNR}}_{M}$	$8.4 \pm 0.1\%$	$8.3 \pm 0.1\%$	$8.4 \pm 0.1\%$

Figure 53. **Logistic regression** on `HMDA - 2017 - TX`, by `sex`

Decision tree prediction set metrics			
	Baseline	**Simple**	**Super**
$\triangle\hat{\mathbf{PR}}$	$3.5 \pm 0.0\%$	$2.2 \pm 0.1\%$	$4.8 \pm 0.0\%$
$\hat{\mathbf{PR}}_F$	$75.2 \pm 0.1\%$	$91.8 \pm 0.0\%$	$79.8 \pm 0.1\%$
$\hat{\mathbf{PR}}_M$	$78.7 \pm 0.1\%$	$94.0 \pm 0.1\%$	$84.6 \pm 0.1\%$
$\triangle\hat{\mathbf{Err}}$	$1.1 \pm 0.1\%$	$0.2 \pm 0.0\%$	$0.5 \pm 0.0\%$
$\hat{\mathbf{Err}}_F$	$17.2 \pm 0.1\%$	$2.6 \pm 0.0\%$	$7.2 \pm 0.1\%$
$\hat{\mathbf{Err}}_M$	$16.1 \pm 0.0\%$	$2.4 \pm 0.0\%$	$6.7 \pm 0.1\%$
$\triangle\hat{\mathbf{FPR}}$	$0.8 \pm 0.1\%$	$0.2 \pm 0.0\%$	$0.3 \pm 0.1\%$
$\hat{\mathbf{FPR}}_F$	$8.8 \pm 0.1\%$	$1.0 \pm 0.0\%$	$2.8 \pm 0.0\%$
$\hat{\mathbf{FPR}}_M$	$8.0 \pm 0.0\%$	$1.2 \pm 0.0\%$	$3.1 \pm 0.1\%$
$\triangle\hat{\mathbf{FNR}}$	$0.3 \pm 0.1\%$	$0.4 \pm 0.0\%$	$0.9 \pm 0.1\%$
$\hat{\mathbf{FNR}}_F$	$8.4 \pm 0.1\%$	$1.6 \pm 0.0\%$	$4.5 \pm 0.1\%$
$\hat{\mathbf{FNR}}_M$	$8.1 \pm 0.0\%$	$1.2 \pm 0.0\%$	$3.6 \pm 0.0\%$
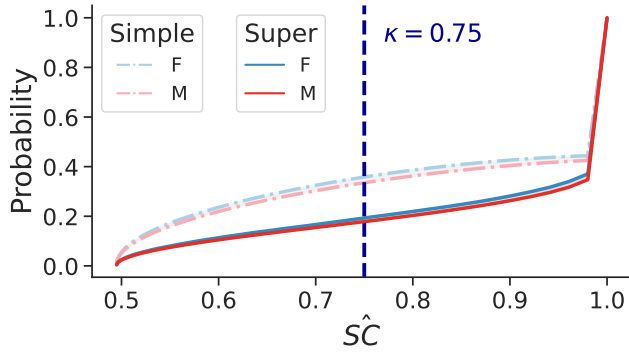
Abstention set metrics		
	Simple	**Super**
$\triangle\hat{\mathbf{AR}}$	$2.0 \pm 0.1\%$	$1.4 \pm 0.0\%$
$\hat{\mathbf{AR}}_F$	$36.2 \pm 0.2\%$	$19.0 \pm 0.1\%$
$\hat{\mathbf{AR}}_M$	$34.2 \pm 0.1\%$	$17.6 \pm 0.1\%$

Figure 54. **Decision trees** on `HMDA - 2017 - TX`, by `sex`

Random forest prediction set metrics			
	Baseline	**Simple**	**Super**
$\triangle\hat{\mathbf{PR}}$	$4.4 \pm 0.0\%$	$4.0 \pm 0.0\%$	$5.2 \pm 0.1\%$
$\hat{\mathbf{PR}}_F$	$74.1 \pm 0.1\%$	$85.5 \pm 0.1\%$	$76.3 \pm 0.2\%$
$\hat{\mathbf{PR}}_M$	$78.5 \pm 0.1\%$	$89.5 \pm 0.1\%$	$81.5 \pm 0.1\%$
$\triangle\hat{\mathbf{Err}}$	$1.0 \pm 0.1\%$	$0.4 \pm 0.0\%$	$0.7 \pm 0.1\%$
$\hat{\mathbf{Err}}_F$	$15.5 \pm 0.1\%$	$4.6 \pm 0.0\%$	$9.3 \pm 0.0\%$
$\hat{\mathbf{Err}}_M$	$14.5 \pm 0.0\%$	$4.2 \pm 0.0\%$	$8.6 \pm 0.1\%$
$\triangle\hat{\mathbf{FPR}}$	$0.3 \pm 0.1\%$	$0.3 \pm 0.0\%$	$0.4 \pm 0.1\%$
$\hat{\mathbf{FPR}}_F$	$7.4 \pm 0.1\%$	$1.7 \pm 0.0\%$	$3.5 \pm 0.1\%$
$\hat{\mathbf{FPR}}_M$	$7.1 \pm 0.0\%$	$2.0 \pm 0.0\%$	$3.9 \pm 0.0\%$
$\triangle\hat{\mathbf{FNR}}$	$0.7 \pm 0.0\%$	$0.8 \pm 0.0\%$	$1.1 \pm 0.1\%$
$\hat{\mathbf{FNR}}_F$	$8.1 \pm 0.1\%$	$3.0 \pm 0.0\%$	$5.8 \pm 0.1\%$
$\hat{\mathbf{FNR}}_M$	$7.4 \pm 0.1\%$	$2.2 \pm 0.0\%$	$4.7 \pm 0.0\%$
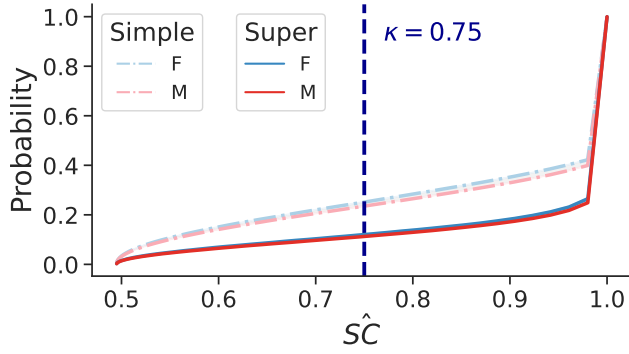
Abstention set metrics		
	Simple	**Super**
$\triangle\hat{\mathbf{AR}}$	$1.7 \pm 0.0\%$	$0.8 \pm 0.2\%$
$\hat{\mathbf{AR}}_F$	$28.4 \pm 0.1\%$	$11.9 \pm 0.2\%$
$\hat{\mathbf{AR}}_M$	$26.7 \pm 0.1\%$	$11.1 \pm 0.0\%$

Figure 55. **Random forests** on `HMDA - 2017 - TX`, by `sex`

E.4.7. DISCUSSION OF EXTENDED RESULTS FOR ALGORITHM 1

Overall, our results support that examining self-consistency and error together provide a much richer picture of model behavior, both with respect to arbitrariness and fairness metric disparities. Particularly in smaller datasets, the learning process produces models with a large degree of variance. As a result, ensembling with confidence can lead to huge abstention rates.

Improving self-consistency by doing a round of variance reduction first and then ensembling with confidence (i.e., super-ensembling) can lead to improvements in error over baselines while having a lower abstention rate. These improvements are typically shared across subgroups, but may not be symmetric; some subgroups may benefit more than others. As a result, even though accuracy increases absolutely for both groups, relative fairness metrics can decrease. This is a different instantiation of the fairness-accuracy trade-off than is often written about, which posits a necessary decrease in accuracy for one subgroup to improve fairness between binarized groups. Our results suggest that it is worth first tuning for accuracy, and then seeing how fairness interventions can balance the benefits across subgroups. Of course, it is possible that doing this could lead to injecting variance back into the model outputs, thereby reducing self-consistency and inducing arbitrariness. We leave this investigation to future work.

Additionally, our results reify that choice of model matters a lot. While overall error rates across model types may be similar, the sources of that error are not necessarily the same. This is an obvious point, relating to bias and variance. However, a lot of fair classification work describes similar performance across logistic regression, decision trees, random forests, SVMs, and MLPs (e.g., Chen et al. (2018)). Looking at self-consistency confirms that this is not the case, with decision trees and random forests in particular exhibiting higher variance, and thus being more amenable to variance reduction and improvements in overall error. As fair classification research transitions to larger benchmarks, it will likely be fruitful to investigate more complex model classes.

We provide run times on our cluster environment in Table 2. We did not select for CPUs with any particular features, and thus the run times are quite variable.

Table 2. These times are recorded for our cluster environment (hh:mm:ss), described in Appendix E.2 for our Algorithm 1 experiments. At the time of running, due to time constraints, the authors had not yet parallelized this part of the code.

Dataset	g	Logistic regression	Decision trees	Random forests
South German Credit	sex	00:42:50	00:25:28	00:34:42
COMPAS	race	00:57:05	00:39:24	00:31:47
Old Adult	sex	01:08:37	01:23:39	00:57:11
Taiwan Credit	sex	00:31:35	01:34:57	01:53:33
New Adult – CA				
Income	sex, race	01:39:53	02:51:13	04:59:07
Employment	sex, race	02:20:15	02:18:16	03:00:15
Public Coverage	sex, race	01:13:33	02:02:57	02:24:08
HMDA – 2017				
NY	sex, race, ethnicity	03:50:52	05:00:19	05:39:44
TX	sex, race, ethnicity	05:18:59	04:10:34	04:18:59

We also provide details on systematic arbitrariness, which we measure using the $\hat{\mathcal{W}}_1$ (Appendix B.3.1). As noted in Appendix B.3.1, since this metric is an average, its measure necessarily changes if we compute it over a different set of levels $\hat{\mathbb{K}}$. To make distances across interventions comparable, we treat CDF values below κ as 0, so that all of the probability mass is on $\hat{\text{SC}} \geq \kappa$. We therefore report two versions of these results, those for no abstention and those that account for abstention at values $< \kappa$. Δ is the difference between $\hat{\mathcal{W}}_{1\,\text{Simple}} - \hat{\mathcal{W}}_{1\,\text{Super}}$. Positive differences indicate cases for which the super-ensembling method decreases the Wasserstein-1 distance between subgroup CDFs; negative differences indicate increases. While in some cases there is an increase, it is worth noting that this aligns with cases for which the $\hat{\mathcal{W}}_1$ distance is very close to 0. Old Adult, highlighted below, is the only dataset that exhibts large amounts of systematic arbitrariness (for decision tress and random forests, in particular; it exhibits the highest amount for logistic regression, but is overall low). Old Adult and New Adult -- Employment (by sex) are two of the only tasks that show any fairness disparities that are $> 3\%$.

Table 3. Empirical Wasserstein-1 ($\hat{\mathcal{W}}_1$) measurements without abstention

Dataset	g	Logistic regression			Decision trees			Random forests		
		Simple	**Super**	Δ	**Simple**	**Super**	Δ	**Simple**	**Super**	Δ
German Credit	sex	0.0181	0.0101	0.0079	0.0162	0.0244	-0.0082	0.0181	0.0175	0.0006
COMPAS	race	0.0073	0.0043	0.0030	0.0189	0.0170	0.0019	0.0073	0.0031	0.0043
Old Adult	sex	0.0206	0.0033	0.0173	0.1386	0.0273	0.1112	0.1266	0.0255	0.1011
Taiwan Credit	sex	0.0028	0.0007	0.0020	0.0223	0.0108	0.0115	0.0240	0.0065	0.0175
New Adult – CA										
Income	sex	0.0009	0.0003	0.0006	0.0138	0.0055	0.0083	0.0089	0.0018	0.0071
	race	0.0003	0.0001	0.0002	0.0170	0.0073	0.0096	0.0163	0.0055	0.0108
Employment	sex	0.0004	0.0002	0.0003	0.0010	0.0011	-0.0001	0.0043	0.0031	0.0013
	race	0.0004	0.0001	0.0004	0.0020	0.0007	0.0013	0.0021	0.0015	0.0006
Public Coverage	sex	0.0004	0.0001	0.0003	0.0029	0.0024	0.0005	0.0045	0.0023	0.0024
	race	0.0010	0.0003	0.0007	0.0200	0.0089	0.0113	0.0235	0.0089	0.0147
HMDA – 2017										
NY	sex	0.0002	0.0004	-0.0002	0.0096	0.0039	0.0056	0.0080	0.0023	0.0056
	race	0.0009	0.0005	0.0005	0.0433	0.0203	0.0231	0.0409	0.0133	0.0276
	ethnicity	0.0005	0.0005	0.0000	0.0229	0.0156	0.0073	0.0248	0.0108	0.0140
TX	sex	0.0001	0.0001	0.0000	0.0153	0.0097	0.0055	0.0113	0.0054	0.0058
	race	0.0001	0.0001	0.0000	0.0010	0.0012	-0.0002	0.0013	0.0007	0.0006
	ethnicity	0.0007	0.0002	0.0004	0.0509	0.0291	0.0219	0.0379	0.0190	0.0188

Table 4. Empirical Wasserstein-1 ($\hat{\mathcal{W}}_1$) measurements with abstention using $\kappa \geq .75$

Dataset	g	Logistic regression			Decision trees			Random forests		
		Simple	**Super**	Δ	**Simple**	**Super**	Δ	**Simple**	**Super**	Δ
German Credit	sex	0.0113	0.0080	0.0034	0.0090	0.0094	-0.0004	0.0084	0.0132	-0.0048
COMPAS	race	0.0035	0.0019	0.0017	0.0039	0.0060	-0.0021	0.0041	0.0019	0.0022
Old Adult	sex	0.0110	0.0020	0.0090	0.0654	0.0155	0.0500	0.0634	0.0139	0.0494
Taiwan Credit	sex	0.0014	0.0005	0.0009	0.0057	0.0059	-0.0002	0.0107	0.0040	0.0067
New Adult – CA										
Income	sex	0.0005	0.0002	0.0004	0.0051	0.0032	0.0019	0.0047	0.0012	0.0035
	race	0.0002	0.0000	0.0002	0.0073	0.0040	0.0033	0.0082	0.0028	0.0053
Employment	sex	0.0002	0.0001	0.0001	0.0005	0.0005	0.0001	0.0020	0.0014	0.0006
	race	0.0002	0.0000	0.0002	0.0012	0.0003	0.0008	0.0008	0.0005	0.0003
Public Coverage	sex	0.0002	0.0001	0.0001	0.0006	0.0012	-0.0006	0.0011	0.0009	0.0002
	race	0.0006	0.0001	0.0005	0.0068	0.0049	0.0019	0.0106	0.0047	0.0059
HMDA – 2017										
NY	sex	0.0001	0.0002	-0.0001	0.0033	0.0020	0.0012	0.0040	0.0013	0.0028
	race	0.0004	0.0002	0.0002	0.0155	0.0111	0.0044	0.0190	0.0076	0.0114
	ethnicity	0.0002	0.0002	0.0001	0.0055	0.0083	-0.0028	0.0081	0.0059	0.0022
TX	sex	0.0000	0.0000	0.0000	0.0061	0.0050	0.0011	0.0058	0.0029	0.0028
	race	0.0000	0.0001	0.0000	0.0004	0.0005	-0.0002	0.0007	0.0004	0.0003
	ethnicity	0.0003	0.0001	0.0003	0.0229	0.0159	0.0070	0.01200	0.0104	0.0095

E.5. Reliability and fairness metrics in `COMPAS` and `South German Credit`

Even before we apply our intervention to improve self-consistency, our results in Section 3 show close-to-parity $\hat{\mathrm{Err}}$, $\hat{\mathrm{FPR}}$, and $\hat{\mathrm{FNR}}$ across subgroups in `COMPAS` (and similarly for `South German Credit`, below). These results are surprising. We run $B = 101$ models to produce estimates of variance and self-consistency, but of course doing this also has the effect of estimating the expected error more generally (with variance representing a portion of that error). Our estimates of expected error for these tasks indicate that the average model produced training on `COMPAS` and `South German Credit`, with respect to popular fairness definitions like Equality of Opportunity and Equalized Odds (Barocas et al., 2019; Hardt et al., 2016) *are in fact baseline close to parity, with **no** fairness intervention applied*. We found this across model types for both datasets, though the story becomes more complicated when we apply techniques to improve self-consistency (discussed in Appendix E.4.7).

Of course, we did not expect this result, as these are two of the *de facto* standard benchmark datasets in algorithmic fairness. They are used in countless other studies to probe and verify algorithmic fairness interventions (Fabris et al., 2022). As a result, we initially thought that our results must be incorrect. We therefore looked at the underlying models in our bootstrap runs to see the error of the underlying models.

We re-ran our baseline experiments with $B = 1001$ and for 100 test/train splits for logistic regression. In Figure 56, we plot the $(100, 100)$ bootstrap models that went into these results. For another view on analogous information, in Table 5, we provide an excerpt of the results for `COMPAS` regarding the underlying 1010 random forest classifiers used to produce Figure 2a.

(a) $\hat{\mathrm{Err}}$ disparity   (b) $\hat{\mathrm{FPR}}$ disparity   (c) $\hat{\mathrm{FNR}}$ disparity
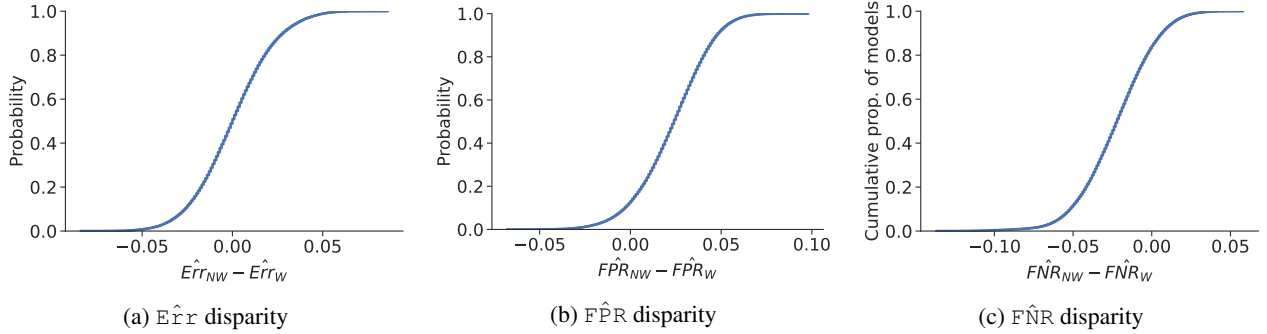
Figure 56. Cumulative distribution of error disparity across 100, 100 logistic regression models trained on `COMPAS`.

Overall, we can see that there is a wide range of error disparities that trend in both directions, with a skew toward higher $\hat{\mathrm{FPR}}$ for $g = \mathrm{NW}$. These results support our claim that training many models is necessary to get an accurate picture of expected error, with implications both for reproducibility of experiments that just train and analyze a small handful of models and for generalizability. There are models that exhibit worse degrees of unfairness in both directions, but they are more unlikely than models that exhibit smaller disparities.

We subset the above results to the 100 models that produce the lowest $\hat{\mathrm{Err}}$, as this is often the selection critera for picking models to post-process. We plot these results below. These top-performing models in fact exhibit (on average) closer-to-parity for $\hat{\mathrm{FPR}}$ and $\hat{\mathrm{FNR}}$.

(a) $\hat{\mathrm{Err}}$ disparity   (b) $\hat{\mathrm{FPR}}$ disparity   (c) $\hat{\mathrm{FNR}}$ disparity

Figure 57. CDF of error disparity across the top 100 logistic regression models (of the 100, 100 models) trained on `COMPAS`.

Table 5. Comparing subgroup error rates in `COMPAS` for different random forest classifiers trained to produce Figure 2a. Each table looks at the top-3 highest differences between subgroups for the specified metri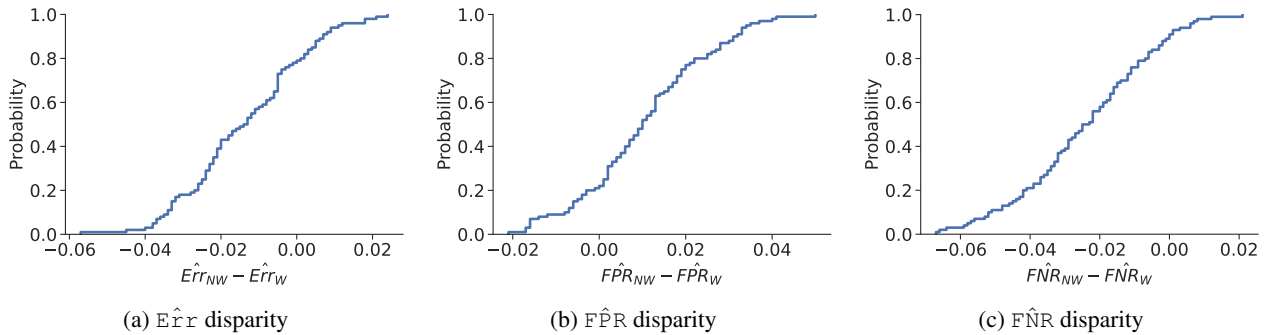c: (a) $\hat{\text{Err}}_{\text{NW}} - \hat{\text{Err}}_{\text{W}}$, when $\hat{\text{Err}}_{\text{NW}} > \hat{\text{Err}}_{\text{W}}$; (b) $\hat{\text{Err}}_{\text{W}} - \hat{\text{Err}}_{\text{NW}}$, when $\hat{\text{Err}}_{\text{W}} > \hat{\text{Err}}_{\text{NW}}$; (c) $\hat{\text{FPR}}_{\text{NW}} - \hat{\text{FPR}}_{\text{W}}$, when $\hat{\text{FPR}}_{\text{NW}} > \hat{\text{FPR}}_{\text{W}}$; (d) $\hat{\text{FPR}}_{\text{W}} - \hat{\text{FPR}}_{\text{NW}}$, when $\hat{\text{FPR}}_{\text{W}} > \hat{\text{FPR}}_{\text{NW}}$; (e) $\hat{\text{FNR}}_{\text{NW}} - \hat{\text{FNR}}_{\text{W}}$, when $\hat{\text{FNR}}_{\text{NW}} > \hat{\text{FNR}}_{\text{W}}$; and, (f) (e) $\hat{\text{FNR}}_{\text{W}} - \hat{\text{FNR}}_{\text{NW}}$, when $\hat{\text{FNR}}_{\text{W}} > \hat{\text{FNR}}_{\text{NW}}$. We highlight the overall error metric in gray, the larger metric (being subtracted from) in blue, the smaller metric (being subtracted) in red, and the difference in the metric between subgroups in purple. Note that run 757 appears twice, which we mark in orange.

(a) The top-3 most unfair models according to subgroup-specific $\hat{\text{Err}}$, when $\hat{\text{Err}}_{\text{NW}} > \hat{\text{Err}}_{\text{W}}$ (i.e., unfair toward NW).

Run #	s	b	$\hat{\text{Err}}$	$\hat{\text{FPR}}$	$\hat{\text{FNR}}$	$\hat{\text{Err}}_{\text{NW}}$	$\hat{\text{FPR}}_{\text{NW}}$	$\hat{\text{FNR}}_{\text{NW}}$	$\hat{\text{Err}}_{\text{W}}$	$\hat{\text{FPR}}_{\text{W}}$	$\hat{\text{FNR}}_{\text{W}}$	$\hat{\text{Err}}_{\text{NW}} - \hat{\text{Err}}_{\text{W}}$
762	8	504	0.374	0.179	0.196	0.405	0.204	0.201	0.315	0.13	0.186	0.09
757	8	464	0.369	0.167	0.202	0.395	0.201	0.193	0.318	0.101	0.218	0.077
328	4	116	0.371	0.165	0.206	0.395	0.181	0.214	0.323	0.134	0.189	0.072

(b) The top-3 most unfair models according to subgroup-specific $\hat{\text{Err}}$, when $\hat{\text{Err}}_{\text{W}} > \hat{\text{Err}}_{\text{NW}}$ (i.e., unfair toward W).

Run #	s	b	$\hat{\text{Err}}$	$\hat{\text{FPR}}$	$\hat{\text{FNR}}$	$\hat{\text{Err}}_{\text{NW}}$	$\hat{\text{FPR}}_{\text{NW}}$	$\hat{\text{FNR}}_{\text{NW}}$	$\hat{\text{Err}}_{\text{W}}$	$\hat{\text{FPR}}_{\text{W}}$	$\hat{\text{FNR}}_{\text{W}}$	$\hat{\text{Err}}_{\text{W}} - \hat{\text{Err}}_{\text{NW}}$
414	5	75	0.376	0.167	0.209	0.352	0.158	0.194	0.422	0.186	0.236	0.07
435	5	180	0.376	0.199	0.177	0.355	0.189	0.166	0.416	0.217	0.198	0.061
413	5	70	0.378	0.189	0.189	0.359	0.188	0.171	0.413	0.191	0.222	0.054

(c) The top-3 most unfair models according to subgroup-specific $\hat{\text{FPR}}$, when $\hat{\text{FPR}}_{\text{NW}} > \hat{\text{FPR}}_{\text{W}}$ (i.e., unfair toward NW).

Run #	s	b	$\hat{\text{Err}}$	$\hat{\text{FPR}}$	$\hat{\text{FNR}}$	$\hat{\text{Err}}_{\text{NW}}$	$\hat{\text{FPR}}_{\text{NW}}$	$\hat{\text{FNR}}_{\text{NW}}$	$\hat{\text{Err}}_{\text{W}}$	$\hat{\text{FPR}}_{\text{W}}$	$\hat{\text{FNR}}_{\text{W}}$	$\hat{\text{FPR}}_{\text{NW}} - \hat{\text{FPR}}_{\text{W}}$
757	8	464	0.369	0.167	0.202	0.395	0.201	0.193	0.318	0.101	0.218	0.1
729	8	240	0.358	0.162	0.197	0.376	0.189	0.187	0.323	0.107	0.216	0.082
791	8	736	0.377	0.171	0.205	0.395	0.198	0.197	0.341	0.118	0.222	0.08

(d) The top-3 most unfair models according to subgroup-specific $\hat{\text{FPR}}$, when $\hat{\text{FPR}}_{\text{W}} > \hat{\text{FPR}}_{\text{NW}}$ (i.e., unfair toward W).

Run #	s	b	$\hat{\text{Err}}$	$\hat{\text{FPR}}$	$\hat{\text{FNR}}$	$\hat{\text{Err}}_{\text{NW}}$	$\hat{\text{FPR}}_{\text{NW}}$	$\hat{\text{FNR}}_{\text{NW}}$	$\hat{\text{Err}}_{\text{W}}$	$\hat{\text{FPR}}_{\text{W}}$	$\hat{\text{FNR}}_{\text{W}}$	$\hat{\text{FPR}}_{\text{W}} - \hat{\text{FPR}}_{\text{NW}}$
639	7	280	0.36	0.187	0.173	0.352	0.174	0.178	0.376	0.212	0.164	0.038
807	9	72	0.381	0.191	0.19	0.372	0.179	0.192	0.398	0.214	0.184	0.035
543	6	264	0.358	0.155	0.203	0.351	0.144	0.206	0.37	0.175	0.196	0.031

(e) The top-3 most unfair models according to subgroup-specific $\hat{\text{FNR}}$, when $\hat{\text{FNR}}_{\text{NW}} > \hat{\text{FNR}}_{\text{W}}$ (i.e., unfair toward NW).

Run #	s	b	$\hat{\text{Err}}$	$\hat{\text{FPR}}$	$\hat{\text{FNR}}$	$\hat{\text{Err}}_{\text{NW}}$	$\hat{\text{FPR}}_{\text{NW}}$	$\hat{\text{FNR}}_{\text{NW}}$	$\hat{\text{Err}}_{\text{W}}$	$\hat{\text{FPR}}_{\text{W}}$	$\hat{\text{FNR}}_{\text{W}}$	$\hat{\text{FNR}}_{\text{NW}} - \hat{\text{FNR}}_{\text{W}}$
246	3	141	0.379	0.166	0.213	0.398	0.169	0.229	0.345	0.161	0.184	0.045
506	6	42	0.367	0.17	0.197	0.386	0.175	0.211	0.332	0.161	0.171	0.04
204	3	15	0.384	0.185	0.199	0.394	0.181	0.213	0.365	0.192	0.173	0.04

(f) The top-3 most unfair models according to subgroup-specific $\hat{\text{FNR}}$, when $\hat{\text{FNR}}_{\text{W}} > \hat{\text{FNR}}_{\text{NW}}$ (i.e., unfair toward W).

Run #	s	b	$\hat{\text{Err}}$	$\hat{\text{FPR}}$	$\hat{\text{FNR}}$	$\hat{\text{Err}}_{\text{NW}}$	$\hat{\text{FPR}}_{\text{NW}}$	$\hat{\text{FNR}}_{\text{NW}}$	$\hat{\text{Err}}_{\text{W}}$	$\hat{\text{FPR}}_{\text{W}}$	$\hat{\text{FNR}}_{\text{W}}$	$\hat{\text{FNR}}_{\text{W}} - \hat{\text{FNR}}_{\text{NW}}$
474	5	375	0.373	0.175	0.199	0.356	0.183	0.174	0.406	0.159	0.247	0.073
401	5	10	0.378	0.189	0.19	0.363	0.197	0.167	0.406	0.173	0.233	0.066
52	1	53	0.367	0.172	0.196	0.351	0.178	0.173	0.397	0.16	0.238	0.065

This detailed view provides insight into how such a result is possible. Broadly speaking, individual runs have roughly similar error;[16] yet, the subgroup-specific error rates that compose the overall error can nevertheless vary widely depending on the underlying training data. This observation aligns with current interest in *model multiplicity* in the algorithmic fairness community (Black et al., 2022b; Watson-Daniels et al., 2022), which imports the idea from Breiman (2001). In this case, as suggested by Table 5, there are models that demonstrate unfairness toward both subgroups with respect to each error rate metric $\hat{Err}$, $\hat{FPR}$, and $\hat{FNR}$. When we move away from attempting to find a *single* model that performs well (accurately or fairly) on COMPAS, and instead consider the information contained across different possible models, we yield the result that the average, expected behavior smooths over the variance in underlying models such that the result is close to fair: ***The average of unfair models with high variance in subgroup error rates is essentially fair.***

Stability analysis. To verify the stability of this result, we re-execute our experiments for increasing numbers of train/test splits S and replicates B. While our results for COMPAS are generally tight for small S (e.g., Figures 2a and 5), this was not the case for German Credit, for which it was difficult to estimate self-consistency consistently. As a result, for COMPAS, we did not expect markedly different results for increased S. Our results for $S = 100, B = 1001$ using logistic regression (Figure 58, Table 6) confirm this intuition.
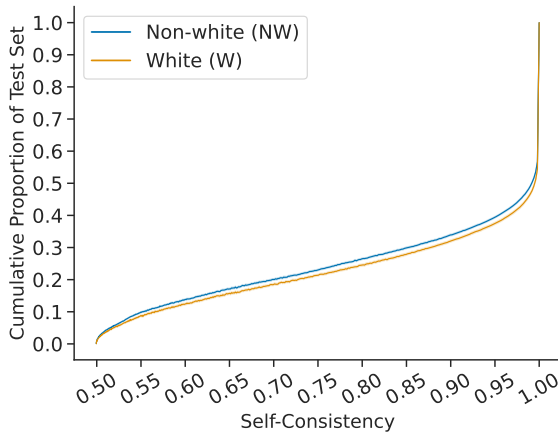

Figure 58. COMPAS split by $g = $ race, $B = 1001, S = 100$

Table 6. Mean $\pm$ STD across $S = 100$ train/test splits $\times\ B = 1001$ runs.

	COMPAS			
	$\hat{Err}$	$\hat{FPR}$	$\hat{FNR}$	$\hat{SC}$
Total	$.333 \pm .008$	$.14 \pm .009$	$.192 \pm .01$	$.883 \pm .004$
$g = $ NW	$.333 \pm .01$	$.148 \pm .011$	$.185 \pm .012$	$.88 \pm .005$
$g = $ W	$.332 \pm .014$	$.125 \pm .013$	$.207 \pm .016$	$.888 \pm .006$

We provide analogous results for German Credit, with $S = 1000, B = 1001$ using random forests (Figure 59, Table 7). It takes an enormous number of runs to produce stable estimates of error and $\hat{SC}$ for German Credit, which indicate statistical equality across groups. Arguably, our results below for $1,001,000$ models still are very high variance (certainly with respect to error metrics). This task really has too few data points (≈ 600) to generalize reliably.
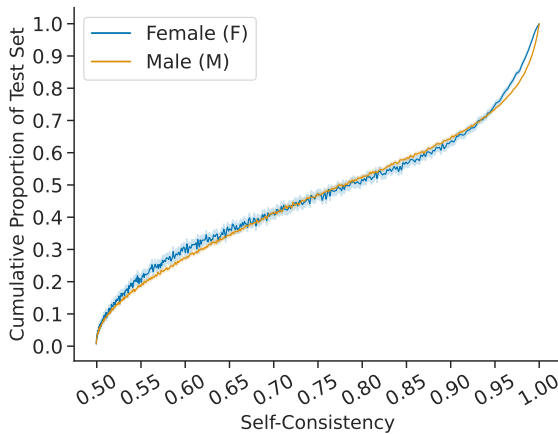

Figure 59. German Credit split by $g = $ sex, $S = 1000, B = 100$

Table 7. Mean $\pm$ STD across $S = 1000$ train/test splits $\times\ B = 1001$ runs.

	South German Credit			
	$\hat{Err}$	$\hat{FPR}$	$\hat{FNR}$	$\hat{SC}$
Total	$.28 \pm .021$	$.173 \pm .028$	$.107 \pm .017$	$.769 \pm .015$
$g = $ F	$.288 \pm .064$	$.183 \pm .072$	$.105 \pm .037$	$.766 \pm .04$
$g = $ M	$.279 \pm .023$	$.171 \pm .029$	$.108 \pm .018$	$.769 \pm .016$

---

[16]This should be taken relatively. In general, COMPAS demonstrates high error; the error is relatively tight given just how much error there is. The error fluctuates depending on the training data, but the average error rate across train/test splits is rather tight, despite the fluctuations in error within the B runs of each split.