Measuring Model Fairness under Noisy Covariates: A Theoretical Perspective

Flavien Prost, Pranjal Awasthi, *

Nick Blumm, Aditee Kumthekar, Trevor Potter

Li Wei, Xuezhi Wang, Ed H. Chi, Jilin Chen, Alex Beutel

Google

USA

ABSTRACT

In this work we study the problem of measuring the fairness of a machine learning model under noisy information. Focusing on group fairness metrics, we investigate the particular but common situation when the evaluation requires controlling for the confounding effect of covariate variables. In a practical setting, we might not be able to jointly observe the covariate and group information, and a standard workaround is to then use proxies for one or more of these variables. Prior works have demonstrated the challenges with using a proxy for sensitive attributes, and strong independence assumptions are needed to provide guarantees on the accuracy of the noisy estimates. In contrast, in this work we study using a proxy for the covariate variable and present a theoretical analysis that aims to characterize weaker conditions under which accurate fairness evaluation is possible. Furthermore, our theory identifies potential sources of errors and decouples them into two interpretable parts γ and ϵ . The first part γ depends solely on the performance of the proxy such as precision and recall, whereas the second part ϵ captures correlations between all the variables of interest. We show that in many scenarios the error in the estimates is dominated by γ via a linear dependence, whereas the dependence on the correlations ϵ only constitutes a lower order term. As a result we expand the understanding of scenarios where measuring model fairness via proxies can be an effective approach. Finally, we compare, via simulations, the theoretical upper-bounds to the distribution of simulated estimation errors and show that assuming some structure on the data, even weak, is key to significantly improve both theoretical guarantees and empirical results.

CCS CONCEPTS

• **Computing methodologies** → **Machine learning**.

KEYWORDS

ML fairness, Statistical parity, Noisy covariates.

*The first two authors contributed equally.

This work is licensed under a Creative Commons Attribution International 4.0 License.

AIES '21, May 19–21, 2021, Virtual Event, USA.

© 2021 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-8473-5/21/05.

<https://doi.org/10.1145/3461702.3462603>

ACM Reference Format:

Flavien Prost, Pranjal Awasthi, and Nick Blumm, Aditee Kumthekar, Trevor Potter, Li Wei, Xuezhi Wang, Ed H. Chi, Jilin Chen, Alex Beutel. 2021. Measuring Model Fairness under Noisy Covariates: A Theoretical Perspective. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society (AIES '21)*, May 19–21, 2021, Virtual Event, USA. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3461702.3462603>

1 INTRODUCTION

As machine learning (ML) systems permeate everyday life it is of utmost importance to understand and correct for the underlying biases present in such systems. As a result there has been a significant interest in recent years in the area of *algorithmic fairness* [7, 10, 25]. An important and often overlooked aspect of addressing biases in AI systems is the challenge of how to actually compute the fairness metric in practice. Given an ML model, fairness metrics typically involve evaluating the disparity in the model performance across different demographics and/or data slices of interest, according to a well defined metric such as *false positive rate*, *calibration* etc. However, often these criteria cannot be exactly evaluated as they require information that is only sparsely available in practical applications. How should one approach measuring fairness metrics in such situations? Our goal is to address this question from a theoretical perspective.

We focus on scenarios that involve measuring group fairness criteria where a practitioner wants to account for a covariate that is not available jointly with the group information. This challenge arises in a variety of applications. For example, consider the case of evaluating a model predicting the toxicity of an online comment, where we want to measure differences in performance of the model across toxic comments directed at different demographic groups, but we only want to compare performance within the same topic or language [11, 19, 30, 31]. In order to measure the performance of the model, one would ideally want access to a corpus of comments, each containing information about the demographic group, topic or language, and the prediction of the model for that comment. However, in practice such joint observations are rarely available. In this case, one would frequently use a model to predict the covariate (topic or language) for each comment, but might wonder how inaccuracies in this model of topic or language influence our measurement of a bias in the toxicity prediction. Another common occurrence is in the context of a recommender system, where a model predicts whether a user would click an item, but a practitioner might want to only consider fairness over high quality items, such as avoiding clickbait or focusing on when the user would actually be satisfied with the item [1, 18, 26, 27]. Again, it is unlikely that a dataset contains

together clicks, item group information, and quality or post-click engagement information; can a practitioner use a model of item quality to accurately estimate the bias in the recommender? Across all of these examples it's important to understand the confidence in our measurements.

Group fairness has received significant attention in recent years resulting in a variety of natural fairness criteria [3] that machine learning systems should satisfy. Our work concerns two common variants, *Statistical Parity* [12] and *Equality of opportunity* [15], which measure the difference of positive rates (or true positive rates) across sub-populations. To model the presence of the confounding or mediating covariate as described in the applications above, one can extend fairness metrics by conditioning on this covariate and this approach has been used in a variety of prior works [4, 7, 17, 23, 24, 32].

We summarize our targeted fairness metric as follows. Consider a machine learning system making a prediction $y \in \{0, 1\}$, a sensitive attribute $\ell \in \{0, 1\}$ and a covariate $v \in \{0, 1\}$. The *Conditional Statistical Parity* [8] metric is defined as:

$$G_{SP} = \mathbb{P}[y=1|v=1, \ell=0] - \mathbb{P}[y=1|v=1, \ell=1]. \quad (1)$$

Similarly, let $y^* \in \{0, 1\}$ be the ground truth label. Then the *Conditional Equal Opportunity* metric is defined as:

$$G_{EO} = \mathbb{P}[y=1|y^*=1, v=1, \ell=0] - \mathbb{P}[y=1|y^*=1, v=1, \ell=1]. \quad (2)$$

For ease of exposition and reducing notation, the rest of the paper focuses on the metric defined in (1). *However, all of our theoretical results trivially generalize to the case of conditional equal opportunity as well by focusing on sub-populations where $y^* = 1$.*

While it is important to incorporate legitimate covariates in fairness metrics, evaluating the metrics accurately requires a large dataset that contains both the group label ℓ , the covariate information v and possibly the label y^* (in the case of equal opportunity). As we discussed before, in practice we might only have an estimate $\hat{v}$ derived from either a proxy attribute in the data that is correlated with v , or a classifier trained to predict v . A natural question that comes up is whether one can output

$$\hat{G}_{SP} = \mathbb{P}[y=1|\hat{v}=1, \ell=0] - \mathbb{P}[y=1|\hat{v}=1, \ell=1] \quad (3)$$

as an estimate of the bias G_{SP} of the model as in (1).

For the sake of clarity, let's return to our example from above of a toxicity model for online comments, which predicts whether a comment contains toxic language ($y = 1$ if toxic) and for which prior works have demonstrated that these models often exhibit biases towards certain demographics [11, 19, 30, 31]. At the same time, the level of toxicity might be significantly influenced by the topic of the comment (i.e., topic is covariate v) and we may want to control for its impact in our fairness analysis. Following the Statistical Parity approach, a reasonable fairness metric would be given by Eq. (1)¹. As both topic and demographic information are typically available for only a small portion of examples, a practitioner might not have access to the joint distribution of (v, ℓ) but a topic classifier $\hat{v}$ might

be available. In that situation, how confident can a practitioner be to use $\hat{G}_{SP}$ as an estimation of its intended metric G_{SP} ?

In the rest of the paper we study conditions under which we can bound the estimation error $|G_{SP} - \hat{G}_{SP}|$. In Section 4, we highlight through simple examples two different sources that will affect the estimation error: (a) the performance of the proxy ($\hat{v}$) and, (b) the joint correlations between variables (y, v and $\hat{v}$) that can be represented as the value of the outcome variable y over the confusion matrix of $(v, \hat{v})$. The bulk of the theory aims to derive error bounds based on the extent to which the data distribution satisfies certain conditions that depend on the above two sources. Throughout the paper, parameters related to the classifier performance will be captured by γ , and parameters related to the joint correlations between variables y, v and $\hat{v}$ will be captured by ϵ . These parameters will always take values in $[0, 1]$.

In Section 5, we derive a bound on the estimation error based solely on the precision and recall of the proxy $\hat{v}$. If $1 - \gamma_A$ is the lower bound on the worst case precision or recall of the proxy on any data slice, i.e., conditioned on $\ell = 0$ or $\ell = 1$, then we show that the error $|G_{SP} - \hat{G}_{SP}|$ can be bounded by $2 \cdot \gamma_A$. However, $2 \cdot \gamma_A$ may not always be small enough for fairness evaluation and our goal will be to progressively refine this bound by exploiting other correlations between variables ($y, v, \hat{v}$ and ℓ).

In Section 6, we identify two cases depending on correlations which might lead to “errors canceling each other”. In the first case (B1), if the precision and recall of $\hat{v}$ are γ_{B1} -close to each other (for $\ell = 0, 1$), then we show that the error is bounded by a linear combination of two terms γ_{B1} and ϵ_{B1} . The term ϵ_{B1} captures how the variables of interest, i.e., y, v and $\hat{v}$ correlated within a data slice ($\ell = 0$ or $\ell = 1$). In the second case (B2), we study how the error behaves as a function of closeness of precision across slices ($\ell = 0$ vs $\ell = 1$), and recall across slices. In this case we again show that the error depends on a linear combination of two terms γ_{B2} and ϵ_{B2} , one capturing relative closeness of precision and recall, and the other how much the correlations between y, v and $\hat{v}$ vary between $\ell = 0$ and $\ell = 1$.

Having identified the various sources of errors above, in Section 7, we provide a refined bound that captures how the various errors interact with each other. As a result we show in Theorem 7.1 that the final estimation error depends linearly in the γ parameters and quadratically in ϵ parameters. Hence, the various errors due to correlations have a multiplying effect leading to a lower order term. Our analysis thereby reveals that the properties of the proxy classifier such as precision and recall play a more important role in the final estimation error, but that the structure of the correlations through ϵ can help reduce the bound on $|G - \hat{G}|$.

Finally, we use a simulation framework to compare our theoretical bounds on estimation errors to their simulated counterparts. We show that our worst-case bounds are met when no assumption is imposed on the data generation process. Furthermore, we show how enforcing some (even weak) structure on the correlations (through the parameter ϵ) can significantly improve both theoretical bounds and empirical distributions, compared to when we rely only on precision and recall.

Our work is similar in spirit to prior works on studying the effect of label noise (outcome variable y^*) on fairness metrics [13, 20],

¹Note, in practice one might prefer to follow the equality of opportunity metric for this or other applications. We use the conditional statistical parity metric throughout the paper for simplicity of notation, but our results easily extend to the conditional equality of opportunity metric as well, and we encourage practitioners to always consider which metric fits best with their application.

as well as noise in the sensitive attribute ℓ [2, 5, 21]. The case of label noise is captured by our theoretical analysis as one can simply take the noisy covariate to be y^* . The case of noise in the sensitive attributes is not directly comparable to our setting. We discuss this more in Section 2.

The rest of the paper is structured as follows. In Section 2 and Section 3 we review related work and discuss preliminaries. In Section 4 we discuss several simple cases where accurate bias estimation is possible provided certain independence assumptions hold. This section is meant to serve as a warm up to help the reader ease into the notation. In Section 5, we build our main theory by first bounding the estimation error solely based on the performance of the $\hat{v}$ classifier, and next identifying, in Section 6, two other conditions that affect the estimation error. We present our main theorem in Section 7 that shows that when all the conditions hold true to different extents, the estimation error can be decomposed into a linear part that captures the properties of the proxy and a quadratic part that captures the correlations. We conclude with experiments on simulated data in Section 8 and discuss conclusions in Section 9.

Note: We would like to point out that in this work we are focused on understanding the theoretical underpinnings of using proxies in evaluation, and *do not* advocate always using the approach of measuring the bias of model via a proxy attribute. As pointed out in prior works [2] the approach could have unintended side effects on the evaluation, and as such practitioners should carefully consider the risks of using such an approach. Our goal is to identify theoretically how errors propagate in the analysis to affect the final outcome. We hope that our work will help practitioners make more informed choices.

2 RELATED WORK

We focus on conditional metrics such as conditional statistical parity and conditional equal opportunity. Several recent works have argued the use of such metrics over or in addition to their unconditioned counterparts. For example, the works of Chouldechova [7], Corbett-Davies et al. [8], Ritov et al. [32] all explore conditioning on the number of prior convictions when measuring the bias of risk assessment tools for recidivism. This naturally leads to the conditional statistical parity and equality of opportunity metrics. The work of Beutel et al. [4] proposes the use of the conditional equal opportunity metric to account for differences in label confidence. Recent works on intersectional fairness consider metrics that involve conditioning on several covariates or arbitrary data slices described by functions of low VC-dimension [17, 22].

The problem of assessing unfair biases and training of fair machine learning models for the above discussed conditional metrics, under noisy label or covariate information has started to receive attention from the research community in recent years. Here we discuss the works most relevant in the context of our results.

Evaluating under uncertainty. The work of Chen et al. [6] studies estimating the *mean demographic disparity* when one has noisy information about the sensitive attribute. They analyze proxies based on threshold based estimators and present conditions under which the method over-estimates or under-estimates the true metric. However, unless strong independence assumptions are made

there is no guarantee that the estimation error will be small. The work of Kallus et al. [21] extends this to other fairness metrics such as equality of opportunity and presents methods to construct uncertainty intervals around the true value of the metric of interest. Again, excluding certain independence assumptions the provided intervals can be large.

The recent work of Awasthi et al. [2] studies what properties of the classifier used to construct the proxy affect the estimation error. They point out that accuracy of the classifier alone is not enough and advocate the use of active learning based algorithms to deal with uncertain information. We would like to point out that the above mentioned works consider noise in the sensitive attribute and are not directly comparable to our setting of noisy covariates. At a technical level, when the noisy variable ℓ is also the one for which we want to access performance gap, i.e., $\ell = 0$ vs. $\ell = 1$, it is not easy to decouple the various sources of error as we do in the case of noisy covariates. As a result we find that in our setting one can get reasonably small bounds even without strong independence assumptions.

The recent work of Fogliato et al. [13] studies how noise in the observed labels (the outcome variable) affects estimates of fairness metrics such as false positive rates, false negative rates, and positive predicted value. In principle the setting of label noise can be captured by our theory. However, the specific assumptions used make our work not directly comparable with that of Fogliato et al. [13]. We provide error bounds that depend on both the precision and recall, as well as other correlations between the variables. The authors in Fogliato et al. [13] only consider how the estimation error depends on the precision (i.e., the noise rate). As a result they need to make strong assumptions such as one-sided noise ($y^* = 1$), and noise affecting only one demographic group. We show on the other hand that by modeling other natural quantities of interest such as recall and correlations, one can avoid independence assumptions and at the same time obtain estimation error guarantees.

Training under uncertainty. Beyond measuring fairness metrics, a recent line of work has also studied the problem of training fair classifiers under noise in labels and/or sensitive attributes. The work of Jiang and Nachum [20] proposes a method to train a fair classifier via a re-weighting scheme given access to noisy labels. The works of Gupta et al. [14], Lamy et al. [29], Zafar et al. [35] show how to incorporate fairness constraints during training when there is noise in the sensitive attributes. Relatedly in training approaches, Hashimoto et al. [16], Lahoti et al. [28] study the extreme scenario where sensitive attributes are entirely unavailable and propose re-weighting approaches to boost low-performing groups; Coston et al. [9] apply transfer learning techniques to overcome missing sensitive attributes. Further recent work studies the trade-off between fairness and accuracy and how it is affected by label noise [34].

3 NOTATION AND PRELIMINARIES

As stated before, we will focus on the conditional statistical parity metric defined in Eq. (1). However all our results continue to hold for the conditional equal opportunity metric, Eq. (2), as well. Let ℓ, v, y denote $\{0, 1\}$ -valued random variables where ℓ is the sensitive attribute, v is the relevant covariate and y is the predicted outcome.

$\mathbb{P}[y=1 v=0, \hat{v}=1, \ell=l]$	$\mathbb{P}[y=1 v=1, \hat{v}=1, \ell=l]$
$\mathbb{P}[y=1 v=0, \hat{v}=0, \ell=l]$	$\mathbb{P}[y=1 v=1, \hat{v}=0, \ell=l]$

Table 1: Value of the outcome y over the confusion matrix of $v, \hat{v}$, conditioned on group $\ell = l$.

Our goal is to approximate

$$G := \mathbb{P}[y=1|v=1, \ell=1] - \mathbb{P}[y=1|v=1, \ell=0]. \quad (4)$$

by computing

$$\hat{G} := \mathbb{P}[y=1|\hat{v}=1, \ell=1] - \mathbb{P}[y=1|\hat{v}=1, \ell=0]. \quad (5)$$

We want to characterize conditions under which the error in our estimates, i.e., $|G - \hat{G}|$ can be bounded. We introduce below notations for two important parameters of our system.

Classifier Performance. We define the conditional precision and recall of the proxy $\hat{v}$ that will play an important role in our analysis.

$$\mathbb{P}[v=1|\hat{v}=1, \ell=l] = 1 - p_l \quad (6)$$

$$\mathbb{P}[\hat{v}=1|v=1, \ell=l] = 1 - r_l. \quad (7)$$

Constraints on the classifier performance will be captured by the variable γ .

Outcome y over the confusion matrix of $\hat{v}$. We will also study how correlations among the variables $y, v, \hat{v}$ affect the error in the measurement of G . In later sections we will quantify these correlations in terms of the value of the outcome y over the confusion matrix of $(v, \hat{v})$ conditioned on $\ell = 0, 1$ that we define in Table 1 for group $\ell = l$. For the sake of brevity we will overload notation and use the term confusion matrix to refer to the Table above. Constraints on the confusion matrix will be captured by the variable ε .

4 SIMPLE CASES

Before proceeding to our main analysis in the next sections, we state a few simple scenarios where accurate bias estimation via $\hat{v}$ is possible. The proofs from this section can be found in the Appendix.

4.1 Case 1: $y \perp \{v, \hat{v}\}|\ell$.

This condition implies that y is independent of the covariate (v) and the proxy ($\hat{v}$) when conditioned on ℓ . If the variables y, v and $\hat{v}$ are viewed as nodes in a graphical model [33], then the above assumption encodes the graph structure in Figure 1. In this case we

Figure 1: A graphical model over three random variables under the assumption that y is independent of $\hat{v}$ and v . All the events are conditioned on ℓ .

prove the following theorem:

THEOREM 4.1. *If the independence assumption encoded in Figure 1 holds, then*

$$G = \hat{G}$$

The theorem is intuitive as the independence assumption makes conditioning on either v or $\hat{v}$ irrelevant. A scenario where the assumption will hold is when the outcome y is completely determined by the value of ℓ , the sensitive attribute. This is typically too strong a condition to hold in practice.

4.2 Case 2: $y \perp \hat{v}|v, \ell$.

Next we consider our first non-trivial case where y is independent of the proxy $\hat{v}$ when conditioned on ℓ and v . This assumption is encoded in Figure 2.

Figure 2: A graphical model over three random variables under the assumption that y is independent of $\hat{v}$ conditioned on v . All the events are conditioned on ℓ .

We will show that if the proxy $\hat{v}$ has high conditional precision p_l , then we will get good estimates via $\hat{G}$.

THEOREM 4.2. *If the independence assumption encoded in Figure 2 holds, and for any $\ell \in \{0, 1\}$ the conditional precision of the proxy $\hat{v}$ is at least $1 - p$, then we have that*

$$|G - \hat{G}| \leq 2p.$$

It is interesting to note that the bound does not depend on the recall of the proxy $\hat{v}$. To get an intuition behind this, notice that the errors in the estimate of G involve both *false positives* ($v = 0, \hat{v} = 1$) and *false negatives* ($v = 1, \hat{v} = 0$). Indeed, we can write each term as:

$$\begin{aligned} \mathbb{P}[y=1|v=1, \ell=0] &= (1 - r_0) \mathbb{P}[y=1|\hat{v}=1, v=1, \ell=0] \\ &\quad + r_0 \mathbb{P}[y=1|\hat{v}=0, v=1, \ell=0] \end{aligned} \quad (8)$$

$$\begin{aligned} \mathbb{P}[y=1|\hat{v}=1, \ell=0] &= (1 - p_0) \mathbb{P}[y=1|v=1, \hat{v}=1, \ell=0] \\ &\quad + p_0 \mathbb{P}[y=1|v=0, \hat{v}=1, \ell=0]. \end{aligned} \quad (9)$$

The estimation error will be small if the two terms are close (similarly for $\ell = 1$).

Due to the independence assumption, we have that $\mathbb{P}[y=1|v=1, \hat{v}=0, \ell=0] = \mathbb{P}[y=1|v=1, \hat{v}=1, \ell=0]$. Therefore the first equation simply becomes $\mathbb{P}[y=1|\hat{v}=1, v=1, \ell=0]$ and the dependency on r_0 disappears. Only false positives ($v = 0, \hat{v} = 1$) dictate the errors which is controlled by the precision of the proxy $\hat{v}$.

One may expect the above condition to hold if the classifier for $\hat{v}$ is trained on a different independent data set and hence its errors

may not be correlated with y . As an extreme case, the condition would hold if for any data point, $\hat{v}$ has a $1 - p$ chance of being correct, independently of other points.

4.3 Case 3: $y \perp v | \hat{v}, \ell$.

Finally, the third case that we consider in this section is the complementary case where y is independent of v when conditioned on ℓ and $\hat{v}$. This is encoded in Figure 3. In this case we will show that

Figure 3: A graphical model over three random variables under the assumption that y is independent of v conditioned on $\hat{v}$. All the events are conditioned on ℓ .

if the proxy $\hat{v}$ has high recall ($\mathbb{P}[\hat{v} = 1 | v = 1, \ell]$) then we will get good estimates via $\hat{G}$.

THEOREM 4.3. *If the independence assumption encoded in Figure 3 holds, and for any $\ell \in \{0, 1\}$ the conditional recall of the proxy $\hat{v}$ is at least $1 - r$, then we have that:*

$$|G - \hat{G}| \leq 2r.$$

Again, the bound above depends only on the recall and the intuition is similar to that of Theorem 4.2. Analyzing again the equations (8) and (9), we can see that (9) is simply equal to $\mathbb{P}[y = 1 | \hat{v} = 1, v = 1, \ell = 0]$, with p_0 disappearing. Therefore only false negatives matter in this case which leads to the constraint on the recall.

Similar to the previous case, one may expect the above condition to hold if the errors of $\hat{v}$ are somewhat random. For example, the condition will hold if every prediction of $\hat{v}$ has a $1 - r$ chance of being correct, independently of other data points.

4.4 Key take-aways.

We presented in this section some simple cases relying on strong independence assumptions where the estimation error from using $\hat{v}$ can be bounded. The goal was to get the reader familiar with the notations but also to illustrate how the estimation error might be affected by two types of system parameters: the performance of the covariate proxy $\hat{v}$ (precision, recall) as well as the general correlations between parameters ($y, v, \hat{v}, \ell$). In the next sections, we present some natural conditions that do not rely on strong assumptions and yet guarantee accurate bias estimation when using $\hat{v}$. In Section 5, we will first bound the estimation error based only on the performance of the classifier $\hat{v}$. Then in Section 6, we will highlight two additional characteristics of the system in terms of ($y, v, \hat{v}, \ell$) that lead to smaller estimation errors. Finally in Section 7, we derive a refined bound that combines all the different considerations.

5 BOUNDS BASED ONLY ON CLASSIFIER PERFORMANCE

Initially, a practitioner might be willing to only make some assumptions on the performance of the proxy $\hat{v}$. In this section, we will study this case and derive the associated bounds.

Main equations. We first write some important equations that will be useful throughout the rest of the paper. We define

$$\delta_0 = \mathbb{P}[y = 1 | v = 1, \ell = 0] - \mathbb{P}[y = 1 | \hat{v} = 1, \ell = 0] \quad (10)$$

From Equation (8) and (9) we get that

$$\begin{aligned} \delta_0 &= (p_0 - r_0) \cdot (\mathbb{P}[y = 1 | \hat{v} = 1, v = 1, \ell = 0]) \\ &\quad + r_0 \mathbb{P}[y = 1 | \hat{v} = 0, v = 1, \ell = 0] \\ &\quad - p_0 \mathbb{P}[y = 1 | \hat{v} = 1, v = 0, \ell = 0] \end{aligned} \quad (11)$$

In a similar manner we can define δ_1 that is conditioned on the slice $\ell = 1$. Eq. (11) represents the estimation error made on the slice $\ell = 0$ and has a natural interpretation. The second term is the error generated by not including *false negative examples* ($\hat{v} = 0, v = 1$) into the $\hat{G}$ term, while the third term is the error caused by including incorrectly, the *false positive examples* ($\hat{v} = 1, v = 0$). The final estimation error is the difference between these two terms:

$$|G - \hat{G}| = |\delta_1 - \delta_0| \quad (12)$$

It is easy to see that the magnitude of each term in δ_0 and δ_1 depends on the performance of the proxy $\hat{v}$. We next investigate this dependence.

5.1 Case A

The first condition that we study is when the proxy $\hat{v}$ has both high conditional precision and conditional recall. In this case we can bound the error in the estimates without any additional assumptions.

THEOREM 5.1. *If for any $\ell \in \{0, 1\}$ the precision and recall of the proxy $\hat{v}$ is at least $1 - \gamma_A$, then we have that*

$$|G - \hat{G}| \leq 2 \cdot \gamma_A.$$

PROOF. Below, we will bound δ_0 and the analysis for δ_1 will be identical.

We can write δ_0 as:

$$\begin{aligned} \delta_0 &= \mathbb{P}[y = 1 | v = 1, \ell = 0] - \mathbb{P}[y = 1 | \hat{v} = 1, \ell = 0] \\ &= (p_0 - r_0) \mathbb{P}[y = 1 | v = 1, \hat{v} = 1, \ell = 0] \\ &\quad + r_0 \mathbb{P}[y = 1 | v = 1, \hat{v} = 0, \ell = 0] \\ &\quad - p_0 \mathbb{P}[y = 1 | v = 0, \hat{v} = 1, \ell = 0]. \end{aligned} \quad (13)$$

Since the third term above is negative, and the probabilities are bounded by 1, we can conclude that δ_0 is at most p_0 . Similarly, the negation, i.e., $-\delta_0$ is bounded by r_0 . Hence we get that

$$\begin{aligned} \left| \mathbb{P}[y = 1 | v = 1, \ell = 0] - \mathbb{P}[y = 1 | \hat{v} = 1, \ell = 0] \right| &\leq \max(p_0, r_0) \\ &\leq \gamma_A \end{aligned}$$

□

Note that precision and recall have to be relatively high so that the bound above can be small enough. For instance, even if the proxy has precision and recall around 0.9, the above bound will be around 0.2, which would likely be too large for fairness evaluation in practical settings. The goal of the next two sections will be to progressively refine this bound by analyzing potential correlations of $(y, v, \hat{v}, \ell)$.

6 ALTERNATE BOUNDS BASED ON OTHER CORRELATIONS.

In many cases requiring high precision and recall on the proxy may not be achievable. Can one still expect the estimation error to be small? We next characterize other conditions based on correlations between the different parameters $(y, v, \hat{v}, \ell)$ that enable us to bound the error. We will study cases B1 and B2, that identify two different canceling effects under which the error might be small, despite not having high precision and recall. In order to understand these it would be beneficial to refer to the confusion matrix as shown in Table 1.

6.1 Intuition for the cases.

Recall δ_ℓ defined in Equation (11) as

$$\begin{aligned} \delta_\ell = & (p_\ell - r_\ell) \cdot (\mathbb{P}[y=1|\hat{v}=1, v=1, \ell=\ell]) \\ & + r_\ell \mathbb{P}[y=1|\hat{v}=0, v=1, \ell=\ell] \\ & - p_\ell \mathbb{P}[y=1|\hat{v}=1, v=0, \ell=\ell] \end{aligned}$$

This equation above represents the estimation error made on the slice ℓ . Additionally, the final estimation error is equal to the difference between the errors made on each slice:

$$|G - \hat{G}| = |\delta_1 - \delta_0|$$

In the previous section, we bounded each term in δ_0 and δ_1 by relying on the performance of the proxy $\hat{v}$ (case A). Below we discuss alternate natural conditions that would suffice for good gap estimation:

- B1. This case captures scenarios when the second and third term in δ_0 (similarly for δ_1) approximately cancel each other, indicating that the contribution of examples wrongly ignored (false negative) is approximately equal to the examples wrongly included (false positive). Additionally, the first term is small.
- B2. This case captures scenarios when δ_0 and δ_1 have similar values and therefore “cancel each other,” indicating that the correlations (between $y, \hat{v}, v$) have a similar behavior for $\ell = 0$ and $\ell = 1$.

In the rest of the section, we will prove theorems that materialize the above intuition rigorously. We will decouple the effect of the performance of the proxy (precision and recall) on the estimation error from the properties of the confusion matrix.

6.2 Case B1: Precision close to recall and closeness of diagonals

In this case we do not assume that the proxy has high precision and recall, but instead simply assume that the precisions and recalls (conditioned on ℓ) are close to each other, up to γ_B . Clearly

this assumption by itself is not enough to guarantee a small error. We also assume that the diagonal entries in the confusion matrix of Table 1 are ε_B -close to each other. This is made precise in the definition below.

Definition 6.1 (Closeness of Diagonals). There exists a small $0 < \varepsilon_{B1} \leq 1$ such that

$$\begin{aligned} & \left| \mathbb{P}[y=1|v=1, \hat{v}=0, \ell=0] - \mathbb{P}[y=1|v=0, \hat{v}=1, \ell=0] \right| \\ & \leq \varepsilon_{B1} \\ & \left| \mathbb{P}[y=1|v=1, \hat{v}=0, \ell=1] - \mathbb{P}[y=1|v=0, \hat{v}=1, \ell=1] \right| \\ & \leq \varepsilon_{B1} \end{aligned}$$

One can expect the above condition to hold in certain scenarios. Notice that for any ℓ , the condition concerns the disagreement region, i.e., $\{v=1, \hat{v}=0\}$, and $\{v=0, \hat{v}=1\}$. In many cases this region could represent “hard to predict” points. The closeness on the diagonal is that the outcome y has similar behavior for both errors types of $\hat{v}$. Returning to the example from Section 1 on predicting toxicity of online comments with the topic being the covariate (v), the classifier $\hat{v}$ may find it hard to predict certain esoteric topics well. It is then natural to expect that the toxicity prediction model may also perform poorly for these topics and its errors may be somewhat random, thereby making ε_{B1} small. There may also be many cases where one can indeed not expect ε_{B1} to be small. However, we will soon see in Section 7 a more refined bound that has a lower order dependence on such errors. Hence, even a weak constraint (on ε_{B1}) suffices for getting smaller bounds.

Under the closeness of diagonal condition we have the following theorem.

THEOREM 6.2. *Let $\varepsilon_{B1} \in [0, 1]$ be the smallest constant such that the closeness of diagonal condition holds with ε_{B1} and that $|r_0 - p_0|, |r_1 - p_1|$ are bounded by γ_{B1} then we have that*

$$|G - \hat{G}| \leq 2(\gamma_{B1} + \varepsilon_{B1}).$$

As Compared to Theorem 5.1, case B1 does not require an absolute bound on the precisions and recalls and quantifies the error in terms of their relative closeness (γ_{B1}). This can be advantageous in some cases; the quantities captured by γ_{B1} might be smaller than γ_A .

PROOF. As before we will establish bounds for both δ_0 and δ_1 . Let's consider δ_0 . From the analysis in the previous section we have that

$$\begin{aligned} & \mathbb{P}[y=1|v=1, \ell=0] - \mathbb{P}[y=1|\hat{v}=1, \ell=0] \\ & = (p_0 - r_0) \mathbb{P}[y=1|v=1, \hat{v}=1, \ell=0] \\ & \quad + r_0 \mathbb{P}[y=1|v=1, \hat{v}=0, \ell=0] \\ & \quad - p_0 \mathbb{P}[y=1|v=0, \hat{v}=1, \ell=0] \\ & \leq (p_0 - r_0) \mathbb{P}[y=1|v=1, \hat{v}=1, \ell=0] \\ & \quad + r_0 \mathbb{P}[y=1|v=1, \hat{v}=0, \ell=0] \\ & \quad - p_0 r_0 \mathbb{P}[y=1|v=1, \hat{v}=0, \ell=0] + \varepsilon r_0 \\ & \leq |p_0 - r_0| + \varepsilon r_0 \leq \gamma_{B1} + \varepsilon_{B1}. \end{aligned}$$

A similar analysis establishes that $\delta_1 \leq \gamma_{B1} + \varepsilon_{B1}$. $\square$

6.3 Case B2: Closeness of precisions/recalls and closeness of models

The previous case identified conditions under which both δ_0 and δ_1 achieve small values, leading to a small estimation error. This error depends on the closeness of precisions and recalls within each slice. However, in practice this may not always hold and one may prefer an alternate condition. We next identify another natural condition that allows for precision and recall within each slice to be arbitrary, and both δ_0 and δ_1 could be large. However, one may still expect the estimation error ($|G - \hat{G}|$) to be small if the errors in δ_0 and δ_1 have a cancelling effect. We present such a condition next.

Definition 6.3 (Model Closeness). There exists $g > 0$ and a $0 < \varepsilon_{B2} \leq 1$ such that for all $b, c \in \{0, 1\}$,

$$\begin{aligned} & \mathbb{P}[y=1|v=b, \hat{v}=c, \ell=1] \\ &= \mathbb{P}[y=1|v=b, \hat{v}=c, \ell=0] + g \pm \varepsilon_{B2}. \end{aligned}$$

In terms of the confusion matrix defined in Table 1 the above condition requires that the entries of the confusion matrix for $\ell = 1$ are noisy translations of the entries for $\ell = 0$, g representing here the gap between groups.

Notice that case B2 allows for arbitrary behavior within each slice (ℓ). Hence, for the estimation error to be small, some assumption is needed on the correlations across slices and the model closeness condition in Definition 6.3 captures that. In practice one may expect the condition to hold if the proxy $\hat{v}$ introduces similar errors across slices. As an example, in the context of the toxicity prediction example, consider a proxy ($\hat{v}$) that was (unintentionally) trained on a corpus consisting of only one group, say $\ell = 1$. When applying $\hat{v}$ to the group $\ell = 0$ one can expect $\hat{v}$ to introduce systematic errors (captured by g in the definition).

In this case we have the following theorem

THEOREM 6.4. Let $\gamma_{B2}, \varepsilon_{B2} \in [0, 1]$ and g be such that the model closeness holds with ε_{B2} and that $|r_0 - r_1|, |p_0 - p_1|$ are bounded by γ_{B2} . Then we have that

$$|G - \hat{G}| \leq 2 \cdot \gamma_{B2} + 3 \cdot \varepsilon_{B2}.$$

Note that in the above Theorem we do not require precisions and recalls within any slice to be high.

PROOF. We would not be able to separately analyze δ_0 and δ_1 since both could be high. Instead we argue about the difference. We have

$$\delta_0 - \delta_1 = T_1 + T_2 + T_3, \quad (14)$$

where

$$\begin{aligned} T_1 &= (p_0 - r_0) \mathbb{P}[y=1|v=1, \hat{v}=1, \ell=0] \\ &\quad - (p_1 - r_1) \mathbb{P}[y=1|v=1, \hat{v}=1, \ell=1], \\ T_2 &= r_0 \mathbb{P}[y=1|v=1, \hat{v}=0, \ell=0] \\ &\quad - r_1 \mathbb{P}[y=1|v=1, \hat{v}=0, \ell=1], \end{aligned}$$

and

$$\begin{aligned} T_3 &= p_1 \mathbb{P}[y=1|v=0, \hat{v}=1, \ell=1] \\ &\quad - p_0 \mathbb{P}[y=1|v=0, \hat{v}=1, \ell=0]. \end{aligned}$$

Since $|r_0 - r_1|, |p_0 - p_1|$ are both bounded by γ_{B2} and model closeness holds we have

$$\begin{aligned} T_1 &= (p_0 - r_0 - p_1 + r_1) \mathbb{P}[y=1|v=1, \hat{v}=1, \ell=0] \\ &\quad - (p_1 - r_1)g \pm \varepsilon_{B2}, \end{aligned}$$

$$\begin{aligned} T_2 &= (r_0 - r_1) \mathbb{P}[y=1|v=1, \hat{v}=0, \ell=0] \\ &\quad - r_1g \pm \varepsilon_{B2}, \end{aligned}$$

and

$$\begin{aligned} T_3 &= (p_1 - p_0) \mathbb{P}[y=1|v=0, \hat{v}=1, \ell=0] \\ &\quad + p_1g \pm \varepsilon_{B2} \end{aligned}$$

Adding the above we get that $T_1 + T_2 + T_3 \leq 2 \cdot \gamma_{B2} + 3 \cdot \varepsilon_{B2}$. Similarly, the negation $\delta_1 - \delta_0$ can be bounded. $\square$

7 SYNTHESIS: A REFINED BOUND.

So far we have presented three natural conditions under which the error in the estimation can be bounded. We recap them below as:

- A. The first bound uses only the performance of the proxy $\hat{v}$, therefore needs high precision and recall.
- B1. The second bound depends on the precision being close to the recall, in addition to the closeness of diagonals.
- B2. The third bound depends on the closeness of precision and recalls between slices and the relationship between the confusion matrices of the two slices.

The bounds earlier concern each of the above conditions in isolation. As a result, they are likely to be wide except unless one of the constraint is very strong, which might not be realistic. In this section we study whether one can combine the power of the three conditions to refine the error bound that naturally adapts to the extent to which the above conditions are satisfied. Towards that end we have the following theorem.

THEOREM 7.1. Let $\gamma_A, (\gamma_{B1}, \varepsilon_{B1})$ and $(\gamma_{B2}, \varepsilon_{B2})$ be the errors up to which conditions A, B1 and B2 above hold. Then we have that

$$\begin{aligned} |G - \hat{G}| &\leq 2 \min(\gamma_A, \gamma_{B1}, \gamma_{B2}) \\ &\quad + \varepsilon_{B2} \cdot (2\gamma_A + \gamma_{B1}) + \varepsilon_{B1} \cdot \gamma_{B1} \end{aligned}$$

This theorem combines the three conditions highlighted before and achieves a refined bound. In particular, the dependence on the parameters governing the precision/recall of the classifier (γ) is linear, whereas in contrast to earlier theorems, the errors due to other correlations (ε) have a multiplying effect. Hence, in practice even a weak estimate of the parameters ε_{B1} and ε_{B2} may suffice to get an estimate of how good $\hat{G}$ is.

PROOF. Recalling the definitions of T_1, T_2 and T_3 from the proof of Theorem 6.4 we get that

$$\begin{aligned} T_1 &= (p_0 - r_0 - p_1 + r_1) \mathbb{P}[y=1|v=1, \hat{v}=1, \ell=0] \\ &\quad - (p_1 - r_1)g \pm \varepsilon_{B2} \cdot (p_1 - r_1), \end{aligned}$$

$$\begin{aligned} T_2 &= (r_0 - r_1) \mathbb{P}[y=1|v=1, \hat{v}=0, \ell=0] \\ &\quad - r_1g \pm \varepsilon_{B2} \cdot r_1, \end{aligned}$$

and

$$T_3 = (p_1 - p_0) \mathbb{P}[y=1|v=0, \hat{v}=1, \ell=0] \\ + p_1 g \pm \varepsilon_{B2} \cdot p_1.$$

Adding up we get that

$$\begin{aligned} T_1 + T_2 + T_3 &\leq (p_0 - r_0 - p_1 + r_1) \mathbb{P}[y=1|v=1, \hat{v}=1, \ell=0] \\ &\quad + (r_0 - r_1) \mathbb{P}[y=1|v=1, \hat{v}=0, \ell=0] \\ &\quad + (p_1 - p_0) \mathbb{P}[y=1|v=0, \hat{v}=1, \ell=0] \\ &\quad + \varepsilon_{B2} \cdot (2\gamma_A + \gamma_{B1}) \\ &\leq (p_0 - r_0 - p_1 + r_1) \mathbb{P}[y=1|v=1, \hat{v}=1, \ell=0] \\ &\quad + (r_0 - r_1) \mathbb{P}[y=1|v=1, \hat{v}=0, \ell=0] \\ &\quad + (p_1 - p_0) \mathbb{P}[y=1|v=1, \hat{v}=0, \ell=0] + \varepsilon_{B2}(2\gamma_A + \gamma_{B1}) \\ &\quad + \gamma_{B1} \cdot \varepsilon_{B1} \\ &\leq (p_0 - r_0 - p_1 + r_1) + \varepsilon_{B2} \cdot (2\gamma_A + \gamma_{B1}) \\ &\quad + \gamma_{B1} \cdot \varepsilon_{B1}. \end{aligned}$$

The bound follows from noticing that

$$p_0 - r_0 - p_1 + r_1 \leq 2 \min(\gamma_A, \gamma_{B1}, \gamma_{B2}).$$

□

Compared to Theorem 5.1, the first term in the bound above may be significantly smaller than $2\gamma_A$. Furthermore, in contrast to Theorem 6.2 and Theorem 6.4, the dependence on ε_{B1} and ε_{B2} is diminished due to multiplication by other error terms. Due to this cancelling effect and the linear dependency on the parameters related to the proxy performance (γ), our theory recommends that the focus of a practitioner should then be on minimizing the first term. We next validate our theory via experiments on simulated data.

8 SIMULATIONS

We proposed in Section 5 an initial bound based solely on the classifier performance and then highlighted that some constraints over the confusion matrix (captured by ε) lead to errors canceling each other.

In this section, we use simulations to further validate our theory. The findings of this section are summarized below.

- We will observe that our theoretical bounds from Section 5.1 that are based solely on the classifier performance represent worst-case scenarios, and that the gap in most cases is significantly smaller than our upper bounds. In particular, the 95th percentile of the distribution of errors is typically more than $2\times$ smaller.
- We will demonstrate the importance of making *some* assumptions on the error parameters via ε , instead of simply relying on classifier performance via γ . The constraint on the parameter values (ε) rules out some of these worst-case scenarios.
- Finally, we will show that even weak constraints over the parameters (with ε) are sufficient to significantly reduce the bound.

Figure 4: Simulated distribution of estimation errors. Each graph represent a different choice of the precision and recall of the proxy. The dashed blue line represent 95th-percentiles and the unbroken red one is the theoretical bound from Theorem 5.1.

8.1 Simulated estimation errors are often smaller than the theoretical bounds.

In order to compute the estimation error, we rely on Equation (13) where we need to define both the classifier performance and the confusion matrix from Table 1. We first illustrate the scenario where we only control for classifier performance (case A) and therefore do not enforce any constraints on the values in the confusion matrix from Table 1.

Setup: We first set classifier-related quantities (p_1, r_1, p_0 and r_0). Note that we display results for a few choice of these parameters, however the observations hold across any choices made for those quantities.

Once we set the above quantities, we generate the confusion matrix, aka the six parameters, $\mathbb{P}[y=1|v=i, \hat{v}=j, \ell=l]$, for $i, j, l \in \{0, 1\}$, by sampling from a uniform prior in $[0, 1]$, $U[0, 1]$. This reflects the situation where no prior structure is made on the confusion matrix. Note that $\mathbb{P}[y=1|v=0, \hat{v}=0, \ell=l]$ are not needed.

Results: The results are displayed in Figure 4. First, we confirm that the bound provided by Theorem 5.1 is correct. More importantly, we observe that these bounds reflect worst-case scenarios and that the errors are typically significantly smaller: indeed, the 95th percentile is approximately $2\times$ smaller than the upper bound. This phenomena is due to the fact that the different terms in Equation (13) cancel each other leading to significantly lower errors than worst-case scenario.

This above setup uses as a uniform prior over parameters and can lead to unrealistic scenarios that might be responsible for worst-case estimation errors. In the next sections will focus on more realistic scenarios by restricting the range of values of the parameters following the principles highlighted by our theory. This essentially constrains the possible values that the confusion matrix in Table 1 can take.

8.2 Importance of assuming structure over the confusion matrix via ε .

In this section, we assume that we have a classifier with relatively good precision and recall (for instance $r_i, p_i \leq 0.1$). The bound from 5.1 is still too large for a good estimate of the error (in this case close to 0.2), and we show that assuming some structure over the correlations is key to getting better estimates.

Setup: To simulate more realistic scenarios, we now enforce some level of structure controlled by ε . As before, we first select a well performing classifier ($p_1 = 0.07$, $r_1 = 0.09$, $p_0 = 0.05$ and $r_0 = 0.1$). We will report experiments only for this choice but the results are consistent across different choices. Then, we set a value ε_{B1} and ε_{B2} , which reflects to what extent the constraints B1 (closeness of diagonals) and B2 (model closeness) hold. For instance, $\varepsilon_{B1} = 0$ means equal parameters and $\varepsilon_{B1} = 1$ is equivalent to having no structure.

We sample the parameters as follows:

- $g \sim U[-1, 1]$, where U denotes the uniform distribution.
- $\mathbb{P}[y = 1|v = 1, \hat{v} = 1, \ell = l] \sim U[0, 1 - g]$, if $g \geq 0$, else $\sim U[g, 1]$.
- $\mathbb{P}[y = 1|v = 1, \hat{v} = 0, \ell = l] \sim U[0, 1 - g]$ if $g \geq 0$, else $\sim U[g, 1]$.
- $\mathbb{P}[y = 1|v = 0, \hat{v} = 1, \ell = l] = \mathbb{P}[y = 1|v = 1, \hat{v} = 0, \ell = l] + U[-\varepsilon_{B2}, \varepsilon_{B2}]$.
- $\mathbb{P}[y = 1|v = i, \hat{v} = j, \ell = l] = \mathbb{P}[y = 1|v = i, \hat{v} = j, \ell = l] + g + U[-\varepsilon_{B2}, +\varepsilon_{B2}]$.

Note that the above sampling might not lead to valid probabilities (outside of $[0, 1]$). If that happens, we restart again and sample all the parameters. This method ensures sampling the confusion matrix in Table 1 that satisfies our constraints.

We compute both the true and estimated gap (G and $\hat{G}$) and report the distribution of errors ($error = |G - \hat{G}|$) for $N = 100000$ independent runs.

Figure 5: Refined bound from Theorem 7.1 dominates other bounds. Each graph represent a different choice of $(\varepsilon_{B1}, \varepsilon_{B2})$. The dashed blue line represents the 95th percentile of the simulated errors. The red (solid) line comes from Theorem 5.1. The pink (resp. purple) line comes from Theorem 6.2 (resp Theorem 6.4). Finally, the green line is given by the refined bound from Theorem 7.1.

Results: We report in Figure 5 the histogram of the errors, as well as the various theoretical bounds. The dashed blue line represents the 95th percentile of the simulated errors. The red (solid) line comes from Theorem 5.1 and represents the case when the practitioner does not make any assumption on the structure of the confusion matrix. The pink (resp. purple) line comes from Theorem 6.2 (resp Theorem 6.4) which focuses on one constraint on the confusion matrix. Finally, the green line is given by the refined bound from Theorem 7.1.

We observe that the bounds from Case A, B1 or B2 are all quite wide (typically above 0.2 in this scenario). Note that their ordering in the plots may differ depending on how we set each constraint. More importantly, the bound from Theorem 7.1 is significantly better than any of the individual ones thereby illustrating that

by considering all the different sources of error, we are able to compute better estimates. In particular, it is important to assume some structure over the confusion matrix, through ε , in order to prevent worst-case scenarios.

8.3 Even weak assumptions on ε are enough to significantly reduce the bound.

The previous experiment highlights the importance of adding some constraints on potential values of the confusion matrix from Table 1. However it might be hard to assume strong conditions (i.e., ε close to 0). In this section, we demonstrate that even loose estimates are enough to significantly improve our estimates.

Figure 6: On the left (resp. right) plot, we set ε_{B2} (resp. ε_{B1}) and vary ε_{B1} (resp. ε_{B2}). The blue lines represent the 95th-percentile of the estimation errors. The refined bound from Theorem 7.1 (dashed orange) dominates the one from Theorem 5.1 (dotted red), even with weak constraints (ε does not have to be small).

Setup: To that purpose, we study how the estimation errors evolve as we progressively strengthen the condition on the correlations. More precisely, we take a good classifier with $p_1 = 0.07$, $r_1 = 0.09$, $p_0 = 0.05$ and $r_0 = 0.1$ (arbitrary values for sake of simplicity). We then set one of $(\varepsilon_{B1}, \varepsilon_{B2})$ to some fixed value (here 0.2) and vary the other one from 0 (very strong condition) to 1 (no condition). Then we use the same approach as in the previous experiment to sample all the remaining variables. We finally analyze the resulting estimation errors, as well as theoretical bounds.

Results: Results are displayed in Figure 6. The blue lines represent the 95th-percentile of the estimation errors from the simulation, the dashed orange line is the theoretical bound from Theorem 7.1 and the dotted red one from case A, i.e., Theorem 5.1.

We can see that, even in this case where the covariate classifier is relatively good ($r_i, p_i \leq 0.1$), the bound from Theorem 5.1 remains quite large (0.2). However the refined bound in Theorem 7.1 leads to a tighter estimate (except in the case when ε_{B2} becomes close to 1 in the right plot). Furthermore, both the theoretical guarantees and empirical results significantly improve when the enforced conditions become quantitatively stronger. More importantly, these graphs demonstrate that Theorem 7.1 does not require strong constraints to achieve a satisfying bound. For instance, the left graph shows that the theoretical bound is still lower than 0.1 if $\varepsilon_{B2} \leq 0.4$ (a relatively weak condition). Additionally the 95th percentile is between .02 and .03, which shows even better estimate in practice.

This demonstrates that *some* assumptions on ε , even if weak, is key in getting good estimates of the error.

9 CONCLUSION

Measuring the fairness of a model might often be restricted by the data available. We analyze in this paper the common scenario of estimating a fairness metric via a proxy covariate. We identify various sources of errors that affect the use of the proxy covariate in measuring model fairness. Our theory demonstrates that the errors can be driven down not only by the performance of the proxy such as precision and recall, but also by the underlying correlations in the data distribution. We provided a refined analysis that combines the various sources of errors and highlights how these errors interact to affect the final bound. More importantly, our work shows that one does not need to make strong independence assumptions in order to study and obtain bounds for the problem of measuring model fairness via proxy covariates.

Finally we demonstrate through simulations that exploiting structure in the correlations, even if loose, is the key to obtaining better bounds on the error in measuring the fairness metrics. We believe that our work will guide practitioners towards making informed choices when designing proxy covariates for fairness evaluation, or obtaining bounds on the errors in their estimation procedures.

10 ACKNOWLEDGEMENTS:

The authors would like to thank Alexander D'Amour for his valuable feedback on this paper.

A PROOFS FROM SECTION 4

PROOF OF THEOREM 4.1. Define,

$$\begin{aligned}\delta_0 &:= \mathbb{P}[y=1|v=1, \ell=0] - \mathbb{P}[y=1|\hat{v}=1, \ell=0] \\ \delta_1 &:= \mathbb{P}[y=1|v=1, \ell=1] - \mathbb{P}[y=1|\hat{v}=1, \ell=1].\end{aligned}$$

Under the independence assumption it is easy to see that

$$\mathbb{P}[y=1|v=1, \ell=0] = \mathbb{P}[y=1|\hat{v}=1, \ell=0].$$

This implies that $\delta_0 = 0$, and similarly that $\delta_1 = 0$. Hence in this case the error in our estimates is zero. $\square$

PROOF OF THEOREM 4.2. We will show that both δ_0 and δ_1 are bounded in magnitude by at most p thereby establishing the theorem. We next show how to bound δ_0 . The analysis for δ_1 will be identical. Recall that

$$\delta_0 = \mathbb{P}[y=1|v=1, \ell=0] - \mathbb{P}[y=1|\hat{v}=1, \ell=0].$$

Under the independence assumption we have

$$\begin{aligned}\delta_0 &= \mathbb{P}[y=1|v=1, \ell=0] \\ &\quad - \mathbb{P}[y=1|v=1, \ell=0] \mathbb{P}[v=1|\hat{v}=1, \ell=0] \\ &\quad - \mathbb{P}[y=1|v=0, \ell=0] \mathbb{P}[v=0|\hat{v}=1, \ell=0] \\ &\leq \mathbb{P}[y=1|v=1, \ell=0] \\ &\quad - \mathbb{P}[y=1|v=1, \ell=0] \mathbb{P}[v=1|\hat{v}=1, \ell=0] \\ &= \mathbb{P}[y=1|v=1, \ell=0] (1 - \mathbb{P}[v=1|\hat{v}=1, \ell=0]) \\ &\leq r\end{aligned}$$

Similarly, we get that

$$\begin{aligned}-\delta_0 &= \mathbb{P}[y=1|v=1, \ell=0] \mathbb{P}[v=1|\hat{v}=1, \ell=0] \\ &\quad + \mathbb{P}[y=1|v=0, \ell=0] \mathbb{P}[v=0|\hat{v}=1, \ell=0] \\ &\quad - \mathbb{P}[y=1|v=1, \ell=0] \\ &= \mathbb{P}[y=1|v=1, \ell=0] (\mathbb{P}[v=1|\hat{v}=1, \ell=0] - 1) \\ &\quad + \mathbb{P}[y=1|v=0, \ell=0] \mathbb{P}[v=0|\hat{v}=1, \ell=0] \\ &= \mathbb{P}[y=1|v=1, \ell=0] \mathbb{P}[v=0|\hat{v}=1, \ell=0] \\ &\quad + \mathbb{P}[y=1|v=0, \ell=0] \mathbb{P}[v=0|\hat{v}=1, \ell=0] \\ &= \mathbb{P}[v=0|\hat{v}=1, \ell=0] \\ &\leq r.\end{aligned}$$

$\square$

PROOF OF THEOREM 4.3. The proof is exactly identical to the proof of Theorem 4.2 above by simply switching the role of v and $\hat{v}$. $\square$

REFERENCES

- [1] Leman Akoglu, Mary McGlohon, and Christos Faloutsos. 2010. Oddball: Spotting anomalies in weighted graphs. In *Pacific-Asia conference on knowledge discovery and data mining*. Springer, 410–421.
- [2] P. Awasthi, A. Beutel, M. Kleindessner, J. Morgenstern, and X. Wang. 2021. Evaluating Fairness of Machine Learning Models Under Uncertain and Incomplete Information. In *ACM Conference on Fairness, Accountability, and Transparency (FAccT)*.
- [3] S. Barocas, M. Hardt, and A. Narayanan. 2018. *Fairness and Machine Learning*. fairmlbook.org. <http://www.fairmlbook.org>.
- [4] Alex Beutel, Jilin Chen, Tulsee Doshi, Hai Qian, Allison Woodruff, Christine Luu, Pierre Kreitmann, Jonathan Bischof, and Ed H Chi. 2019. Putting fairness principles into practice: Challenges, metrics, and improvements. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*. 453–459.
- [5] Jiahao Chen, Nathan Kallus, Xiaojie Mao, Geoffry Svacha, and Madeleine Udell. 2019. Fairness under unawareness: Assessing disparity when protected class is unobserved. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*.
- [6] J. Chen, N. Kallus, X. Mao, G. Svacha, and M. Udell. 2019. Fairness Under Unawareness: Assessing Disparity When Protected Class Is Unobserved. In *Conference on Fairness, Accountability, and Transparency (ACM FAT*)*.
- [7] Alexandra Chouldechova. 2017. Fair prediction with disparate impact: A study of bias in recidivism prediction instruments. *Big data* 5, 2 (2017), 153–163.
- [8] Sam Corbett-Davies, Emma Pierson, Avi Feller, Sharad Goel, and Aziz Huq. 2017. Algorithmic decision making and the cost of fairness. *CoRR* abs/1701.08230 (2017). arXiv:1701.08230 <http://arxiv.org/abs/1701.08230>
- [9] Amanda Coston, Karthikeyan N Ramamurthy, Dennis Wei, Kush R Varshney, Skyler Speakman, Zairah Mustahsan, and Supriyo Chakraborty. 2019. Fair Transfer Learning with Missing Protected Attributes. In *AAAI / ACM Conference on Artificial Intelligence, Ethics, and Society*.
- [10] W. Dieterich, C. Mendoza, and T. Brennan. 2016. *COMPAS Risk Scales: Demonstrating Accuracy Equity and Predictive Parity*. Technical Report. Northpointe Inc. <https://www.equivant.com/response-to-propublica-demonstrating-accuracy-equity-and-predictive-parity/>.
- [11] Lucas Dixon, John Li, Jeffrey Sorensen, Nithum Thain, and Lucy Vasserman. 2018. Measuring and Mitigating Unintended Bias in Text Classification. In *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society* (New Orleans, LA, USA) (AIES '18). ACM, New York, NY, USA, 67–73. <https://doi.org/10.1145/3278721.3278729>
- [12] Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard Zemel. 2012. Fairness through awareness. In *Proceedings of the 3rd innovations in theoretical computer science conference*. 214–226.
- [13] Riccardo Fogliato, Alexandra Chouldechova, and Max G'Sell. 2020. Fairness Evaluation in Presence of Biased Noisy Labels. In *International Conference on Artificial Intelligence and Statistics*. PMLR, 2325–2336.
- [14] Maya Gupta, Andrew Cotter, Mahdi M Fard, and Serena Wang. 2018. Proxy Fairness. (2018). arXiv:1806.11212 [cs.LG].
- [15] Moritz Hardt, Eric Price, and Nathan Srebro. 2016. Equality of opportunity in supervised learning. In *Advances in Neural Information Processing Systems*.
- [16] T. Hashimoto, M. Srivastava, H. Namkoong, and P. Liang. 2018. Fairness Without Demographics in Repeated Loss Minimization. In *International Conference on Machine Learning (ICML)*. Code available on <https://bit.ly/2sFkDpE>.

- [17] Úrsula Hébert-Johnson, Michael P Kim, Omer Reingold, and Guy N Rothblum. 2017. Calibration for the (Computationally-Identifiable) Masses. *CoRR* abs/1711.08513 (2017). *arXiv preprint arXiv:1711.08513* (2017).
- [18] Benjamin Horne and Sibel Adali. 2017. This just in: Fake news packs a lot in title, uses simpler, repetitive content in text body, more similar to satire than real news. In *Proceedings of the International AAAI Conference on Web and Social Media*, Vol. 11.
- [19] Hossein Hosseini, Sreeram Kannan, Baosen Zhang, and Radha Poovendran. 2017. Deceiving Google's Perspective API Built for Detecting Toxic Comments. *CoRR* abs/1702.08138 (2017). *arXiv:1702.08138* <http://arxiv.org/abs/1702.08138>
- [20] Heinrich Jiang and Ofir Nachum. 2020. Identifying and correcting label bias in machine learning. In *International Conference on Artificial Intelligence and Statistics*. PMLR, 702–712.
- [21] Nathan Kallus, Xiaojie Mao, and Angela Zhou. 2020. Assessing Algorithmic Fairness with Unobserved Protected Class Using Data Combination. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*.
- [22] Michael Kearns, S. Neel, Roth A., and Z. S. Wu. 2018. Preventing Fairness Gerrymandering: Auditing and Learning for Subgroup Fairness. In *International Conference on Machine Learning (ICML)*.
- [23] Michael Kearns, Seth Neel, Aaron Roth, and Zhiwei Steven Wu. 2018. Preventing fairness gerrymandering: Auditing and learning for subgroup fairness. In *International Conference on Machine Learning*. PMLR, 2564–2572.
- [24] Niki Kilbertus, Mateo Rojas Carulla, Giambattista Parascandolo, Moritz Hardt, Dominik Janzing, and Bernhard Schölkopf. 2017. Avoiding discrimination through causal reasoning. In *Advances in Neural Information Processing Systems*. 656–666.
- [25] Jon Kleinberg, Sendhil Mullainathan, and Manish Raghavan. 2017. Inherent Trade-Offs in the Fair Determination of Risk Scores. In *Innovations in Theoretical Computer Science Conference (ITCS)*.
- [26] Srijan Kumar and Neil Shah. 2018. False information on web and social media: A survey. *arXiv preprint arXiv:1804.08559* (2018).
- [27] Aditee Ajit Kumthekar, Aniruddh Nath, Ed H. Chi, Ed H. Chi, Jilin Chen, Li Wei, Lichan Hong, Mahesh Sathiamoorthy, Shawn Andrews, Xinyang Yi, and Zhe Zhao. 2019. Recommending What Video to Watch Next: A Multitask Ranking System.
- [28] Preethi Lahoti, Alex Beutel, Jilin Chen, Kang Lee, Flavien Prost, Nithum Thain, Xuezhi Wang, and Ed H. Chi. 2020. Fairness without Demographics through Adversarially Reweighted Learning. *arXiv:2006.13114* [cs.LG]
- [29] Alexandre L Lamy, Ziyuan Zhong, Aditya K Menon, and Nakul Verma. 2019. Noise-tolerant fair classification. In *Advances in Neural Information Processing Systems*.
- [30] Hamdy Mubarak, Kareem Darwish, and Walid Magdy. 2017. Abusive Language Detection on Arabic Social Media. In *Proceedings of the First Workshop on Abusive Language Online*. Association for Computational Linguistics, Vancouver, BC, Canada, 52–56. <https://doi.org/10.18653/v1/W17-3008>
- [31] Ji Ho Park, Jamin Shin, and Pascale Fung. 2018. Reducing Gender Bias in Abusive Language Detection. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, Brussels, Belgium, 2799–2804. <https://doi.org/10.18653/v1/D18-1302>
- [32] Ya'acov Ritov, Yuekai Sun, and Ruofei Zhao. 2017. On conditional parity as a notion of non-discrimination in machine learning. *arXiv preprint arXiv:1706.08519* (2017).
- [33] Martin J Wainwright and Michael Irwin Jordan. 2008. *Graphical models, exponential families, and variational inference*. Now Publishers Inc.
- [34] Michael Wick, Swetasudha Panda, and Jean-Baptiste Tristan. 2019. Unlocking fairness: a trade-off revisited. (2019).
- [35] Muhammad B Zafar, Isabel Valera, Manuel G Rodriguez, and Krishna Gummadi. 2017. Fairness Constraints: Mechanisms for Fair Classification. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*.