

Equality of opportunity in travel behavior prediction with deep neural networks and discrete choice models

Yunhan Zheng
Shenhao Wang*
Jinhua Zhao

Massachusetts Institute of Technology
77 Mass Ave, Cambridge, Massachusetts, U.S.

Abstract

Although researchers increasingly adopt machine learning to model travel behavior, they predominantly focus on prediction accuracy, ignoring the ethical challenges embedded in machine learning algorithms. This study introduces an important missing dimension - computational fairness - to travel behavior analysis. It highlights the accuracy-fairness tradeoff instead of the single dimensional focus on prediction accuracy in the contexts of deep neural network (DNN) and discrete choice models (DCM). We first operationalize computational fairness by *equality of opportunity*, then differentiate between the bias inherent in data and the bias introduced by modeling. The models inheriting the inherent biases can risk perpetuating the existing inequality in the data structure, and the biases in modeling can further exacerbate it. We then demonstrate the prediction disparities in travel behavior modeling using the 2017 National Household Travel Survey (NHTS) and the 2018-2019 My Daily Travel Survey in Chicago. Empirically, DNN and DCM reveal consistent prediction disparities across multiple social groups: both over-predict the false negative rate of frequent driving for the ethnic minorities, the low-income and the disabled populations, and falsely predict a higher travel burden of the socially disadvantaged groups and the rural populations than reality. Comparing DNN with DCM, we find that DNN can outperform DCM in prediction disparities because of DNN's smaller misspecification error. To mitigate prediction disparities, this study introduces an absolute correlation regularization method, which is evaluated with synthetic and real-world data. The results demonstrate the prevalence of prediction disparities in travel behavior modeling, and the disparities still persist regarding a variety of model specifics such as the number of DNN layers, batch size and weight initialization. Since these prediction disparities can exacerbate social inequity if prediction results without fairness adjustment are used for transportation policy making, we advocate for careful consideration of the fairness problem in travel behavior modeling, and the use of bias mitigation algorithms for fair transport decisions.

*Corresponding Author; Email: shenhao@mit.edu

Keywords: Machine learning; deep neural network; travel behavior; discrete choice models; fairness in artificial intelligence

1 Introduction

In recent years, a growing body of literature has adopted machine learning models, particularly deep neural networks (DNNs), to predict travel behavior. The common practice of machine learning is to identify the best model by fitting the training data and being evaluated on the test data, with the objectives of performance optimization and output prediction in various scenarios [1, 2]. Comparing DNN with traditional logit models, previous studies have shown that DNN has higher predictive power and typically makes fewer mistakes in predictions compared to multinomial logit models (MNL) [3, 4]. Specifically, DNN is powerful owing to factors such as the relaxation of linear relationships among variables [5], automatic feature learning [5] and the ability to accommodate various data structures [6, 7].

However, machine learning also poses tremendous ethical challenges. Many studies have found that machine learning models can produce much worse prediction results for disadvantaged groups such as black people, women and low-income populations, leading to unfair treatment of these populations. For example, software based on machine learning algorithms to predict future criminals is biased against people of color [8]. Research on online advertisement showed that ads (e.g. PeopleSmart ads, public records ads) for public records on a person appear disproportionately higher for black-identifying names [9]. Literature concerning text classification demonstrated gender bias in word embeddings trained on Google News, which systematically associates men with computer programmers and women with homemakers [10]. Although the fairness problem in machine learning has been revealed across a large number of contexts, thus far no study has examined the computational fairness issue in travel behavior modeling. In fact, fairness has been an enduring topic in the transportation field [11, 12, 13]. The traditional transport fairness research either adopts a highly qualitative approach or presents quantitative metrics without integrating fairness into algorithms. In this computational era, it is critical to demonstrate the risks of naively adopting models without fairness concerns and showcase the integration of fairness metrics into modeling practices and policy decisions.

Motivated by these research gaps, this research investigates how to measure, evaluate, and mitigate prediction disparities of travel behavior models regarding protected status - race, gender, income, medical condition and region. We take the following three steps. First, we introduce *equality of opportunity* to measure computational fairness in travel behavioral

modeling based on the fairness theory in machine learning research. Second, we identify and measure prediction disparities in travel behavior modeling using binary logistic regression (BLR) and DNN. Third, building upon the approach by Beutel et al. [14], we adopt an absolute correlation regularization method to mitigate the prediction biases, and evaluate the performance of the new model with bias mitigation. The second and the third steps are deployed on both synthetic and real-world datasets. Experiments conducted on the synthetic datasets show how the prediction disparity varies with data structure, number of predictors, sample size and the degree of model misspecification. The fairness analysis is then deployed on the 2017 National Household Travel Survey (NHTS) dataset which has wide coverage of geographic areas and populations with different characteristics, as well as the large sample size and the diversity of input features. It is also deployed on a regional dataset—the 2018-19 Chicago travel survey data. Our analyses show that the prediction disparity issues are prevalent in travel behavior predictions, and our findings are robust to change of hyperparameters such as the number of DNN layers, batch size and weight initialization. Also, our analyses on the nation-wide travel data (NHTS) and the regional data (Chicago travel survey) suggest similar prediction disparity issues and both demonstrate the effectiveness of our bias mitigation approach. The code for our experiments is available online at: <https://github.com/zhengyunhan/Equality-of-Opportunity-Travel-Behavior.git>.

Prediction disparities could bias transport policy decisions unfavorably against socially disadvantaged groups, such as low-income and ethnic minority populations. For example, when the estimate of African-American communities' demand for transit has a higher error rate, transit agencies would make more mistakes when considering adding bus routes and investing in new transit stations for minority neighborhoods. Echoing Title VI, the Department of Transportation regulations in 49 CFR Part 21 are designed to ensure that "no person in the United States, based on race, color, or national origin, is excluded from participation in, denied the benefits of, or otherwise subjected to discrimination under any program that DOT financially assists" [15]. Therefore, people from different groups of interest (such as race, gender, income) should be treated fairly by the models and algorithms which have been widely deployed to inform transportation project planning and policy making.

The paper is organized as follows: the next section reviews travel behavior modeling and fairness in machine learning. Section 3 introduces data and models, where we introduce methods for measuring and mitigating prediction disparities, and also illustrate how the methods are tied to fairness theories in machine learning. Section 4 presents the results,

which include the simulation experiments on the synthetic data, the quantification of prediction disparity across various dependent variables and protected variables in the NHTS data and the Chicago travel survey data, as well as the results of bias mitigation for both synthetic and real-world data. Section 5 and 6 summarize the key findings and discuss the study limitations and future research directions.

2 Literature Review

Discrete choice modeling has been used in numerous travel behavior studies. With the development of artificial intelligence, machine learning techniques such as DNN have become increasingly popular and been adopted widely to achieve higher prediction accuracy [16, 17, 18]. Researchers adopted DNN to predict travel mode choices [5, 19], car ownership [1], traffic accidents [20], travellers' decision rules [21], driving behaviors [22], traffic flows [23, 24] and many other decisions in the transportation field. Some researchers focused on comparing the performance of various machine learning models such as DCM, DNN, SVM, and random forest on travel behavior prediction [25, 26, 27, 28, 29]. However, nearly all of these studies evaluate the performance of different models in terms of accuracy, stability and interpretability.

The previous machine learning literature has presented a great body of evidence showing that a model can act discriminatorily on one population in a variety of settings including, but not limited to, criminal risk assessment [8, 30], clinical care [31, 32], online advertisement delivery [33, 9], text classification [34, 10] and credit risk evaluation [35, 36]. Approaches to understand and address unfairness in machine learning models include formalizations of fairness in machine learning [37, 38], designing fairness-enhancing algorithms [39, 40, 41] and solving fairness concerns in real-world industries [42, 43]. In contrast to the rich literature in computer science focusing on fairness analysis, the authors cannot identify any study that has investigated the fairness issue in travel behavior prediction.

On the other hand, fairness has long been a crucial concern in transportation planning, and there has been numerous transport equity analysis [44, 45, 46, 47]. Among the variety of equity-focused transportation studies, most carried out fairness discussions in terms how transportation infrastructure, mobility services, opportunities and rights are distributed among the population [11, 12]. Research in this domain usually involves the cost and benefit analysis for specific groups or individuals related to a specific transportation project

[12]. The analysis of cost may include transportation cost [13] and environmental cost [48], such as the air and noise pollution produced by transportation-related construction. The benefit analysis focuses on the benefits people receive in terms of accessibility, mobility and economic vitality [44, 49]. In these studies, fairness is usually evaluated based on the distributions of costs and benefits among different demographics, neglecting the fact that the machine-learned models deployed to estimate travel demand - which is critical for cost and benefit analysis - itself can be unfair [50, 38]. Therefore, instead of focusing on substantive fairness which emphasizes resource allocation and decision making, we take a step back and examine bias exhibited in the model estimation results. As such, this study aims to enrich the fairness discussion by focusing on machine-learned models, investigating the fairness issue in the modeling process and analyzing the biases that arise in the modeling results.

Among a variety of fairness definitions, this paper uses equality of opportunity to define fairness. Equality of opportunity is a relaxation of the fairness measure equality of odds, and equality of odds states that the protected and unprotected groups should have equal rates for true positives and false positives [51, 52]. Since achieving equal rates for both measures (true positives and false positives) in practice is usually hard, equality of opportunity is adopted instead, stating that protected and unprotected groups should have equal true positive rates [51, 52]. A similar fairness definition - “reciprocal equality of opportunity” - is also adopted in this study, which requires that the protected and unprotected groups have equal true negative rates [53].

Equality of opportunity is a type of “disparate impact” analysis which evaluates fairness based on model impact (results) - specifically whether policies or practices have a disproportionately adverse impact on protected classes [50]. This fairness notion is chosen as it is inherently connected with the notion of equality of opportunity in the traditional transportation equity literature. In the traditional transport equity literature, equality of opportunity focuses on applications and resource allocations. It asserts that the education, employment and consumer opportunities accessible to residents should be equal between different groups [12]. Violation of equality of opportunity in travel behavior predictions can consequently affect the transportation resources accessible by different populations, thus perpetuating inequality of opportunity in reality.

In addition to equality of opportunity, another widely-adopted fairness measure focusing on disparate impact is “demographic parity” [51, 52], which is achieved when the likelihood of

a positive outcome is the same regardless of whether the person is in the protected group. For example, when studying gender fairness in predicting the usage of public transit, demographic parity is achieved when the proportion of people predicted as using public transit frequently is the same between male and female. Equality of opportunity is preferred to demographic parity since the latter fails to account for discrimination which is explainable in terms of legitimate grounds [54].

Apart from disparate impact analysis, another strand of research - called “disparate treatment” analysis - evaluates fairness in terms of treatment rather than modeling result to see if the decisions are made (partly) based on the subject’s sensitive attribute information [55]. This type of fairness includes “fairness through unawareness” and “counterfactual fairness”. In the definition of fairness through unawareness, an algorithm is fair as long as any protected attributes are not explicitly used in the decision-making process [56, 57]. On the contrary, counterfactual fairness deems a predictor to be fair if its output remains the same when the protected attribute is flipped to its counterfactual value [38, 58]. Disparate treatment emphasizes explicit formal classification and intentional discrimination. Therefore, in many machine learning modeling cases where there is no discriminatory intent, disparate impact doctrine is more suited to analyzing unintended biases in data mining compared with disparate treatment doctrine.

The various fairness definitions can also be categorized based on whether they are individual-focused or group-focused. Group fairness, such as equality of opportunities, requires a fair model to treat different groups equally, whereas individual fairness refers to the rule that deems a predictor fair if it produces similar outputs for similar individuals [38, 59]. Though we use equality of opportunities as the fairness definition in this research, the above-mentioned fairness definitions can be adopted for future studies in this area.

3 Data and Methods

3.1 Equality of opportunity as the definition of fairness

This study measures fairness by equality of opportunity, mathematically denoted as $P(\hat{y} = 1|z = 0, y = 1) = P(\hat{y} = 1|z = 1, y = 1)$, in which y represents the binary travel behavioral outcomes, \hat{y} represents the predicted values, z represents the protected variable such as race and gender. Intuitively, equality of opportunity requires predicted travel behavior to be

conditionally independent of the protected attributes given that the real outcome is positive [52]. Taking racial disparity as an example, equality of opportunity implies that the predicted travel demand is independent of the travelers being in minority or majority groups, thus achieving a socially non-discriminatory predictive performance [60]. A related concept is referred to as reciprocal equality of opportunity, denoted as $P(\hat{y} = 1|z = 0, y = 0) = P(\hat{y} = 1|z = 1, y = 0)$, implying that the predicted travel behavioral outcome is conditionally independent of the protected attributes given that the real outcome is negative [53].

3.2 Data and Variables

In this study, numerical experiments are conducted on two datasets: a group of synthetic datasets and the 2017 National Household Travel Survey data. For each of these two datasets, both BLR and DNN are deployed for model estimations. The experiments on the synthetic data are used to show how the covariance between a protected variable and an explanatory variable may lead to disparate results, and how this prediction disparity varies with the covariance of these two variables, the sample size, and the number of input variables. We then test our bias mitigation method on the synthetic data. For the NHTS data, we examine the prediction disparity for a series of protected (sensitive) variables (e.g. race, gender, income, medical conditions and urban-rural divide) and different dependent variables. Our bias mitigation method is later tested on this real-world dataset as well.

3.3 Models and Bias Measurement

In this study, we adopt two models, BLR and DNN, for model predictions, which are later evaluated by fairness metrics for different demographics.

3.3.1 Binary Logistic Regression (BLR)

As a classic travel behavior modeling method, BLR has been widely deployed to predict the probability of a certain outcome. The outcome probability is defined as follows:

$$P(y_i = 1|\mathbf{x}_i) = \frac{1}{1 + \exp(-(\alpha + \mathbf{x}_i^\top \boldsymbol{\beta}))}$$

where y_i identifies the dependent variable for individual i , \mathbf{x}_i represents the vector of all the independent variables, α is the intercept, $\boldsymbol{\beta}$ is the vector of parameters associated with attribute \mathbf{x}_i , which is estimated by the negative log likelihood loss function.

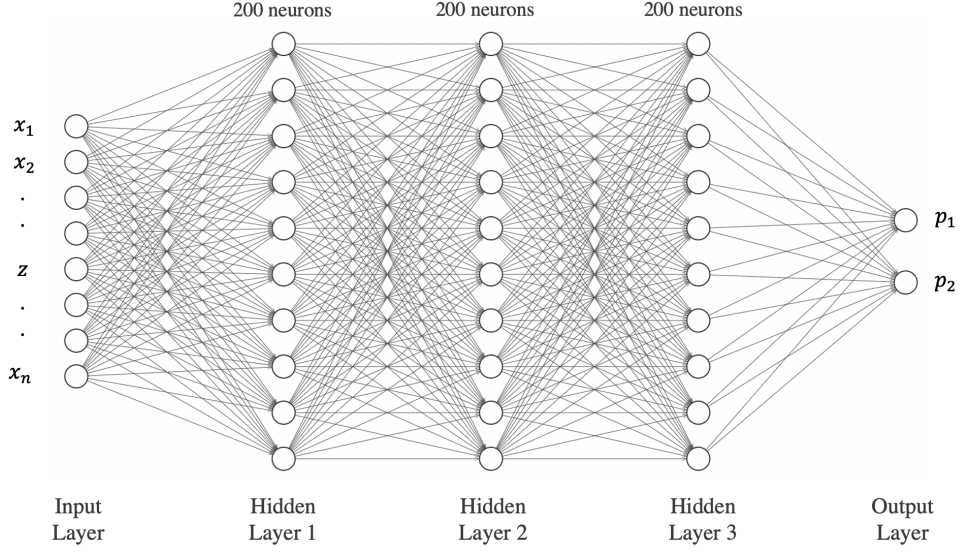


Figure 1: DNN architecture

3.3.2 Deep Neural Network Modeling (DNN)

DNN usually outperforms traditional methods regarding prediction accuracy, because of the non-linear transformation. The outcome probability using DNN can be expressed as follows [27, 26]:

$$P(y_i = 1|\mathbf{x}_i) = \frac{1}{1 + \exp(-\Phi(\mathbf{x}_i, \boldsymbol{\beta}))}$$

where $\Phi(\mathbf{x}_i, \boldsymbol{\beta})$ represents a layer-by-layer transformation: $\Phi(\mathbf{x}_i, \boldsymbol{\beta}) = (g_m \circ \dots \circ g_2 \circ g_1)(\mathbf{x}_i; \boldsymbol{\beta})$, in which each $g_l(\mathbf{x}_i^\top \boldsymbol{\beta}) = \text{ReLU}(\mathbf{x}_i^\top \boldsymbol{\beta} + b_l)$ denotes one standard module in DNN which consists of linear and rectified linear unit (ReLU) transformations. A dropout layer with a dropout rate of 0.01 is applied after each standard module to prevent overfitting. The architecture of the DNNs used in this study is shown in Figure 1, which includes 3 hidden layers and 200 neurons in each layer. Noted that the protected variable z is also included in \mathbf{x} as an explanatory variable.

3.3.3 Bias Measurement

After applying BLR and DNN models for a specific prediction task, we measure the fairness metrics for the prediction result. Based on the fairness definition of equality of opportunity, unfairness occurs when the machine-learned models offer much worse results for some demo-

graphic groups than others [52]. The degree of unfairness is measured by the false positive rate (FPR) gap or the false negative rate (FNR) gap between two groups depending on the specific context. The two fairness metrics are calculated as:

$$\text{False Positive Rate (FPR) Gap} = \frac{FP_{z=0}}{TN_{z=0} + FP_{z=0}} - \frac{FP_{z=1}}{TN_{z=1} + FP_{z=1}} \quad (1)$$

$$\text{False Negative Rate (FNR) Gap} = \frac{FN_{z=0}}{TP_{z=0} + FN_{z=0}} - \frac{FN_{z=1}}{TP_{z=1} + FN_{z=1}} \quad (2)$$

In the above expressions, TP_z , FP_z , TN_z , and FN_z represent the number of true positives, false positives, true negatives and false negatives in class z , with $z = 0$ representing the disadvantaged group. For example, when examining racial bias in predicting frequent usage of rideshare, $FN_{z=0}$ represents the number of individuals in the minority group who frequently use rideshare but are wrongly categorized as not doing so, and $FP_{z=0}$ represents the number of individuals in the minority group who do not frequently use rideshare but are wrongly categorized as doing so. Higher false negative rate or false positive rate for the minority group intuitively suggests that the algorithm makes more mistakes on the ethnic minority group with regard to predicting whether an individual uses rideshare frequently, which might lead to significant mismatch between demand and supply of TNC service in minority neighborhoods. Therefore, our definition captures the essential intuition in transport equity discussions. For these two fairness metrics, lower absolute value is better.

3.4 Bias Mitigation

Adapted from the work of Beutel et al. [14], we mitigate prediction disparity by adding a regularization term to the loss function. While Beutel et al. [14] used correlation between the output distributions of two groups as their regularization term since the outputs in their studies are continuous scores, we use the correlation between the predicted probability distributions of two groups instead, which makes the regularization term differentiable in our classification tasks. This regularization loss term helps shrink the difference of prediction disparity across groups towards zero. Compared with other approaches which generally come with notable engineering concerns, this approach is lightweight, can be easily adapted to real-world systems and has achieved good empirical results [14]. The loss function is specified as:

$$\min_p (1 - \lambda)L_{primary} + \lambda|Corr(p(\mathbf{x}), z|y = q)| \quad (3)$$

where

$$L_{primary} = \sum_{i=1}^N [-y_i \log(p(\mathbf{x}_i)) - (1 - y_i) \log(1 - p(\mathbf{x}_i))] \quad (4)$$

$$Corr(p(\mathbf{x}), z | y = q) = \frac{\sum_{i \in S_q} (p(\mathbf{x}_i) - \overline{p(\mathbf{x}_i)})(z_i - \overline{z_i})}{(\sqrt{\sum_{i \in S_q} (p(\mathbf{x}_i) - \overline{p(\mathbf{x}_i)})^2} + \epsilon) * (\sqrt{\sum_{i \in S_q} (z_i - \overline{z_i})^2} + \epsilon)} \quad (5)$$

$$S_q = \{i | y_i = q\}; \epsilon = e^{-20} \quad (6)$$

In the above equation, \mathbf{x}_i is the vector representing the explanatory variables. y_i denotes the true outcome and $p(\mathbf{x}_i)$ represents the estimated probability of the output $y_i=1$ using the explanatory variables. q is equal to 1 if the false negative bias is examined and is equal to 0 if the false positive bias is examined. z_i denotes the value of the protected variable. ϵ is added to prevent the denominator from becoming zero. S_q represents the set of samples with $y_i = q$, which is used to compute the correlation loss.

$L_{primary}$ is a negative log likelihood loss function for the DNN. The correlation term is a penalty added to the model based on the distribution of the predictions. By reducing the correlation between $p(\mathbf{x}_i)$ and z_i conditioning on $y_i = q$, it minimizes the conditional dependence between the distribution of the predicted probabilities and the group membership determined by the protected variable. λ is a parameter controlling the tradeoff between primary loss and fairness loss. When $\lambda = 0$, no bias mitigation is employed. In this research, we demonstrate the effectiveness of using this fairness-adjusted loss function in both BLR and DNN.

4 Experiments

In this section, the experiment results for the synthetic dataset, the NHTS data and the Chicago travel survey data are reported. For these three datasets, we first conduct BLR and DNN for model estimations, then examine fairness issues in the prediction results, and lastly apply our bias mitigation methods. DNN is implemented using the TensorFlow library in Python. BLR is implemented using the scikit-learn library when examining the fairness issues while using the TensorFlow library for bias mitigation, since TensorFlow allows us to modify the loss function.

Experiments conducted in TensorFlow all use the mini-batch stochastic gradient descent method with a batch size of 1,000 and a step size of 0.0001 in each training. We draw samples without replacement to generate the mini-batches within an epoch. After a mini-batch is generated, the algorithm calculates the prediction loss and updates the coefficients. The model that produces the lowest training loss among the 50 epochs is chosen and later performs prediction over the test data. We ran 5 trials of 5-fold validation for each experiment.

4.1 Synthetic Experiment

In our simulations, we consider the type of bias that arises when the true predictor and the protected variable are highly correlated. In this case, the true predictor of the outcome also happens to serve as a reliable proxy for class membership in the training set [50, 61]. For example, if the usage of rideshare is positively associated with income, and ethnic minorities in the training data tend to have lower income, the algorithm will tend to predict low usage of rideshare for the minority population, even if the true contributing factor for rideshare usage is income rather than race. This type of bias is inherent in the existing population inequality and has been less emphasized in previous literature.¹

4.1.1 Data generation process

In the synthetic dataset, each data point can be represented by a tuple $(z_i, x_i, \mathbf{k}_i, y_i)$, where z represents the protected variable (e.g. race, gender), x is the explanatory variable that is correlated with z (e.g. income), \mathbf{k} is a vector of explanatory variables which does not include x and z . y represents the binary outcome.

First, x and all the elements in \mathbf{k} are drawn from the independent standard normal distributions. z is generated as a binary variable that is positively correlated with x and is independent with \mathbf{k} . The label y is drawn from a binomial distribution with probability $Pr(y = 1) = \frac{1}{1 + \exp(-V)}$. The systematic utility function V takes x and \mathbf{k} as the input variables.

¹Besides this inherent bias, other sources of bias could exist, such as the imbalanced training data problem. This problem arises when the disadvantaged group has insufficient training data or is misrepresented, in which case the model will either fail to learn a correct statistical pattern or favor the majority group during the estimates, since the training data in the disadvantaged groups often misrepresent the true population when they are insufficient [62, 63, 31, 64, 65].

We test two scenarios with the true utility function taking a linear form and a quadratic form respectively. Let $V = \alpha + \mathbf{w}\phi(x, \mathbf{k})$. In the first scenario, $\phi(x, \mathbf{k})$ takes the linear transformation: $\phi(x, \mathbf{k}) = [x, k_1, k_2, \dots, k_d]$. The weight for x is set to 1, and the weights for other explanatory variables are chosen from the set $\{-0.5, 0.5\}$ with equal probabilities. In the second scenario, $\phi(x, \mathbf{k})$ takes the quadratic transformation: $\phi(x, \mathbf{k}) = [x, k_1, \dots, k_d, x^2, k_1^2, \dots, k_d^2]$. The weights for x and x^2 are set as 1 and 0.5, and the weights for other explanatory variables take $\{-0.5, 0.5\}$ with equal probabilities. Our data generation process makes sure that the mean value of z is 0.5, which means that there are approximately equal numbers of $z = 0$ and $z = 1$, thus giving us a balanced dataset. Detailed descriptions of the data generation process can be found in Appendix A.

We examine the cases where $Cov(z, x)$ is non-negative and x positively affects V . We therefore use $z = 0$ to mimic the disadvantaged population and use $z = 1$ to mimic the privileged population, since in the real world the privileged population (e.g. the ethnic majority group) is often positively correlated with the factor (e.g. income) that has a positive effect on the utility of an advanced mobility service (e.g. the utility of using a ride-hailing service).

4.1.2 Fairness measurement results

In the prediction phase, we use z , x and \mathbf{k} as the explanatory variables for both the logit model and the DNN model, so these two choice models are defined as $Logit(z, x, \mathbf{k})$ and $DNN(z, x, \mathbf{k})$.

First, we want to examine how fairness measures and accuracy vary regarding the correlation between the sensitive attribute z and the explanatory variable x , the sample size and the number of explanatory variables in the data generation process. Therefore, we run experiments along these three dimensions, and when each dimension is examined, we set the other two dimensions as the default values. The default values for $Cov(z, x)$, sample size and number of explanatory variables are 0.2, 10^6 and 5. For each experiment, three datasets are randomly generated based on the above data generation process.²

²Occasionally, the data generation process in Scenario 2 produced datasets with highly imbalanced outcomes (e.g. when the minority class accounts for less than 30% of the total samples). In that case, we would drop that imbalanced dataset and generate another one. We iterated this process until all the datasets are roughly balanced (when the minority class accounts for 40%-50% of the total samples)

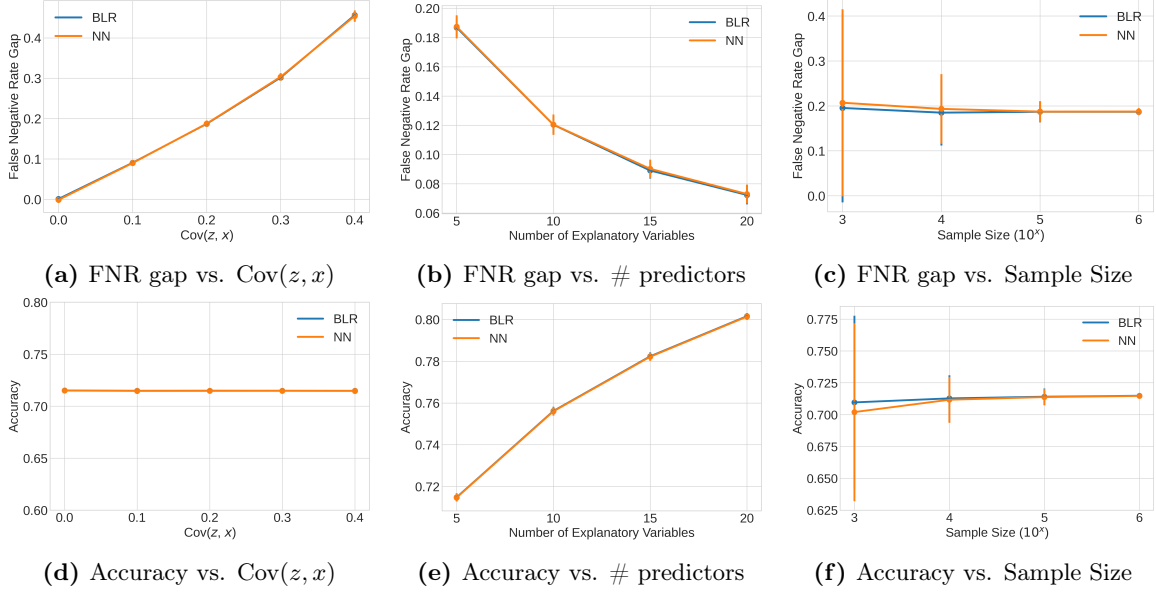


Figure 2: Fairness metric and accuracy with different parameters (BLR vs. DNN): true model taking the linear form; estimation models: $Logit(z, x, \mathbf{k})$ and $DNN(z, x, \mathbf{k})$. $Logit(z, x, \mathbf{k})$ adopts a linear specification, so both models contain the true model.

Figure 2 presents the results of the linear data generating process, while in Figure 3, the true data generating process is quadratic in variables. In both figures, the first row shows the FNR gap (see Equation 2) between the disadvantaged group (defined as $z = 0$) and the privileged group (defined as $z = 1$) as the measure of fairness and the second row shows the prediction accuracy. The x-axis of the first, second and third columns respectively represent the covariance between z and x , the number of explanatory variables in the data generation process and the sample size. Each figure plots both the BLR and DNN results, which are represented by the blue and orange colors. The figures plot the values averaged across 5 trials of 5-fold validation in 3 datasets for each experiment; the error bar indicates the standard deviation multiplied by 1.96, which approximates the confidence interval of the estimations.

As shown in Figure 2, the results of BLR and DNN mostly overlap since they both recover the true linear model. Figure 2a shows that the FNR gap increases with $Cov(z, x)$, indicating that as x becomes more positively correlated with z , the algorithm is more likely to falsely associate the disadvantaged population ($z=0$) with negative outcomes, even if their real outcomes are actually positive. Prediction disparity is a metric relatively independent

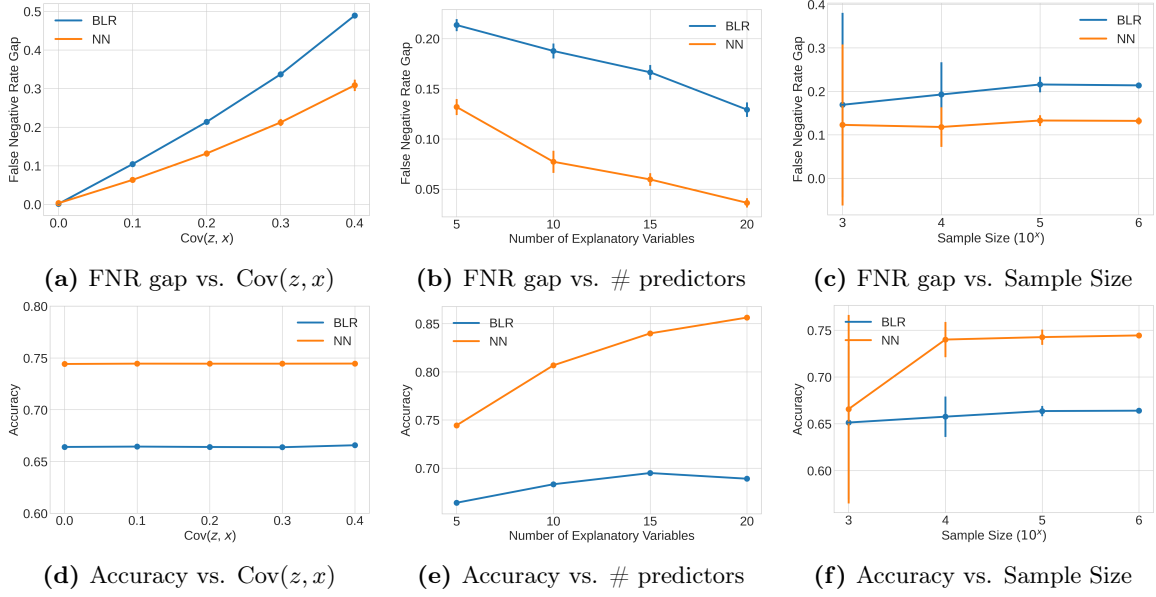


Figure 3: Fairness metric and accuracy with different parameters (BLR vs. DNN): true model taking the quadratic form; estimation models: $Logit(z, x, \mathbf{k})$ and $DNN(z, x, \mathbf{k})$. $Logit(z, x, \mathbf{k})$ follows a linear model specification, so it has the misspecification error.

of the predictive performance, as there is no difference in prediction accuracy with different $Cov(z, x)$ (Figure 2d); it is also not due to the imbalanced training data problem, as the outcome variable, the protected variable and all the explanatory variables are balanced in the training data. The prediction disparity is purely inherent in the relationship among variables in the data. Figure 2b shows that the FNR gap decreases with the number of explanatory variables, which is probably because increasing the number of predictors dilutes the influence of x on the outcome. Figure 2c shows that the variance of the fairness and accuracy estimations decrease as the sample size increases.

Figure 3 shows the results of prediction fairness and accuracy when the true data generation model takes a quadratic form. In this case, the BLR with linear specification has the model misspecification error while DNN does not. Figure 3a shows that prediction disparity still increases with the increase of $Cov(z, x)$, and DNN is always associated with smaller FNR compared with BLR for $Cov(z, x) > 0$. This result indicates that the model misspecification not only induces more prediction error, but also harms prediction fairness. Figure 3b shows that though increasing the number of explanatory variables can reduce the prediction disparity, the magnitude of the prediction disparity caused by model misspecification was not

reduced. Figure 3c indicates that the fairness prediction result becomes more stable as the sample size increases.

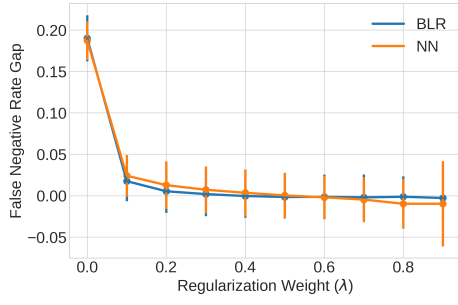
4.1.3 Bias mitigation results

To address prediction disparity, we apply the bias mitigation method illustrated in Section 3.4 to the synthetic datasets with $Cov(z, x) = 0.2$, sample size of 100,000 and number of predictors set to 5. For each regularization weight λ , we run the training procedure 5 times for each of the 3 datasets with 5-fold cross-validation and report the average results in Figure 4 for Scenario 1 and Figure 5 for Scenario 2. The error bars in the figures indicates the standard deviation multiplied by 1.96, which approximates the confidence interval.

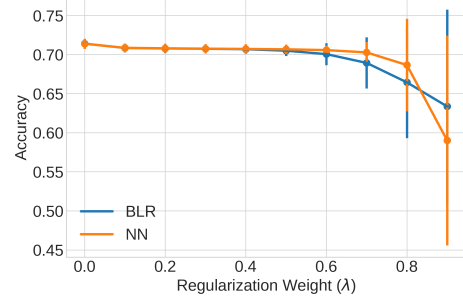
Figure 4a and 5a show that in both scenarios, applying the regularization even with a small weight (e.g. $\lambda = 0.1$) can substantially reduce the prediction disparity and this finding holds for both BLR and DNN. Given the model misspecification for BLR in Scenario 2, Figure 5a shows that our method is still effective in reducing the prediction bias to as low as zero.

Figure 4b and 5b report the corresponding model accuracy as the regularization weight varies. The results show that when $\lambda < 0.7$, the accuracy decreases only slightly. These results suggest that improvement in prediction fairness can be achieved with a minimal cost of accuracy.

We also test our mitigation method on the dataset with an imbalanced distribution over the protected variable. For the data generated in Scenario 2, we randomly drop half of the data in the disadvantaged group to make the ratio of number of samples in the disadvantaged and the privileged groups roughly 1:2. Figure 6 shows the bias mitigation results with this new dataset. It shows that the increase in bias mitigation weight leads to a reduction in the mean FNR gap for both DNN and BLR, and the fairness-accuracy trade-off still exists. However, for DNN, though the FNR gap continues to decrease with increasing λ , the absolute value of the FNR gap becomes larger when $\lambda > 0.2$, indicating that the model may adjust too much to mitigate the bias when the bias mitigation weight is too large.

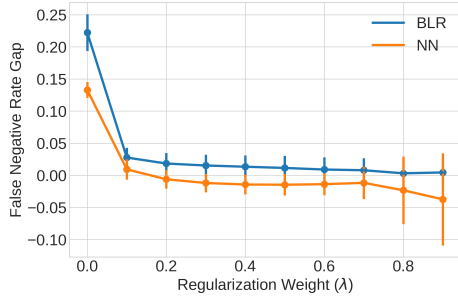


(a) FNR gap vs. Regularization Weight

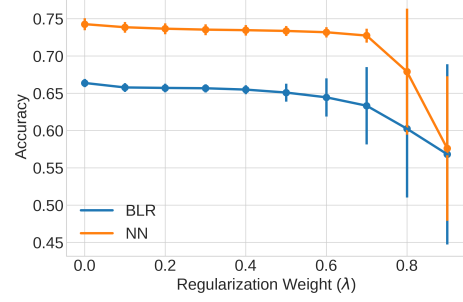


(b) Accuracy vs. Regularization Weight

Figure 4: Fairness and accuracy by bias mitigation weight (λ): true model taking the linear form (Scenario 1)

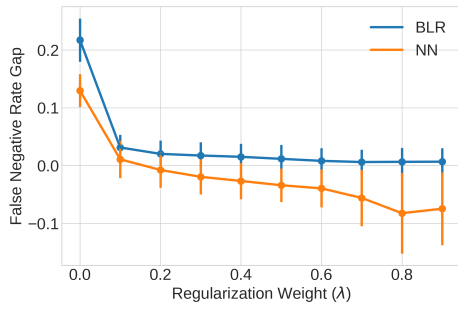


(a) FNR gap vs. Regularization Weight

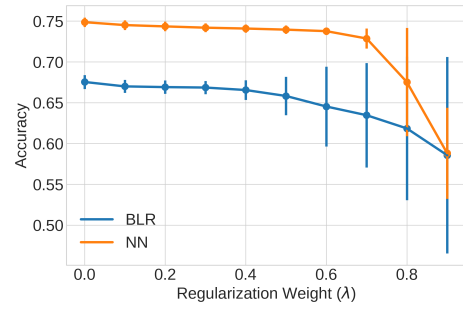


(b) Accuracy vs. Regularization Weight

Figure 5: Fairness and accuracy by bias mitigation weight (λ): true model taking the quadratic form (Scenario 2)



(a) FNR gap vs. Regularization Weight



(b) Accuracy vs. Regularization Weight

Figure 6: Fairness and accuracy by bias mitigation weight (λ): imbalanced dataset

4.2 The NHTS Dataset

4.2.1 Data and variables

The NHTS data are collected directly from a stratified random sample of U.S. households. The richness of the dataset enables us to examine fairness in predictions with varying dependent variables and protected attributes. Protected attributes are the variables we want to protect against in the model prediction process, which include race, gender, income, medical condition and urban-rural divide in this study. In terms of “race”, we define the ethnic minority group as the non-white population. In terms of the variable “income”, we identify low-income households based on the combination of household size and last year’s household income following the 2017 Health and Human Services poverty guidelines [66]. An individual is deemed to have a “medical condition” if he or she answered “yes” to the question “do you have a condition or handicap that makes it difficult to travel outside of the home?” in the survey. Regarding the protected variable “region”, the question “household in urban area? Answer ‘yes’ or ‘no’.” is used to identify whether the individual is an urban or rural resident.

The dependent variables examined in this study can be categorized into two groups: the first group contains four variables indicating the “yes” or “no” answers to “usually work from home”, “have the option of working from home”, “agree that travel is a financial burden” and “agree that gas price affects travel”; the second group contains four variables indicating the high frequent usage of four travel modes: bike, car, bus and rideshare. These eight dependent variables are all binary variables, with “yes” taking the value 1. A detailed description of these variables can be found in Appendix B.

The distributions of the dependent variables except “travel burden” and “gas price impact” are highly skewed, and previous research has found that when the outcome class sizes are highly imbalanced, the classification algorithms tend to strongly favor the majority outcome class, resulting in very low or even no detection of the minority outcome class [67]. Therefore, for each of the six imbalanced dependent variables, we balance the data to facilitate training by downsampling the majority class. The summary statistics of the independent and dependent variables as well as the distributions of two groups of dependent variables by different protected attributes are reported in Appendix B.

4.2.2 Bias mitigation method with sample weights

As previously mentioned, one source of bias in the data is that the training data might not be representative of the overall population. Luckily, NHTS contains the sample weight³ [68] for each individual, which can be used to address the representation bias. We incorporate the sample weights in the training and evaluation phases. Weighted accuracy and weighted fairness metrics are used for model evaluation. To be specific, the weighted accuracy is calculated as:

$$\text{Weighted Accuracy} = \frac{\sum_i^N 1(\hat{y}_i = y_i)w_i}{\sum_i^N w_i} \quad (7)$$

where w_i represents the sample weight for sample i , y_i is the label and \hat{y}_i is the predicted outcome. Similarly, the sample weight is applied for each sample when calculating fairness metrics (FNR and FPR). N denotes the sample size.

Corresponding to the weighted evaluation metrics, the sample weights are also applied in the loss function during the training process. The weighted loss function is written as:

$$\min_p (1 - \lambda)L_{\text{primary}} + \lambda|\text{Corr}(p(\mathbf{x}), z|y = q)| \quad (8)$$

where

$$L_{\text{primary}} = \sum_{i=1}^N \frac{w_i L_i}{\sum_i^N w_i}; L_i = -y_i \log(p(\mathbf{x}_i)) - (1 - y_i) \log(1 - p(\mathbf{x}_i)) \quad (9)$$

$$\text{Corr}(p(\mathbf{x}), z|y = q) = \frac{\sum_{i \in S_q} w_i (p(\mathbf{x}_i) - m(p(\mathbf{x}_i)))(z_i - m(z_i))}{(\sqrt{\sum_{i \in S_q} w_i (p(\mathbf{x}_i) - m(p(\mathbf{x}_i)))^2} + \epsilon) * (\sqrt{\sum_{i \in S_q} w_i (z_i - m(z_i))^2} + \epsilon)} \quad (10)$$

$$m(p(\mathbf{x}_i)) = \frac{\sum_{i \in S_q} w_i p(\mathbf{x}_i)}{\sum_{i \in S_q} w_i}; m(z_i) = \frac{\sum_{i \in S_q} w_i z_i}{\sum_{i \in S_q} w_i} \quad (11)$$

$$S_q = \{i | y_i = q\}; \epsilon = e^{-20} \quad (12)$$

BLR is implemented through the scikit-learn library when evaluating fairness issues and through TensorFlow when mitigating the bias. For DNN, all of the experiments are implemented in TensorFlow. Each of the experiments conducted in TensorFlow uses the mini-batch gradient descent method with a step size of 0.0001 during training. The best model among the 5000 epochs is chosen and later performs prediction over the test data. 5-fold

³which is primarily calculated as the inverse of the probability of selection of the person in the given sampling stratum from the sampling frame

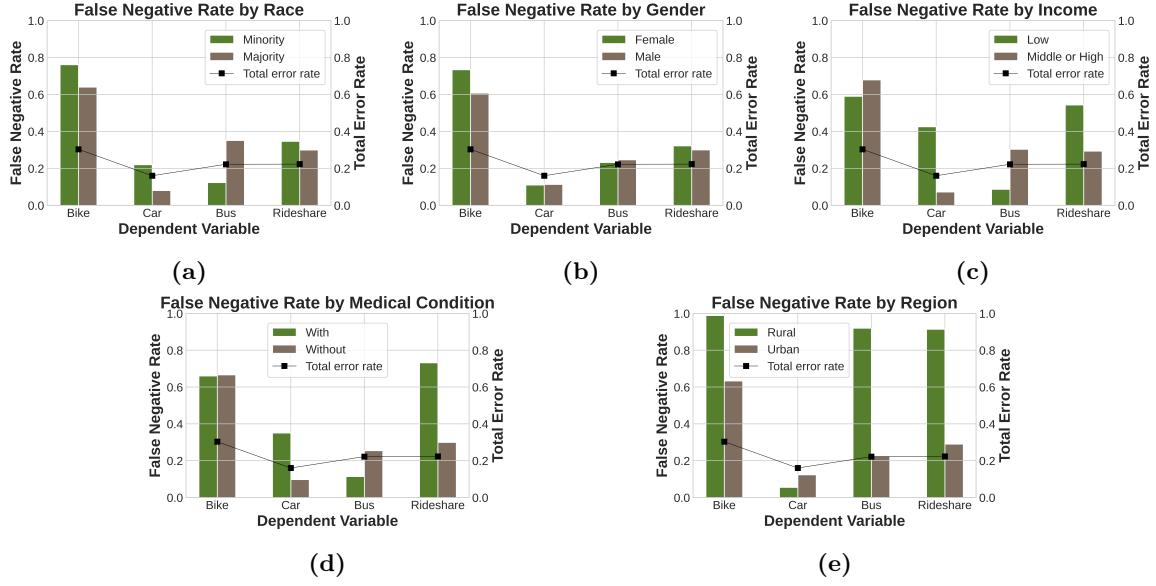


Figure 7: Disparity of prediction accuracy (BLR): frequent usage of bike, car, bus and rideshare

validation is conducted for each experiment.

4.2.3 Fairness issues in the adoption of BLR and DNN

The comparison of prediction accuracy with respect to various protected variables are presented by the bar charts in Figure 7 and 8 for BLR and in Figure 9 and 10 for DNN. Each bar chart depicts the prediction accuracy of two populations grouped by a specific protected variable (race, gender, income, medical condition, region) side by side. The height gap of two adjacent bars shows the prediction disparity for that protected variable. Figure 7 and 9 illustrate the prediction results of the dependent variables regarding travel mode usage by BLR and DNN. Figure 8 and 10 plot the prediction results of the dependent variables “work from home”, “work from home option”, “travel burden” and “gas price impact” by BLR and DNN. The dependent variables are specified on the x-axis of the bar charts.

The y-axis of the bar charts represents one of the error rates: FPR or FNR. We examine FNR for the first group of dependent variables since we are concerned about cases where active users of a certain travel mode are not identified (Figure 7 and 9), and FPR for the second group of dependent variables (Figure 8 and 10) since we want to focus on instances which have negative outcomes but are wrongly identified as positive (e.g. people who do not

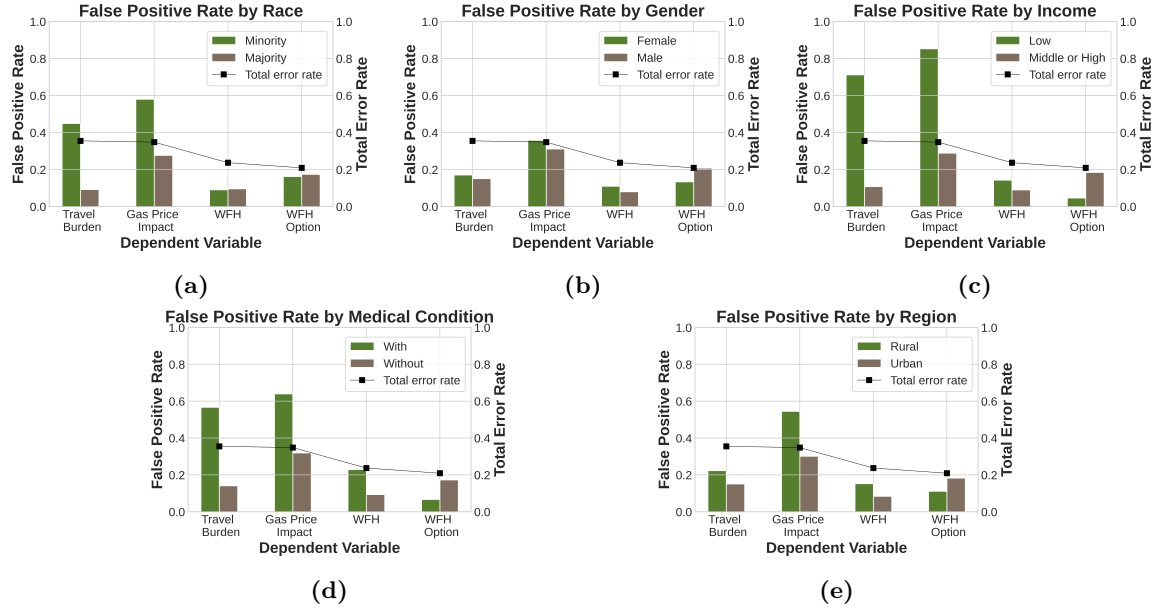


Figure 8: Disparity of prediction accuracy (BLR): work from home, work from home option, travel burden, gas price impact

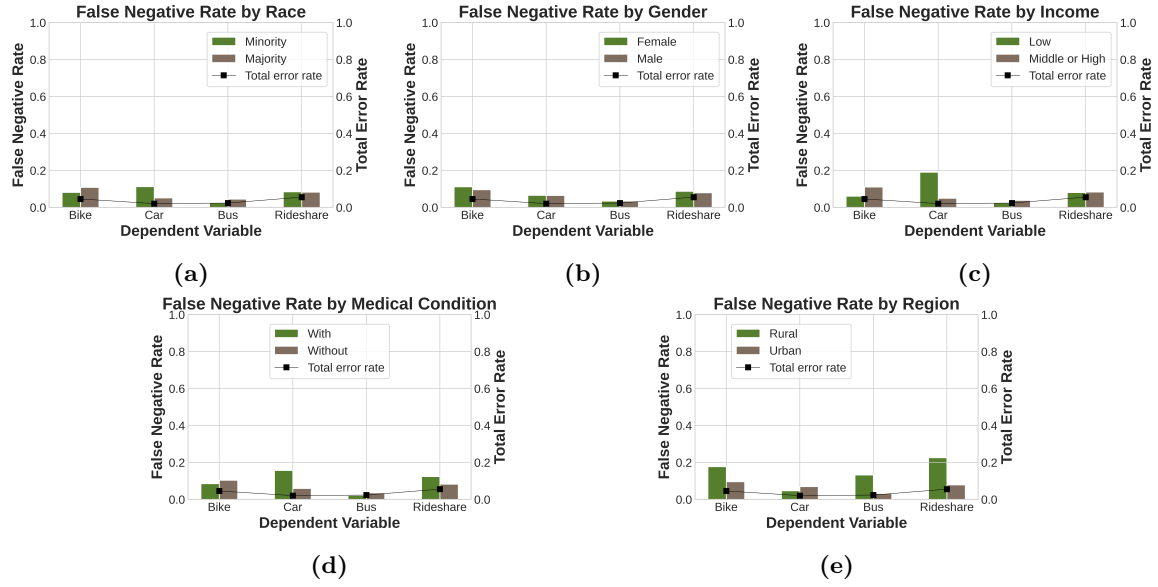


Figure 9: Disparity of prediction accuracy (DNN): frequent usage of bike, car, bus and rideshare

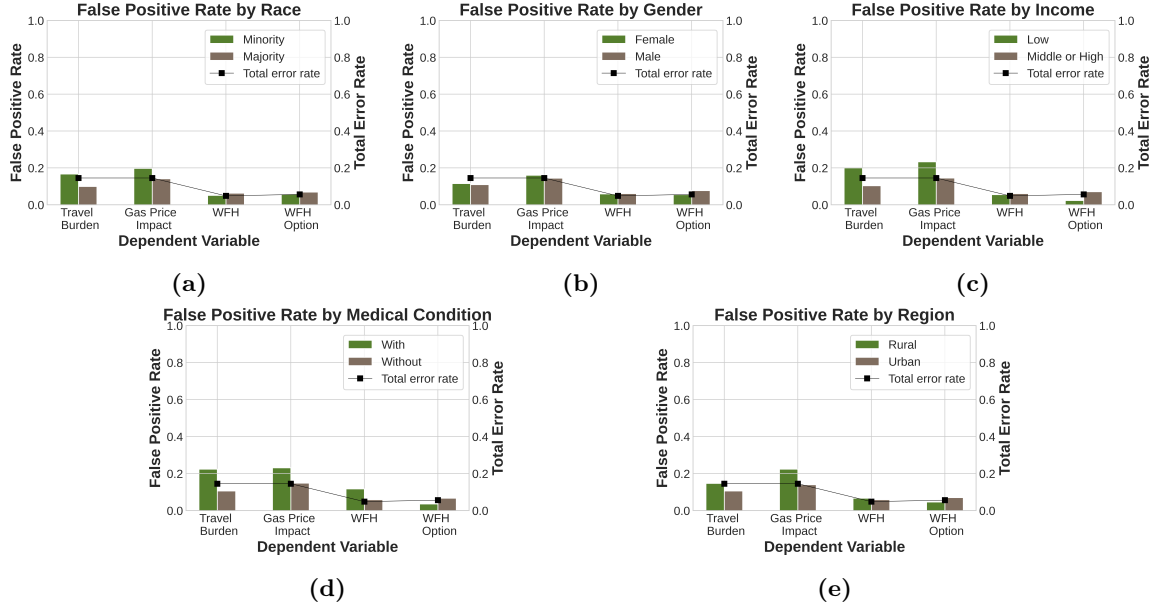


Figure 10: Disparity of prediction accuracy (DNN): work from home, work from home option, travel burden, gas price impact

have the option of working from home but are mistakenly identified as having the option). In each bar chart, the height of a bar represents the magnitude of the class-specific FNR or FPR. The total error rate is also presented which refers to one minus the weighted accuracy for all samples.

First, we focus on the BLR results. Figure 7 and 8 show that prediction disparities are common with the implementation of BLR. Figure 7 presents the false negative bias across different populations regarding the frequent usage of different travel modes. The plots show that except for the fairly consistent prediction accuracy between male and female, considerable disparity in prediction accuracy exists for all other protected variables. Racial bias is large in the prediction of frequent usage of car and bus; income bias and bias towards medical condition are considerable in the predictions of car, bus and rideshare. Among all the protected attributes, the regional disparity is the largest with respect to predicting the frequent usage of bike, bus and rideshare. As a result, the proportion of rural residents that use bike, bus and rideshare frequently is underestimated compared with that of the urban residents. If policy makers use the modeling results to inform transportation resource allocations such as the planning of bike lanes and bus routes without considering the prediction bias, the rural area will very likely be under-served.

In terms of the dependent variables, when estimating high frequent usage of rideshare, we find that for all protected variables, the disadvantaged group always has higher FNR than the other group with either BLR or DNN. This finding indicates that the communities where these disadvantaged social groups (ethnic minorities, women, low-income groups, people who have medical conditions and rural residents) dominate could experience less ride sharing service if decision makers use these data for rideshare demand estimation and do not account for the prediction disparity.

Figure 8 illustrates significant racial bias, income bias and health-related bias for the predictions of “travel burden” and “gas price impact”. In other words, the ethnic minority population, the low-income population and people with health conditions are more likely to be predicted as “regarding travel as a financial burden” and “agreeing that price of gasoline affects travel” when the true outcome is actually negative. These biases could lead to disadvantageous consequences for vulnerable populations. For example, given that banks are less likely to loan to someone if they perceive that person to be under financial stress, if the ethnic minorities suffer from higher FPR in travel burden prediction, they are more likely to be rejected when applying for loans.

Next, we compare the results between BLR and DNN. The prediction disparity and the total error rate of BLR are considerably larger than those of DNN in all scenarios, showing that the prediction disparity of BLR largely comes from model misspecifications. By reducing model misspecification, DNN can not only increase prediction accuracy, but also improve computational fairness. These findings are consistent with our simulation results as illustrated in Figure 3, which show the positive relationship between the magnitude of the fairness gap and the degree of model misspecifications.

However, the prediction results of DNN (Figure 9 and Figure 10) also show that even if the prediction achieves very high accuracy ($>94\%$ for all dependent variables in Figure 9 and $>85\%$ for all dependent variables in Figure 10), the fairness gap still exists, and for some sensitive variables the prediction disparity can be higher than 15% (e.g. the FNR gap for the frequent usage of rideshare prediction between rural and urban residents). This finding tells us that deploying complex models aimed at improving prediction accuracy cannot guarantee to eliminate prediction biases. Therefore, a dedicated method should be deployed to mitigate unintended bias rather than merely attempting to improve the model prediction power.

Besides, with the exception of the racial difference in bike frequency prediction, the signs of the fairness gaps for other protected variable and dependent variable combinations are the same between BLR and DNN. The consistency between the two models suggest that the prediction disparity may reveal bias inherent in the existing inequality in society which is baked into the data. For example, in the predictions of mode usage, it is found that the socially disadvantaged groups such as the ethnic minority group and people having medical conditions tend to suffer from higher FNR in predictions of car and rideshare usage, probably because they are negatively associated with certain factors (e.g. income) that positively contribute to the usage of these modes which are more expensive compared with other modes.

We also conduct the sensitivity analysis by changing the number of layers in DNN, the batch size and the weight initialization. The results are presented in Appendix D, which show that the fairness gaps generally still exist with these variants, and in most cases, changing these hyperparameters does not change the signs of the fairness gaps.

4.2.4 Bias mitigation results

We adopt the absolute correlation regularization method to mitigate the unintended bias for both BLR and DNN models. The method is applied to mitigate two types of prediction disparities: the FNR gap for estimating the frequent rideshare usage between rural and urban residents, and the FPR gap for the prediction of “travel burden” between the ethnic minority group and the majority group, as the initial prediction disparity in these two prediction tasks is substantial compared with those in other prediction tasks. In Figure 11 and 12, plot (a) reports the average prediction disparity and plot (b) reports the average prediction accuracy of 5-fold cross-validation with varied bias mitigation weight (λ). In both plots, higher weight indicates larger penalty for prediction bias.

Both Figure 11 and 12 demonstrate the effectiveness of our bias mitigation method. The prediction disparity decreases as we increase λ , and this effect is particularly substantial for BLR. The blue curves in these two figures show that the fairness gap diminishes sharply even if only a small degree of bias mitigation weight is applied.

In Figure 11, we find that increasing λ from 0 to 0.2 can reduce the FPR gap from 62.4% to -6.4% for BLR and from 14.5% to 4.8% for DNN, only at the expense of reducing the overall accuracy from 77.7% to 76.6% for BLR and from 94.5% to 90.6% for DNN, and a

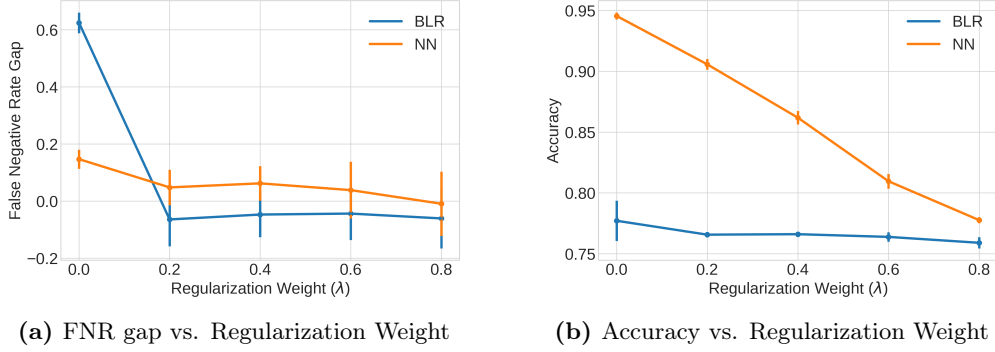


Figure 11: Fairness and accuracy by bias mitigation weight (λ): regional bias in the prediction of frequent rideshare usage

similar effect is found in Figure 12. These results indicate that with the adoption of our bias mitigation method, we can reduce the fairness gap in BLR to a similar level of that in DNN. In Appendix C, we also show the convergence of the loss functions in training.

Appendix E reports the change of the F1-score gap and the other fairness gap (the FPR gap for the frequent rideshare usage prediction and the FNR gap for the travel burden prediction) after these bias mitigations. The results indicate that the absolute values of both the FNR gap and the FPR gap decrease in these two prediction tasks after bias mitigation. The absolute F1-score gap also drops when BLR is applied. When DNN is applied, the absolute F1-score gap drops only for the travel burden prediction, but not for the frequent rideshare usage prediction.

We also report the change of prediction disparity for other protected variables in Appendix F. The results show that mitigating the bias for one variable has the possibility of increasing prediction disparity for other variables. Therefore, future research is needed to explore ways to mitigate bias for multiple protected variables. One natural extension is to add the correlation regularization terms for all protected variables of interest to the objective function during training, thus protecting fairness for various protected variables.

4.3 The Chicago Daily Travel Survey

4.3.1 Data and method

To test the validity of our research findings, we further apply the fairness evaluation and mitigation methods to the Chicago travel survey dataset. The 2018-19 My Daily Travel Sur-

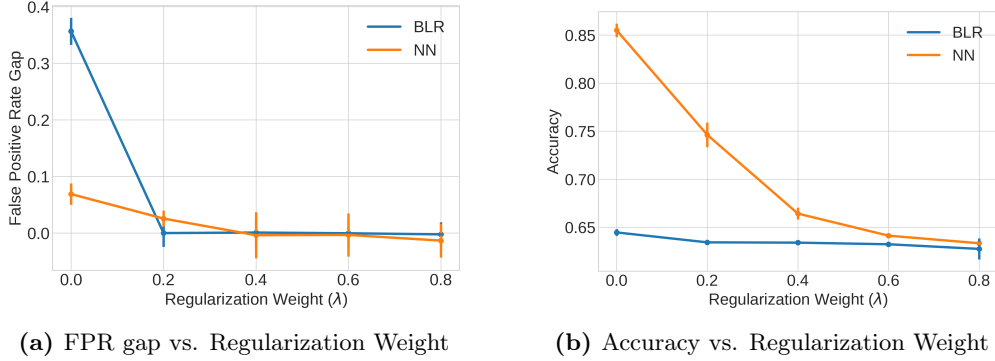


Figure 12: Fairness and accuracy by bias mitigation weight (λ): racial bias in the prediction of travel burden

vey in Chicago⁴ was produced by the Chicago Metropolitan Agency for Planning (CMAP), and the data was collected between August 2018 and April 2019. The dataset captures and represents key information regarding demographic, land-use, and travel behavior in the Chicago region.

We use this data to predict people’s commuting mode. Among the 13,730 total observations, the weighted percentage of using auto as the travel mode to work is 73.1%.

Therefore, we examine whether an individual uses an auto or non-auto travel mode to work as the dependent variable. The protected variables examined are race, income and disability. In terms of “race”, we define the ethnic minority group as the non-white population. In terms of the variable “income”, we identify low-income households based on the combination of household size and last year’s household income following the 2019 U.S. Federal Poverty Guidelines [69]. For “disability”, we compare the modeling results between disabled individuals and nondisabled individuals. Descriptive statistics of the dependent variable, the explanatory variables and the protected variables are presented in Appendix G.

4.3.2 Experiments and results

We adopt the same models that have been applied to the NHTS data to measure fairness and mitigate bias for work travel mode prediction using the Chicago travel survey data. BLR is implemented through the scikit-learn library. DNN with 3 hidden layers is implemented

⁴The data can be found on the Chicago Metropolitan Agency for Planning Data Hub: <https://datahub.cmap.illinois.gov/dataset/mydailytravel-2018-2019-public>

in TensorFlow using the mini-batch gradient descent method with step size 0.001 during training. The best model among the 1000 epochs is chosen and later performs prediction over the test data. The 5-fold validation is conducted for each experiment.

Table 1 shows that the disadvantaged groups (ethnic minority, low-income and disable populations) all have higher FNR than the privileged groups when predicting whether the work travel mode is auto, as the FNR gaps are all positive with both BLR and DNN. This is consistent with the findings in NHTS which show that these disadvantaged groups all have higher FNR when predicting frequent car usage.

Then, we apply the bias mitigation for “race”, “income” and “disability” separately, and report the bias mitigation results for each of the protected variables in Table 2. The mitigation weights of 0.6, 0.8 and 0.9 are applied to reduce bias for each of the three protected variables. The results indicate that in general, increasing the bias mitigation weight can reduce the absolute value of the FNR gap with a drop in accuracy, showing that the bias mitigation method is effective in reducing the prediction disparity.

Model	Accuracy	FNR gap		
		Race	Income	Disability
BLR	0.813 (0.012)	0.028 (0.009)	0.145 (0.05)	0.085 (0.079)
DNN	0.917 (0.001)	0.087 (0.015)	0.135 (0.116)	0.038 (0.036)

Note: This table reports the means and standard deviations (in parentheses) of the accuracy and FNR gap of 5-fold cross-validation for BLR and DNN (with 3 hidden layers).

Table 1: Fairness and accuracy results for the work travel mode prediction on the test set of the Chicago travel survey using BLR and DNN

λ	Race		Income		Disability	
	FNR gap	Accuracy	FNR gap	Accuracy	FNR gap	Accuracy
0.0	0.087 (0.015)	0.917 (0.001)	0.135 (0.116)	0.917 (0.001)	0.038 (0.036)	0.917 (0.001)
0.6	0.033 (0.029)	0.867 (0.005)	0.091 (0.078)	0.867 (0.004)	0.046 (0.09)	0.889 (0.004)
0.8	0.014 (0.025)	0.838 (0.0)	0.022 (0.032)	0.839 (0.003)	-0.023 (0.063)	0.856 (0.004)
0.9	0.001 (0.015)	0.827 (0.002)	0.032 (0.026)	0.824 (0.001)	0.019 (0.104)	0.836 (0.003)

Note: the prediction model is DNN with 3 hidden layers. λ represents the bias mitigation weight. The means and standard deviations (in parentheses) are calculated based on the 5-fold cross-validation.

Table 2: The results of bias mitigation for different protected variables (the Chicago travel survey)

5 Conclusions

This study investigates equality of opportunity as a measurement of computational fairness in travel behavior modeling with BLR and DNN. Fairness has long been a critical concern in transportation studies. However, past transportation equity studies generally evaluated equity based on the cost and benefit analysis for different populations and neglected the potential fairness problem in machine-learned models. This research aims to enrich the transportation equity research by emphasizing the integration of fairness metrics into the modeling process to improve fair transportation decision making and resource allocation.

First, we use the concept of equality of opportunity to define computational fairness in travel behavior modeling, which is measured by the gap in true positive rates between two groups of populations. This definition is connected to the equality of opportunity in traditional transportation equity literature, since choice probabilities in discrete choice models are a natural metric for opportunities. The unfair prediction of choice probabilities in modeling can lead to the allocation of inadequate transportation resources to disadvantaged neighborhoods.

Then, we conduct simulated experiments to show two sources of prediction disparities: the bias inherent in the data structure and the bias in the modeling process. Prediction disparity increases as the correlation between the protected variables and the explanatory variable(s) increases, exhibiting the bias inherent in the data structure but not caused by modelers. The inherent bias illustrates how prediction disparities may still exist without the presence of human errors. When the true model specification cannot be captured by BLR but can be captured by DNN, DNN can produce lower prediction bias than BLR, owing to its ability to capture the complex relationships among variables. This type of bias is the algorithmic bias that can be mitigated by choosing a more fitting model, thus is a sort of human bias.

Next, we conduct computational fairness analysis on the NHTS and the Chicago travel survey datasets using BLR and DNN. The results reveal the prevalence of prediction disparities in travel behavior modeling particularly when BLR is adopted, and these prediction disparities are consistent with the change of the hyperparameters including the number of DNN layers, the batch size and the weight initialization. Though the magnitudes of unfairness are different, the signs of the fairness gap (e.g. which group has higher FNR/FPR) are the same between BLR and DNN in the vast majority of our prediction cases, probably reflect-

ing the bias inherent in the existing population disparity. For the NHTS analysis, we find that the regional disparity in predictions of frequent usage of bike, bus and rideshare is the largest among all the protected attributes, indicating that bike, bus and rideshare services allocated by the predictions will be withheld from rural residents compared with urban residents. When estimating the high frequent usage of rideshare, the disadvantaged group has higher FNR than the other group for all protected variables, indicating that these disadvantaged social groups might receive insufficient ride sharing service if the TNC companies use these data for rideshare demand estimation but do not consider prediction disparity. For the Chicago travel survey analysis, the results show that the disadvantaged groups have higher FNR when predicting whether the travel mode is auto, which is consistent with the findings in the NHTS analysis.

Finally, we adopt the absolute correlation regularization method to mitigate the bias in the BLR and DNN prediction using the synthetic, NHTS and the Chicago travel survey datasets. The method proves to be effective for both BLR and DNN. When the initial fairness gap is large, applying our bias mitigation method with only a small weight of regularization can considerably improve the fairness result. Though there is an accuracy-fairness trade-off in most predictions, we can achieve substantial improvement in prediction fairness with only a small reduction in accuracy by careful selection of the bias mitigation weight. The synthetic experiments show that the bias mitigation works whether the bias is caused by inherent or modeling bias, and whether the dataset is imbalanced or not. The NHTS experiments also show that through mitigating the bias for one fairness metric (e.g. the FNR gap), generally other fairness metrics (e.g. the FPR gap and the F1-score gap) can also be improved. However, mitigating bias for one protected variable may increase the bias for other protected variables, which necessitates future research into solutions to improving fairness for multiple protected variables simultaneously.

All in all, we argue that researchers and policy makers should be aware of the normative aspect in seemingly value-neutral machine learning predictions, since the prediction disparities can lead to severely unfair treatment of already marginalized social groups. Only after acknowledging the existence of the biases can policy makers start to adopt effective remedies.

6 Discussions and Future Work

This study reveals the need to incorporate the equality of opportunity metric into the travel demand modeling framework, which is especially important when forecasting mode split under various transportation service deployment scenarios. It is also important when predicting the market share of new travel modes, such as automated vehicles, electric mobility and mobility-as-a-service packages. Due to the discriminatory behaviors exhibited by the prediction algorithms, socially disadvantaged groups may be under-served by the providers of these services. In the context of personalized on-demand mobility service, the preferences of the passengers towards various travel mode alternatives are captured by the choice modeling results, and the operators incorporate these predicted choice probabilities into their optimization framework for vehicle fleet allocation in order to maximize their profits [70, 71]. Given the potential existence of computational bias, the passengers from the disadvantaged group may not get access to their desired travel mode, which consequently exacerbates the social inequity problem.

Inherent and modeling biases are only two sources of prediction disparities. However, future studies should investigate other factors that also lead to prediction disparity and pinpoint more granular sources of bias in travel predictions. Previous literature has shown that prediction disparities can stem from various sources such as sampling procedure, feature quality and algorithms, but no study has examined these factors in the context of travel behavior modeling [50, 72, 61]. While we use equality of opportunity as the fairness metric in this research, scholars can also consider other metrics to measure computational fairness. Another future research direction is to study fairness issues with machine learning models other than DNN and BLR. In addition, while we try to address computational fairness through an algorithmic approach, future studies can explore how to combine our computational perspective with other remedies such as taking a participatory effort to evaluate fairness in the modeling results for various transportation programs.

This study demonstrates that adjusting BLR and DNN models to reduce prediction disparity may lead to lower prediction accuracy, but how to trade off accuracy and fairness essentially involves a value judgement, on which the authors do not take a stance. In fact, this needs to be analyzed on the basis of different normative principles, and may involve some ethical debates about what should be considered “fair” in a given context. In addition, mitigating unintended prediction biases can involve more costs. Policy makers should

determine whether the gain in computational fairness can outweigh the efforts of improving computational fairness, such as employing bias mitigation algorithms or collecting data that represent socially marginalized groups. Further cost and benefit analysis is necessary to bridge the gap between algorithmic fairness tools and social justice needs in the real world.

7 Acknowledgements

We thank Singapore-MIT Alliance for Research and Technology (SMART) for partially funding this research.

8 Author Contribution Statement

The authors confirm contribution to the paper as follows: study conception and design: Y.Z., S.W., J.Z.; data collection: Y.Z., S.W.; analysis and interpretation of results: Y.Z.; draft manuscript preparation: Y.Z., S.W.; research supervision: J.Z. All authors reviewed the results and approved the final version of the manuscript.

9 Conflict of Interests

The authors do not have any conflicts of interest to declare.

References

- [1] Paredes M, Hemberg E, O'Reilly UM, Zegras C. Machine learning or discrete choice models for car ownership demand estimation and prediction? In: 2017 5th IEEE International Conference on Models and Technologies for Intelligent Transportation Systems (MT-ITS). IEEE; 2017. p. 780–785.
- [2] Cheng L, Chen X, De Vos J, Lai X, Witlox F. Applying a random forest method approach to model travel mode choice behavior. *Travel behaviour and society*. 2019;14:1–10.
- [3] Fernández-Delgado M, Cernadas E, Barro S, Amorim D. Do we need hundreds of classifiers to solve real world classification problems? *The journal of machine learning research*. 2014;15(1):3133–3181.

- [4] Karlaftis MG, Vlahogianni EI. Statistical methods versus neural networks in transportation research: Differences, similarities and some insights. *Transportation Research Part C: Emerging Technologies*. 2011;19(3):387–399.
- [5] Wang S, Wang Q, Zhao J. Deep neural networks for choice analysis: Extracting complete economic information for interpretation. *Transportation Research Part C: Emerging Technologies*. 2020;118:102701.
- [6] Krizhevsky A, Sutskever I, Hinton GE. Imagenet classification with deep convolutional neural networks. In: *Advances in neural information processing systems*; 2012. p. 1097–1105.
- [7] Goodfellow I, Bengio Y, Courville A. *Deep learning*. MIT press; 2016.
- [8] Angwin J, Larson J, Mattu S, Kirchner L. Machine bias. *ProPublica*, May. 2016;23:2016.
- [9] Sweeney L. Discrimination in Online Ad Delivery. *ACM Queue*. 2013.
- [10] Bolukbasi T, Chang KW, Zou JY, Saligrama V, Kalai AT. Man is to computer programmer as woman is to homemaker? debiasing word embeddings. In: *Advances in neural information processing systems*; 2016. p. 4349–4357.
- [11] Ramjerdi F. Equity Measures and Their Performance in Transportation. *Transportation Research Record*. 2006;1983(1):67–74. Available from: <https://doi.org/10.1177/0361198106198300110>.
- [12] Litman T. Evaluating transportation equity. *World Transport Policy & Practice*. 2002;8(2):50–65.
- [13] Martens K. Substance precedes methodology: on cost–benefit analysis and equity. *Transportation*. 2011;38:959–974.
- [14] Beutel A, Chen J, Doshi T, Qian H, Woodruff A, Luu C, et al. Putting fairness principles into practice: Challenges, metrics, and improvements. In: *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*; 2019. p. 453–459.
- [15] The Department of Transportation. The Department of Transportation Title VI Program; 2019. Available from: <https://www.transportation.gov/mission/departement-transportation-title-vi-program>.

- [16] Cui Y, He Q, Khani A. Travel Behavior Classification: An Approach with Social Network and Deep Learning. *Transportation Research Record*. 2018;2672:68–80.
- [17] Wang S, Mo B, Zhao J. Deep neural networks for choice analysis: Architecture design with alternative-specific utility functions. *Transportation Research Part C: Emerging Technologies*. 2020;112:234–251.
- [18] Allahviranloo M, Recker W. Daily activity pattern recognition by using support vector machines with multiple classes. *Transportation Research Part B: Methodological*. 2013;58:16–43.
- [19] Cantarella GE, de Luca S. Multilayer feedforward networks for transportation mode choice analysis: An analysis and a comparison with random utility models. *Transportation Research Part C: Emerging Technologies*. 2005;13(2):121–155.
- [20] Zhang Z, He Q, Gao J, Ni M. A deep learning approach for detecting traffic accidents from social media data. *Transportation research part C: emerging technologies*. 2018;86:580–596.
- [21] van Cranenburgh S, Alwosheel A. An artificial neural network based approach to investigate travellers’ decision rules. *Transportation Research Part C: Emerging Technologies*. 2019;98:152–166.
- [22] Huang X, Sun J, Sun J. A car-following model considering asymmetric driving behavior based on long short-term memory neural networks. *Transportation research part C: emerging technologies*. 2018;95:346–362.
- [23] Polson NG, Sokolov VO. Deep learning for short-term traffic flow prediction. *Transportation Research Part C: Emerging Technologies*. 2017;79:1–17.
- [24] Wu Y, Tan H, Qin L, Ran B, Jiang Z. A hybrid deep learning based traffic flow prediction method and its understanding. *Transportation Research Part C: Emerging Technologies*. 2018;90:166–180.
- [25] Cheng L, Chen X, Vos JD, Lai X, Witlox F. Applying a random forest method approach to model travel mode choice behavior. *Travel Behaviour and Society*. 2019;14:1–10.
- [26] Wang S, Wang Q, Bailey N, Zhao J. Deep neural networks for choice analysis: A statistical learning theory perspective. *Transportation Research Part B: Methodological*. 2021;148:60–81.

- [27] Omrani H. Predicting Travel Mode of Individuals by Machine Learning. *Transportation Research Procedia*. 2015;840–849.
- [28] Karlaftis MG, Vlahogianni EI. Statistical methods versus neural networks in transportation research: Differences, similarities and some insights. *Transportation Research Part C: Emerging Technologies*. 2011;19:387–399.
- [29] Sekhar CR, Minal, Madhu E. Mode Choice Analysis Using Random Forrest Decision Trees. *Transportation Research Procedia*. 2016;17:644–652.
- [30] Chouldechova A. Fair prediction with disparate impact: A study of bias in recidivism prediction instruments. *Big data*. 2017;5(2):153–163.
- [31] Rajkomar A, Hardt M, Howell MD, Corrado G, Chin MH. Ensuring Fairness in Machine Learning to Advance Health Equity. *Annals of Internal Medicine*. 2018;169:866.
- [32] Goodman SN, Goel S, Cullen MR. Machine learning, health disparities, and causal reasoning. *Annals of internal medicine*. 2018;169(12):883–884.
- [33] Datta A, Tschantz MC, Datta A. Automated experiments on ad privacy settings: A tale of opacity, choice, and discrimination. *Proceedings on privacy enhancing technologies*. 2015;2015(1):92–112.
- [34] Dixon L, Li J, Sorensen J. Measuring and Mitigating Unintended Bias in Text Classification. *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*. 2018.
- [35] Deku SY, Kara A, Molyneux P. Access to consumer credit in the UK. *The European Journal of Finance*. 2016;22(10):941–964.
- [36] Lee MSA. Context-conscious fairness in using machine learning to make decisions. *AI Matters*. 2019;5(2):23–29.
- [37] Corbett-Davies S, Goel S. The Measure and Mismeasure of Fairness: A Critical Review of Fair Machine Learning. *CoRR*. 2018;abs/1808.00023. Available from: <http://arxiv.org/abs/1808.00023>.
- [38] Gajane P. On formalizing fairness in prediction with machine learning. *CoRR*. 2017;abs/1710.03184. Available from: <http://arxiv.org/abs/1710.03184>.
- [39] Friedler SA, Scheidegger C, Venkatasubramanian S, Choudhary S, Hamilton EP, Roth D. A comparative study of fairness-enhancing interventions in machine learning. In:

- Proceedings of the Conference on Fairness, Accountability, and Transparency - FAT* '19. ACM Press; 2019. p. 329–338. Available from: <http://dl.acm.org/citation.cfm?doid=3287560.3287589>.
- [40] Kamishima T, Akaho S, Sakuma J. Fairness-aware Learning through Regularization Approach. In: 2011 IEEE 11th International Conference on Data Mining Workshops; 2011. p. 643–650.
 - [41] Zhang BH, Lemoine B, Mitchell M. Mitigating Unwanted Biases with Adversarial Learning. In: Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society. ACM; 2018. p. 335–340. Available from: <https://dl.acm.org/doi/10.1145/3278721.3278779>.
 - [42] Kenneth Holstein, Jennifer Wortman Vaughan, Hal Daumé, Miro Dudik, Hanna Wallach. Improving Fairness in Machine Learning Systems: What Do Industry Practitioners Need? In: Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems - CHI '19. ACM Press; 2019. p. 1–16. Available from: <http://dl.acm.org/citation.cfm?doid=3290605.3300830>.
 - [43] Gunduz D, de Kerret P, Sidiropoulos ND, Gesbert D, Murthy CR, van der Schaar M. Machine Learning in the Air. *IEEE Journal on Selected Areas in Communications*. 2019;37(10):2184–2199. Available from: <https://ieeexplore.ieee.org/document/8839651/>.
 - [44] Bills TS. Looking beyond the mean for equity analysis: Examining distributional impacts of transportation improvements. *Transport Policy*. 2017;54:61–69. Available from: <https://doi.org/10.1016/j.tranpol.2016.08.003>.
 - [45] Golub A. Welfare and Equity Impacts of Gasoline Price Changes under Different Public Transportation Service Levels. *Journal of Public Transportation*. 2010;13:1–21.
 - [46] Linovski O, Baker DM, Manaugh K. Equity in practice? Evaluations of equity in planning for bus rapid transit. *Transportation Research Part A: Policy and Practice*. 2018;113:75–87.
 - [47] Delbosc A, Currie G. Using Lorenz curves to assess public transport equity. *Journal of Transport Geography*. 2011;19:1252–1259.
 - [48] Schweitzer L, Valenzuela A. Environmental Injustice and Transportation: The Claims and the Evidence. *Journal of Planning Literature*. 2004;18:383–398.

- [49] Martens K, Golub A, Robinson G. A justice-theoretic approach to the distribution of transportation benefits: Implications for transportation planning practice in the United States. *Transportation Research Part A: Policy and Practice*. 2012;46:684–695.
- [50] Barocas S, Selbst AD. Big data’s disparate impact. *Calif L Rev*. 2016;104:671.
- [51] Mehrabi N, Morstatter F, Saxena N, Lerman K, Galstyan A. A survey on bias and fairness in machine learning. *arXiv preprint arXiv:190809635*. 2019.
- [52] Hardt M, Price E, Srebro N. Equality of Opportunity in Supervised Learning. *Advances in Neural Information Processing Systems*. 2016:3315–3323.
- [53] Beutel A, Chen J, Zhao Z, Chi EH. Data decisions and theoretical implications when adversarially learning fair representations. *arXiv preprint arXiv:170700075*. 2017.
- [54] Binns R. Fairness in machine learning: Lessons from political philosophy. In: *Conference on Fairness, Accountability and Transparency*; 2018. p. 149–159.
- [55] Zafar MB, Valera I, Ródriguez MG, Gummadi KP. Fairness Constraints: Mechanisms for Fair Classification. In: Singh A, Zhu J, editors. *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*. vol. 54 of *Proceedings of Machine Learning Research*. Fort Lauderdale, FL, USA: PMLR; 2017. p. 962–970. Available from: <http://proceedings.mlr.press/v54/zafar17a.html>.
- [56] Grgic-Hlaca N, Zafar MB, Gummadi KP, Weller A. The case for process fairness in learning: Feature selection for fair decision making. In: *NIPS Symposium on Machine Learning and the Law*. vol. 1; 2016. p. 2.
- [57] Dwork C, Hardt M, Pitassi T, Reingold O, Zemel R. Fairness through awareness. In: *Proceedings of the 3rd innovations in theoretical computer science conference*; 2012. p. 214–226.
- [58] Kusner MJ, Loftus J, Russell C, Silva R. Counterfactual fairness. In: *Advances in neural information processing systems*; 2017. p. 4066–4076.
- [59] John PG, Vijaykeerthy D, Saha D. Verifying Individual Fairness in Machine Learning Models. *arXiv preprint arXiv:200611737*. 2020.
- [60] Wadsworth C, Vera F, Piech C. Achieving Fairness through Adversarial Learning: an Application to Recidivism Prediction. *arXiv preprint arXiv:180700199*. 2018.

- [61] Zliobaite I. Fairness-aware machine learning: a perspective. arXiv preprint arXiv:170800754. 2017.
- [62] Yao S, Huang B. Beyond parity: Fairness objectives for collaborative filtering. arXiv preprint arXiv:170508804. 2017.
- [63] Gu X, Angelov PP, Soares EA. A self-adaptive synthetic over-sampling technique for imbalanced classification. *International Journal of Intelligent Systems*. 2020;35(6):923–943.
- [64] Ferryman K, Pitcan M. Fairness in precision medicine. *Data & Society*. 2018;1.
- [65] Calders T, Žliobaitė I. Why unbiased computational processes can lead to discriminative decision procedures. In: *Discrimination and privacy in the information society*. Springer; 2013. p. 43–57.
- [66] Cochran N. Annual update of the HHS poverty guidelines. Department of Health and Human Services. 2017.
- [67] Byon E, Shrivastava AK, Ding Y. A classification procedure for highly imbalanced class sizes. *IIE Transactions*. 2010;42(4):288–303.
- [68] Federal Highway Administration. 2017 NHTS Data User Guide. US Department of Transportation Rockville, MD; 2018.
- [69] of Health UD, Services H, et al.. Office of the Assistant Secretary for Planning and Evaluation.(2019, January 12). 2019 poverty guidelines; 2019.
- [70] Atasoy B, Ikeda T, Song X, Ben-Akiva ME. The concept and impact analysis of a flexible mobility on demand system. *Transportation Research Part C: Emerging Technologies*. 2015;56:373–392.
- [71] Danaf M, Atasoy B, De Azevedo CL, Ding-Mastera J, Abou-Zeid M, Cox N, et al. Context-aware stated preferences with smartphone-based travel surveys. *Journal of choice modelling*. 2019;31:35–50.
- [72] Suresh H, Gutttag JV. A framework for understanding unintended consequences of machine learning. arXiv preprint arXiv:190110002. 2019.

Appendix

A The Synthetic Data Generation Process

The label y is generated based on the following equations:

$$U_i = V_i + \varepsilon_i = f(x_i, \mathbf{k}_i) + \varepsilon_i \quad (13)$$

$$Pr(y_i = 1) = \frac{1}{1 + \exp(-V_i)} \quad (14)$$

Equation 13 denotes the true utility function specification, where ε is the extreme value distributed random term. Equation 14 is the Sigmoid function calculating the probability of the binary outcomes. In simulations, for each individual i we draw y_i from a binomial distribution which takes value 1 with $Pr(y_i = 1)$. Note that this is where sampling errors may occur.

$$(a_i, x_i) \sim N\left(0, \begin{pmatrix} 1 & Cov_{ax} \\ Cov_{ax} & 1 \end{pmatrix}\right) \quad (15)$$

$$z_i = \begin{cases} 1, & \text{if } a_i \geq 0 \\ 0, & \text{if } a_i < 0 \end{cases} \quad (16)$$

$$\mathbf{k}_i \sim N(0, I) \quad (17)$$

z , x and \mathbf{k} are derived based on Equation 15, 16 and 17. a is an intermediate variable for the creation of z . a , x and all of the elements in \mathbf{k} are drawn from a multivariate Gaussian distribution with zero-mean and unit-variance. The above variables drawn from this distribution are all independent, except that a , x are correlated with the covariance being Cov_{ax} . z takes value 1 if $a \geq 0$ and takes value 0 otherwise. Therefore, the mean value of z is 0.5, indicating that the number of positive values and negative values of z is approximately the same. In our simulations, the value of Cov_{ax} varies across 0, 0.25, 0.5, 0.75, 1. The corresponding values of $Cov(z, x)$ are 0, 0.1, 0.2, 0.3, 0.4, each of which is empirically calculated as the value of $Cov(z, x)$ averaged across the 3 independently generated datasets used in the simulations for each Cov_{ax} .

$f(x_i, \mathbf{k}_i)$ in Equation 13 indicates the true systematic utility function. We consider two

scenarios:

$$f(x_i, \mathbf{k}_i) = \alpha + x_i\beta_{x_1} + \mathbf{k}_i^\top \boldsymbol{\beta}_{\mathbf{k}_1} \quad (18)$$

$$f(x_i, \mathbf{k}_i) = \alpha + x_i\beta_{x_1} + x_i^2\beta_{x_2} + \mathbf{k}_i^\top \boldsymbol{\beta}_{\mathbf{k}_1} + (\mathbf{k}_i \odot \mathbf{k}_i)^\top \boldsymbol{\beta}_{\mathbf{k}_2} \quad (19)$$

In Scenario 1, the utility specification takes a linear form (Equation 18). We set α to 0 and β_{x_1} to 1. Each entry of $\boldsymbol{\beta}_{\mathbf{k}_1}$ takes $\{-0.5, 0.5\}$ values with equal probabilities. As such, x will have the strongest positive influence on V compared with other explanatory variables.

In Scenario 2, the utility specification takes a quadratic form (Equation 19). We set α to -0.5, β_{x_1} to 1 and β_{x_2} to 0.5. Each entry of $\boldsymbol{\beta}_{\mathbf{k}_1}$ and $\boldsymbol{\beta}_{\mathbf{k}_2}$ takes $\{-0.5, 0.5\}$ values with equal probabilities.

This setup for both scenarios makes sure that $E(V)$ is equal to 0, so the average numbers of outcomes y taking the value 1 and 0 are approximately the same.

B Descriptive Statistics (NHTS Survey)

1. *Dependent variables: travel behavior and attitude indicators*

- Work from home: "Do you usually work from home?" (Yes=1)
- Work from home option: "Do you have the option of working from home?" (Yes=1)
- Travel burden: "Do you agree that travel is a financial burden?" ("Strongly agree" or "Agree"= 1)
- Gas price impact: "Do your agree that gas price affects travel?" ("Strongly agree" or "Agree"= 1)

2. *Dependent variables: variables regarding travel mode usage*

- Frequent usage of bike: "Do you use bike for travel daily or a few times a week?" (Yes=1)
- Frequent usage of car: "Do you use car for travel daily or a few times a week?" (Yes=1)
- Frequent usage of bus: "Do you use bus for travel daily or a few times a week?" (Yes=1)

- Frequent usage of rideshare: "Do you use rideshare at least once in the past 30 days?" (Yes=1)

Variable	Mean	Std.	Min	Median	Max
<i>Socio-demographics:</i>					
Age	45.834	17.372	6.0	45.0	92.0
*Gender (Male=1)	0.479	0.500	0.0	0.0	1.0
*Ethnic minority (Yes=1)	0.240	0.427	0.0	0.0	1.0
*Low income household (Yes=1)	0.130	0.337	0.0	0.0	1.0
*Medical condition (Yes=1)	0.065	0.246	0.0	0.0	1.0
Driver status (Yes=1)	0.915	0.280	0.0	1.0	1.0
Education: Less than high school (Yes=1)	0.068	0.252	0.0	0.0	1.0
Education: High school graduate (Yes=1)	0.190	0.392	0.0	0.0	1.0
Education: College degree (Yes=1)	0.293	0.455	0.0	0.0	1.0
Education: Bachelor's degree (Yes=1)	0.245	0.430	0.0	0.0	1.0
Home ownership (Yes=1)	0.661	0.473	0.0	1.0	1.0
Primary activity: absent (Yes=1)	0.027	0.163	0.0	0.0	1.0
Primary activity: homemaker (Yes=1)	0.073	0.260	0.0	0.0	1.0
Primary activity: unemployed (Yes=1)	0.035	0.184	0.0	0.0	1.0
Primary activity: retired (Yes=1)	0.158	0.365	0.0	0.0	1.0
Primary activity: going to school (Yes=1)	0.067	0.250	0.0	0.0	1.0
Born in the U.S. (Yes=1)	0.861	0.346	0.0	1.0	1.0
More than one job (Yes=1)	0.075	0.263	0.0	0.0	1.0
Health level (Excellent=1, Poor=5)	2.142	0.960	1.0	2.0	5.0
Job: Sales or service (Yes=1)	0.168	0.373	0.0	0.0	1.0
Job: Clerical or administrative support (Yes=1)	0.070	0.256	0.0	0.0	1.0
Job: Manufacturing type (Yes=1)	0.087	0.282	0.0	0.0	1.0
Work for pay (Yes=1)	0.069	0.253	0.0	0.0	1.0
Level of physical activity (Never/rarely=1, Vigorous=3)	2.173	0.594	1.0	2.0	3.0
In public or private school (Yes=1)	0.024	0.154	0.0	0.0	1.0
Full-time worker (Yes=1)	0.504	0.500	0.0	1.0	1.0
<i>Household Variables:</i>					
Number of drivers in the HH	2.038	0.945	0.0	2.0	9.0

Note: the sample weights are used to compute the summary statistics; (*) denotes the independent variables that are also treated as the protected variables based on which we examine the prediction disparities.

Table B1: Summary statistics of the explanatory and dependent variables

Variable	Mean	Std.	Min	Median	Max
Last year's household income (K \$)	84.009	64.924	5.0	62.5	250.0
Number of household members	2.915	1.435	1.0	3.0	13.0
Number of household vehicles	2.187	1.305	0.0	2.0	12.0
No child in the HH (Yes=1)	0.352	0.478	0.0	0.0	1.0
Youngest child's age < 15 in the HH (Yes=1)	0.343	0.475	0.0	0.0	1.0
Youngest child's age < 21 in the HH (Yes=1)	0.102	0.303	0.0	0.0	1.0
Number of workers in household	1.495	1.004	0.0	1.0	7.0
<i>Built Environment:</i>					
*Household in urban area (Yes=1)	0.836	0.371	0.0	1.0	1.0
Northeast Region (Yes=1)	0.178	0.382	0.0	0.0	1.0
Midwest Region (Yes=1)	0.217	0.412	0.0	0.0	1.0
Gas Price (\$)	2.395	0.208	2.0	2.4	3.0
% of renter-occupied housing in the block group	31.226	22.877	2.0	20.0	97.5
Housing units/sq.m in the block group (in thousands)	3.004	5.409	0.0	1.5	30.0
Second City (Yes=1)	0.201	0.401	0.0	0.0	1.0
Suburban (Yes=1)	0.235	0.424	0.0	0.0	1.0
Small Town (Yes=1)	0.195	0.396	0.0	0.0	1.0
Number of workers/sq.m in the census tract (in thousands)	1.751	1.703	0.0	1.5	5.0
MSA population > 1 million, without rail (Yes=1)	0.282	0.450	0.0	0.0	1.0
MSA population < 1 million (Yes=1)	0.299	0.458	0.0	0.0	1.0
Population size of the MSA (in millions)	1.959	1.643	0.0	2.0	4.0
In an urban area (Yes=1)	0.741	0.438	0.0	1.0	1.0
In an urban cluster (Yes=1)	0.095	0.293	0.0	0.0	1.0
Area surrounded by urban areas (Yes=1)	0.000	0.022	0.0	0.0	1.0
Urban area size (in millions)	1.018	0.910	0.0	0.8	2.0
<i>Travel Pattern and Internet Usage:</i>					
Flexible work time (Yes=1)	0.307	0.461	0.0	0.0	1.0
Travel day began at home location (Yes=1)	0.938	0.241	0.0	1.0	1.0
Frequent internet use (Yes=1)	0.955	0.208	0.0	1.0	1.0
<i>Dependent Variables:</i>					
Work from home option (Yes=1)	0.184	0.388	0.0	0.0	1.0
Work from home (Yes=1)	0.116	0.321	0.0	0.0	1.0
Travel is a financial burden (Yes=1)	0.396	0.489	0.0	0.0	1.0
Gas price affects travel (Yes=1)	0.478	0.500	0.0	0.0	1.0
Frequent usage of bike (Yes=1)	0.066	0.248	0.0	0.0	1.0
Frequent usage of bus (Yes=1)	0.063	0.243	0.0	0.0	1.0
Frequent usage of car (Yes=1)	0.926	0.262	0.0	1.0	1.0
Frequent usage of rideshare (Yes=1)	0.114	0.317	0.0	0.0	1.0

Note: the sample weights are used to compute the summary statistics; (*) denotes the independent variables that are also treated as the protected variables based on which we examine the prediction disparities.

Table B2: (Cont.) Summary statistics of the explanatory and dependent variables

Protected Variable		Dependent Variable			
Type	Value	WFH	WFHO	TB	GPI
Race	Minority	29.73%	31.25%	46.6%	56.59%
	Majority	34.02%	33.79%	33.02%	42.98%
Gender	Male	33.52%	37.67%	34.21%	43.66%
	Female	33.19%	28.92%	35.95%	46.37%
Income	Low	37.62%	15.53%	56.81%	67.23%
	Middle or High	33.1%	34.29%	33.14%	43.07%
Medical Condition	With	49.72%	31.54%	49.2%	57.53%
	Without	32.99%	33.41%	34.1%	44.19%
Region	Urban	32.71%	34.9%	33.63%	42.71%
	Rural	35.9%	26.76%	40.67%	53.85%
Sample Size		160110	204603	703647	703647
Number of Positives		53400	68296	247300	317432

Note: “WFH” stands for “work from home”, “WFHO” stands for “work from home option”, “TB” stands for “travel burden”, “GPI” stands for “gas price impact”; each percentage number indicates the proportion of positives to the total number of the corresponding dependent variable in the subset of the corresponding protected variable; since the outcome distributions of “WFH” and “WFHO” are highly skewed, for these two variables the data is balanced so that the ratio of the major outcome class to the minor outcome class instances is 2:1.

Table B3: Summary statistics of dependent variables by protected variables in the data set

Protected Variable		Frequent Usage of			
Type	Value	Bus	Bike	Car	Rideshare
Race	Minority	55.03%	30.91%	48.35%	37.28%
	Majority	27.1%	33.83%	71.28%	32.56%
Gender	Male	33.03%	35.99%	66.65%	35.09%
	Female	33.67%	30.96%	66.6%	31.75%
Income	Low	62.76%	39.06%	29.27%	20.76%
	Middle or High	28.54%	32.82%	73.98%	34.34%
Medical Condition	With	46.37%	26.01%	38.92%	13.31%
	Without	32.21%	33.88%	69.98%	34.53%
Region	Urban	37.81%	35.8%	63.82%	37.96%
	Rural	10.79%	23.09%	78.85%	11.43%
Sample Size		72235	130939	95323	181132
Number of Positives		24106	43712	63504	60378

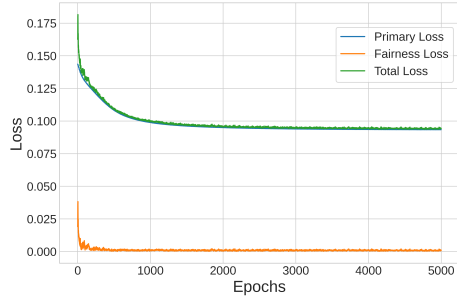
Note: since the distributions of the travel mode usage are highly skewed, for each dependent variable the data is balanced so that the ratio of the major outcome class to the minor outcome class instances is 2:1.

Table B4: (Cont.) Summary statistics of dependent variables by protected variables in the data set

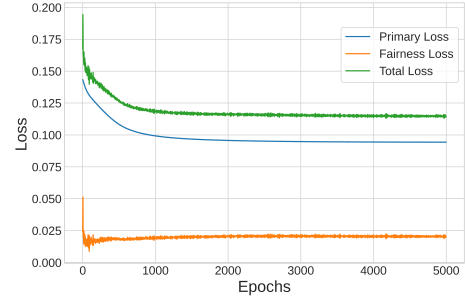
C Convergence of Loss Values in the Training Process

The convergence of loss values during the training process is shown in Figure C1, where we try to mitigate the prediction disparity between rural and urban residents when predicting frequent usage of rideshare using DNN. The bias mitigation weight is 0.8. Primary loss refers to $(1 - \lambda)$ times the cross-entropy loss which is used to increase the prediction accuracy, whereas fairness loss refers to λ times the correlation loss which is applied to mitigate the prediction bias. The total loss is computed as the sum of the primary and the fairness loss.

From Figure C1a, we can see that in the first few epochs, the reduction of total loss is mainly driven by the reduction of fairness loss as the values of both losses drop sharply. As the fairness loss drops to nearly zero, the algorithm then primarily tries to mitigate the prediction loss, while the fairness loss remains very small in the following training steps. The testing data shows similar trends of loss values during the training process (Figure C1b), except that the fairness loss never drops to the near-zero level. All in all, Figure C1 illustrates how our bias mitigation method works, and how the algorithm manages to



(a) Training data



(b) Testing data

Figure C1: Change of loss values in the training process. Protected variable: urban-rural divide; dependent variable: the frequent usage of rideshare; mitigation weight (λ): 0.8.

substantially reduce fairness loss at early stage of training.

D Sensitivity Analysis

Model	FNR gap				FPR gap			
	Bike	Bus	Car	Rideshare	TB	GPI	WFH	WFHO
(Sensitive Attribute = Race)								
Main: logit	0.121 (0.018)	-0.228 (0.02)	0.14 (0.016)	0.047 (0.017)	0.357 (0.012)	0.303 (0.015)	-0.004 (0.01)	-0.012 (0.008)
Main: DNN	-0.027 (0.007)	-0.02 (0.007)	0.058 (0.013)	0.002 (0.009)	0.071 (0.007)	0.056 (0.01)	-0.011 (0.005)	-0.01 (0.005)
2 layers DNN	-0.024 (0.013)	-0.032 (0.007)	0.053 (0.017)	-0.003 (0.008)	0.108 (0.01)	0.091 (0.015)	-0.019 (0.005)	-0.02 (0.005)
4 layers DNN	-0.022 (0.008)	-0.013 (0.01)	0.053 (0.009)	-0.002 (0.01)	0.046 (0.008)	0.036 (0.011)	-0.007 (0.006)	-0.012 (0.005)
Batch size=5000	-0.033 (0.016)	-0.046 (0.009)	0.041 (0.005)	0.011 (0.019)	0.041 (0.009)	0.048 (0.006)	-0.009 (0.004)	-0.035 (0.006)
He initialization	-0.025 (0.01)	-0.024 (0.008)	0.053 (0.022)	-0.001 (0.008)	0.066 (0.008)	0.057 (0.004)	-0.011 (0.004)	-0.014 (0.007)
(Sensitive Attribute = Gender)								
Main: logit	0.127 (0.026)	-0.014 (0.008)	-0.004 (0.011)	0.022 (0.032)	0.02 (0.01)	0.047 (0.006)	0.029 (0.006)	-0.074 (0.008)
Main: DNN	0.017 (0.01)	-0.002 (0.006)	-0.001 (0.011)	0.01 (0.009)	0.008 (0.004)	0.015 (0.007)	-0.001 (0.006)	-0.022 (0.005)
2 layers DNN	0.033 (0.015)	-0.003 (0.01)	-0.004 (0.003)	0.011 (0.01)	0.012 (0.005)	0.018 (0.007)	0.003 (0.007)	-0.029 (0.007)
4 layers DNN	0.009 (0.013)	-0.002 (0.006)	-0.003 (0.005)	0.008 (0.009)	0.007 (0.004)	0.011 (0.004)	-0.003 (0.008)	-0.016 (0.004)
Batch size=5000	0.031 (0.013)	-0.007 (0.007)	0.002 (0.008)	0.016 (0.015)	-0.001 (0.003)	0.006 (0.004)	-0.004 (0.004)	-0.063 (0.008)
He initialization	0.015 (0.008)	-0.004 (0.008)	-0.003 (0.01)	0.011 (0.008)	0.012 (0.003)	0.016 (0.003)	-0.002 (0.006)	-0.021 (0.004)
(Sensitive Attribute = Income)								
Main: logit	-0.09 (0.039)	-0.216 (0.009)	0.353 (0.022)	0.25 (0.023)	0.605 (0.013)	0.564 (0.003)	0.053 (0.014)	-0.139 (0.011)
Main: DNN	-0.051 (0.009)	-0.019 (0.009)	0.133 (0.018)	0.003 (0.011)	0.101 (0.011)	0.084 (0.016)	-0.006 (0.011)	-0.049 (0.007)
2 layers DNN	-0.076 (0.009)	-0.03 (0.006)	0.13 (0.012)	0.018 (0.019)	0.165 (0.021)	0.15 (0.02)	-0.012 (0.011)	-0.068 (0.009)
4 layers DNN	-0.031 (0.004)	-0.014 (0.01)	0.127 (0.017)	-0.007 (0.01)	0.072 (0.005)	0.054 (0.007)	-0.002 (0.014)	-0.036 (0.007)
Batch size=5000	-0.087 (0.012)	-0.044 (0.014)	0.092 (0.018)	0.037 (0.034)	0.042 (0.01)	0.046 (0.011)	-0.001 (0.012)	-0.119 (0.016)
He initialization	-0.041 (0.005)	-0.02 (0.008)	0.114 (0.028)	0.002 (0.006)	0.089 (0.007)	0.076 (0.009)	-0.006 (0.007)	-0.046 (0.005)

Note:

1. means and standard deviations (in parentheses) of the FNR gap/FPR gap and accuracy of 5-fold cross-validation trained with different models.
2. "Main: logit" is the BLR model, "Main: DNN" is the default DNN model tested in Section 4.2. For this default DNN model, the number of hidden layers is 3, and it is trained with a batch size of 500,000 and the Xavier uniform weight initialization. The other four models are adapted from the default DNN model: "2 layers DNN" and "4 layers DNN" change the number of DNN hidden layers to 2 and 4 while keeping other hyperparameters fixed. "Batch size=5000" changes the batch size to 5000. "He initialization" change the weight initializer to He normal initializer.

Table D1: Fairness and accuracy results of different models

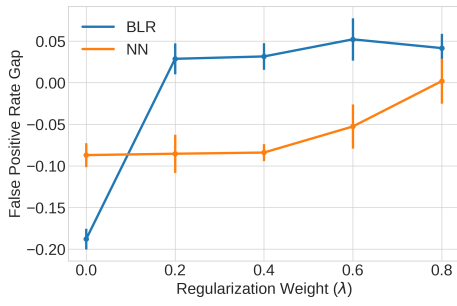
Model	FNR gap				FPR gap			
	Bike	Bus	Car	Rideshare	TB	GPI	WFH	WFHO
(Sensitive Attribute = Medical Condition)								
Main: logit	-0.006 (0.054)	-0.14 (0.023)	0.253 (0.044)	0.433 (0.055)	0.425 (0.02)	0.321 (0.025)	0.137 (0.057)	-0.106 (0.009)
Main: DNN	-0.02 (0.026)	-0.016 (0.006)	0.091 (0.027)	0.043 (0.025)	0.117 (0.014)	0.083 (0.012)	0.065 (0.042)	-0.028 (0.014)
2 layers DNN	-0.026 (0.024)	-0.029 (0.005)	0.089 (0.033)	0.061 (0.026)	0.154 (0.02)	0.106 (0.01)	0.053 (0.023)	-0.023 (0.01)
4 layers DNN	-0.011 (0.023)	-0.018 (0.005)	0.075 (0.01)	0.022 (0.02)	0.086 (0.01)	0.064 (0.01)	0.065 (0.04)	-0.017 (0.02)
Batch size=5000	-0.01 (0.034)	-0.037 (0.006)	0.057 (0.019)	0.116 (0.069)	0.055 (0.012)	0.05 (0.008)	0.05 (0.031)	-0.076 (0.018)
He initialization	-0.016 (0.031)	-0.017 (0.01)	0.082 (0.027)	0.042 (0.026)	0.115 (0.022)	0.075 (0.013)	0.057 (0.03)	-0.025 (0.009)
(Sensitive Attribute = Region)								
Main: logit	0.356 (0.016)	0.695 (0.046)	-0.068 (0.012)	0.624 (0.018)	0.073 (0.01)	0.244 (0.008)	0.068 (0.011)	-0.073 (0.017)
Main: DNN	0.076 (0.023)	0.116 (0.027)	-0.026 (0.014)	0.146 (0.028)	0.041 (0.005)	0.082 (0.01)	0.009 (0.008)	-0.021 (0.01)
2 layers DNN	0.088 (0.027)	0.168 (0.026)	-0.018 (0.009)	0.222 (0.04)	0.053 (0.004)	0.112 (0.013)	0.018 (0.009)	-0.03 (0.013)
4 layers DNN	0.053 (0.012)	0.079 (0.033)	-0.027 (0.007)	0.089 (0.02)	0.037 (0.008)	0.063 (0.004)	0.01 (0.005)	-0.018 (0.005)
Batch size=5000	0.112 (0.047)	0.173 (0.027)	-0.022 (0.006)	0.391 (0.056)	0.026 (0.015)	0.05 (0.017)	0.009 (0.005)	-0.039 (0.013)
He initialization	0.073 (0.01)	0.13 (0.039)	-0.023 (0.012)	0.136 (0.03)	0.038 (0.007)	0.079 (0.008)	0.017 (0.006)	-0.023 (0.003)
Model	Accuracy							
	Bike	Bus	Car	Rideshare	TB	GPI	WFH	WFHO
Main: logit	0.697 (0.006)	0.778 (0.007)	0.84 (0.004)	0.777 (0.008)	0.645 (0.002)	0.652 (0.003)	0.763 (0.003)	0.791 (0.003)
Main: DNN	0.955 (0.002)	0.977 (0.001)	0.98 (0.001)	0.946 (0.001)	0.854 (0.004)	0.856 (0.002)	0.952 (0.001)	0.944 (0.001)
2 layers DNN	0.928 (0.001)	0.964 (0.001)	0.968 (0.001)	0.923 (0.001)	0.795 (0.003)	0.796 (0.002)	0.928 (0.002)	0.919 (0.002)
4 layers DNN	0.971 (0.001)	0.984 (0.0005)	0.986 (0.001)	0.966 (0.001)	0.893 (0.003)	0.893 (0.003)	0.971 (0.001)	0.966 (0.001)
Batch size=5000	0.949 (0.004)	0.965 (0.002)	0.969 (0.001)	0.866 (0.007)	0.896 (0.004)	0.897 (0.004)	0.94 (0.007)	0.867 (0.007)
He initialization	0.959 (0.001)	0.978 (0.001)	0.979 (0.0004)	0.952 (0.001)	0.871 (0.004)	0.867 (0.005)	0.959 (0.002)	0.951 (0.001)

Table D2: (Cont.) Fairness and accuracy results of different models

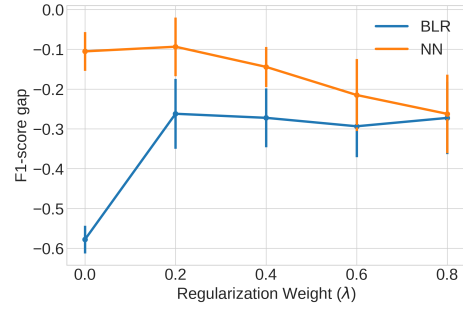
E Other Fairness Measures after Bias Mitigation

Figure E1a indicates that when mitigating the absolute value of the FNR gap in terms of “region” for the frequent rideshare usage prediction, the absolute value of the FPR gap also decreases (i.e., the value becomes closer to zero). Similarly, Figure E2a shows that when mitigating the absolute value of the FPR gap in terms of "race" for the travel burden prediction, the absolute value of the FNR gap also decreases.

In terms of the F1-score⁵, the results show that the bias mitigation also leads to drops in the absolute F1-score gaps in Figure E2b. However, in Figure E1b, the absolute F1-score gap only drops when BLR is applied, but not for DNN.

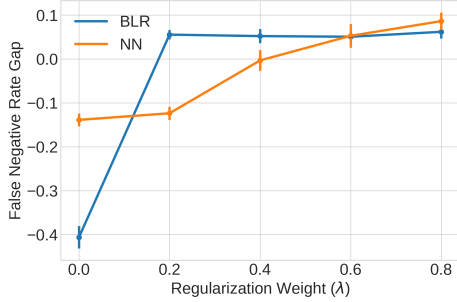


(a) FPR Gap vs. Regularization Weight for FNR

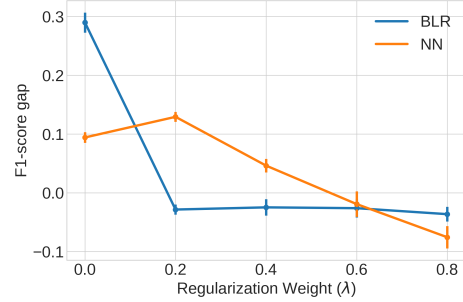


(b) F1-score Gap vs. Regularization Weight for FNR

Figure E1: FPR gap and F1-score gap by bias mitigation weight (λ) for FNR: regional gap in the prediction of frequent rideshare usage



(a) FNR gap vs. Regularization Weight for FPR



(b) F1-score gap vs. Regularization Weight for FPR

Figure E2: FNR gap and F1-score gap by bias mitigation weight (λ) for FPR: racial gap in the prediction of travel burden

⁵The F1-score conveys the balance between recall and precision. The F1-score gap is calculated as:

$$\text{F1-score gap} = \frac{TP_{z=0}}{TP_{z=0} + 0.5 \cdot (FP_{z=0} + FN_{z=0})} - \frac{TP_{z=1}}{TP_{z=1} + 0.5 \cdot (FP_{z=1} + FN_{z=1})}$$

F Effects of Bias Mitigation on Other Protected Variables

λ	FNR gap					Accuracy
	Race	Gender	Income	Medical Condition	Region	
0.0	0.047 (0.017)	0.022 (0.032)	0.25 (0.023)	0.433 (0.055)	0.624 (0.018)	0.777 (0.008)
0.2	0.069 (0.025)	0.049 (0.035)	0.295 (0.024)	0.47 (0.055)	-0.064 (0.048)	0.766 (0.001)
0.4	0.07 (0.025)	0.047 (0.035)	0.285 (0.037)	0.471 (0.056)	-0.047 (0.04)	0.766 (0.001)
0.6	0.077 (0.02)	0.055 (0.036)	0.325 (0.027)	0.488 (0.045)	-0.044 (0.047)	0.764 (0.002)
0.8	0.11 (0.02)	0.066 (0.033)	0.373 (0.028)	0.48 (0.047)	-0.061 (0.054)	0.759 (0.002)

Table F1: Change of FNR gaps for various protected variables with the increase of bias mitigation weight (λ) regarding "region" when predicting the frequent rideshare usage BLR.

λ	FPR gap					Accuracy
	Race	Gender	Income	Medical Condition	Region	
0.0	0.357 (0.012)	0.02 (0.01)	0.605 (0.013)	0.425 (0.02)	0.073 (0.01)	0.645 (0.002)
0.2	-0.0002 (0.012)	0.01 (0.008)	0.561 (0.022)	0.43 (0.019)	0.138 (0.01)	0.634 (0.001)
0.4	0.001 (0.012)	0.006 (0.01)	0.552 (0.019)	0.434 (0.018)	0.142 (0.015)	0.634 (0.001)
0.6	-0.001 (0.013)	0.002 (0.007)	0.514 (0.069)	0.448 (0.028)	0.168 (0.03)	0.632 (0.001)
0.8	-0.002 (0.011)	-0.004 (0.014)	0.424 (0.113)	0.437 (0.024)	0.224 (0.077)	0.627 (0.006)

Table F2: Change of FPR gaps for various protected variables with the increase of bias mitigation weight (λ) regarding "race" when predicting the travel burden using BLR.

λ	FNR gap					Accuracy
	Race	Gender	Income	Medical Condition	Region	
0.0	0.002 (0.009)	0.01 (0.009)	0.003 (0.011)	0.043 (0.025)	0.146 (0.028)	0.946 (0.001)
0.2	0.008 (0.008)	0.015 (0.01)	0.029 (0.011)	0.048 (0.029)	0.048 (0.032)	0.906 (0.002)
0.4	0.026 (0.017)	0.03 (0.01)	0.083 (0.035)	0.163 (0.039)	0.062 (0.031)	0.862 (0.003)
0.6	0.056 (0.019)	0.044 (0.028)	0.231 (0.038)	0.308 (0.061)	0.038 (0.051)	0.81 (0.003)
0.8	0.071 (0.021)	0.047 (0.038)	0.288 (0.03)	0.494 (0.069)	-0.009 (0.058)	0.778 (0.001)

Table F3: Change of FNR gaps for various protected variables with the increase of bias mitigation weight (λ) regarding "region" when predicting the frequent rideshare usage using 3-layer DNN.

λ	FPR gap					Accuracy
	Race	Gender	Income	Medical Condition	Region	
0.0	0.071 (0.007)	0.008 (0.004)	0.101 (0.011)	0.117 (0.014)	0.041 (0.005)	0.854 (0.004)
0.2	0.025 (0.007)	0.013 (0.005)	0.234 (0.014)	0.202 (0.016)	0.086 (0.011)	0.746 (0.007)
0.4	-0.004 (0.021)	0.007 (0.011)	0.469 (0.029)	0.352 (0.014)	0.148 (0.014)	0.664 (0.003)
0.6	-0.003 (0.02)	0.012 (0.011)	0.561 (0.016)	0.381 (0.031)	0.168 (0.012)	0.641 (0.001)
0.8	-0.014 (0.015)	0.008 (0.009)	0.459 (0.042)	0.328 (0.02)	0.2 (0.02)	0.633 (0.001)

Table F4: Change of FPR gaps for various protected variables with the increase of bias mitigation weight (λ) regarding "race" when predicting the travel burden using 3-layer DNN.

G Descriptive Statistics (the Chicago Travel Survey)

	Mean	Std.	Min	Median	Max
<i>Independent Variables:</i>					
*Ethnic minority (Yes=1)	0.221	0.415	0.0	0.0	1.0
*Low income household (Yes=1)	0.056	0.230	0.0	0.0	1.0
*Disability status (Yes=1)	0.024	0.152	0.0	0.0	1.0
Age	41.712	14.247	16.0	40.0	87.0
Education level (1-6: low-high)	4.410	1.480	1.0	5.0	6.0
Number of jobs	1.117	0.395	1.0	1.0	9.0
Have a valid driver's license (Yes=1)	0.944	0.230	0.0	1.0	1.0
Gender (female=1)	0.529	0.499	0.0	1.0	1.0
Count of household members	3.155	1.481	1.0	3.0	12.0
Count of household vehicles	2.052	1.201	0.0	2.0	10.0
Home ownership (Yes=1)	0.687	0.464	0.0	1.0	1.0
Count of years lived at current address	11.040	10.100	1.0	8.0	120.0
Distance to the work place (km)	17.191	33.700	0.0	11.2	1272.2
Last year's household income (K \$)	98.023	52.445	7.5	87.5	175.0
<i>Dependent Variable:</i>					
Travel mode to work (auto=1, else=0)	0.731	0.444	0.0	1.0	1.0
Sample size: 13,730					

Note: the sample weights are used to compute the summary statistics; (*) denotes the independent variables that are also treated as the protected variables based on which we examine the prediction disparities.

Table G1: Summary statistics of the explanatory and dependent variables in the Chicago travel survey dataset