A Burden Shared is a Burden Halved: A Fairness-Adjusted Approach to Classification

Bradley Rava¹, Wenguang Sun², Gareth M. James³ and Xin Tong⁴

Abstract

We investigate the fairness issue in classification, where automated decisions are made for individuals from different protected groups. In high-consequence scenarios, decision errors can disproportionately affect certain protected groups, leading to unfair outcomes. To address this issue, we propose a fairness-adjusted selective inference (FASI) framework and develop data-driven algorithms that achieve statistical parity by controlling the false selection rate (FSR) among protected groups. Our FASI algorithm operates by converting the outputs of black-box classifiers into R-values, which are both intuitive and computationally efficient. These R-values serve as the basis for selection rules that are provably valid for FSR control in finite samples for protected groups, effectively mitigating the unfairness in group-wise error rates. We demonstrate the numerical performance of our approach using both simulated and real data.

Keywords: Calibration by group; Fairness in machine learning; False selection rate; Selective Inference; Statistical parity.

1 Introduction

In a broad range of applications, artificial intelligence (AI) systems are rapidly replacing human decision-making. Many of these scenarios are sensitive in nature, where the AI's decision, correct or not, can directly impact one's social or economic status. A few examples include a bank determining credit card limits, stores using facial recognition systems to detect shoplifters, and hospitals attempting to identify which of their patients has a specific disorder. Unfortunately, despite their supposedly unbiased approach to decision-making, there has been increasing evidence that AI algorithms often fail to treat equally people of different genders, races, religions, or other protected attributes. Whether this is due to the historical bias in one's training data, or otherwise, it is important, for both legal and policy reasons, that we make ethical use of data and ensure that decisions are made fairly for everyone regardless of their protected attributes.

Despite the significant efforts in developing supervised learning algorithms to improve the prediction accuracy, making reliable and fair decisions in the classification setting remains a critical and challenging problem for two main reasons. Firstly, AI algorithms are often required to make classifications on all new observations without a careful assessment of associated uncertainty or ambiguity. This limitation highlights the need for a more flexible framework to handle

¹University of Sydney Business School.

²Center for Data Science, Zhejiang University. Author for correspondence: wgsun@zju.edu.cn

³Goizueta Business School, Emory University.

⁴Department of Data Sciences and Operations, University of Southern California.

intrinsically difficult classification tasks where a definitive decision carries high stakes. Such a framework should enable decision-makers to wait and gather additional information with greater confidence before making a final decision. Secondly, modern machine learning models, such as neural networks, are often highly complex, making it challenging, if not impossible, to explicitly quantify the uncertainty associated with their outputs or to provide guarantees on the fairness of the decisions. Therefore, developing methods that can ensure both risk control and fairness is crucial for AI systems to be reliable and trustworthy.

This article develops a "fairness-adjusted selective inference" (FASI) framework to address the critical issues of uncertainty assessment, error rate control and statistical parity in classification. We provide an *indecision* option for observations which cannot be selected into any classes with confidence. These observations will then be separately evaluated. This practice often aligns with the policy objectives in many real world scenarios. For example, incorrectly classifying a low-risk individual as a recidivist or rejecting a well-deserving candidate for the loan request is much more expensive than turning the case over for a more careful review. A mis-classification is an error, the probability of which must be controlled to be small as its consequence can be severe. By contrast, the cost of an indecision is usually much less. For example, the ambiguity can be mitigated by collecting additional contextual knowledge of the convicted individual or requesting more information from the loan applicant. Under the selective inference (Benjamini 2010) framework, we only make definitive decisions on a selected subset of all individuals; the less consequential indecision option is considered as a wasted opportunity rather than an error. A natural error rate notion under this framework is the False Selection Rate (FSR), which is defined as the expected fraction of erroneous classifications among the selected subset of individuals. The goal is to develop decision rules that aim to control and equalize the FSR across protected groups, while minimizing the total wasted opportunities.

A critical issue is that a classification rule that controls the overall FSR may have disparate impacts on different protected groups. We illustrate the point using the COMPAS data set (Angwin et al. 2016, Dieterich et al. 2016). The COMPAS algorithm has been widely used in the US to help inform courts about a defendant's recidivism likelihood, i.e., the likelihood of a convicted criminal recommitting a crime, so any prediction errors could have significant implications. The left panel of Figure 1 shows the *False Selection Proportions* (FSP), i.e. the fraction of individuals who did not recommit a crime among those who were classified as recidivists. The classification rule was constructed via a Generalized Additive Model (GAM)¹ (Hastie et al. 2009, James et al. 2023) to achieve the target FSR of 25%. We first split the COMPAS data into distinct training and test sets. The GAM was fitted using the training data set, and subsequently applied to the test set to predict whether a defendant was a recidivist.

We can see that the green bar, which provides the overall FSP for all races, is close to the target value. Moreover, the rule appears to be "fair" for all individuals, regardless of their protected attributes, in the sense that the *same* threshold has been applied to the confidence scores (i.e. estimated class probabilities) produced by the *same* GAM fit. However, the blue and orange bars show that the FSPs for different racial groups differ significantly from 25%, which is clearly not a desirable situation.

This article introduces a new notion of fairness that requires parity in FSR control across various protected groups. This aligns with the social and policy goals in a range of decisionmaking scenarios such as selecting recidivists or determining risky loan applicants, where the burden of erroneous classifications should be shared equally among different genders and races. However, the development of effective and fair FSR rules is challenging. First, controlling the error rate associated with a classifier, such as one built around the GAM procedure, critically

¹Although a GAM was utilized for illustration purposes, we emphasize that the same issue can arise regardless of the specific machine learning algorithm employed.



Figure 1: The selection of recidivists from a pool of criminal defendants (Broward County, Florida). The target FSR is 25%. Left: the unadjusted approach. Right: the proposed FASI approach.

depends on the accuracy of the scores. However, the assessment of the accuracy/uncertainty of these scores largely remains unknown. Second, we wish to provide practitioners with theoretical guarantees on the parity and validity for FSR control, regardless of the algorithm being used, including complex black-box classifiers.

To address these issues, we have developed a data-driven FASI algorithm specifically designed to control the FSRs of protected groups below a user-specified level α . The right panel of Figure 1 illustrates the FSPs of FASI on the recidivism data. All individual FSPs are controlled at 25% approximately. FASI works by converting the confidence scores from a black-box learning algorithm to an R-value, which is intuitive, easy to compute, and comparable across different protected groups. We then show that selecting all observations with R-value no greater than α will result in an FSR of approximately α . Hence, we can directly use this R-value to assign new observations a class label or, for observations with high R-values, assign them to the *indecision* class.

This paper makes several contributions. Firstly, we introduce a novel notion of fairness within the selective inference framework, incorporating an indecision option. In high-consequence situations, it is sensible to exercise caution, by either withholding or separately evaluating such cases until additional evidence is gathered. This reduces the risk of making definitive decisions without sufficient support, thus promoting cautious and fair decisions in these complex scenarios. Secondly, a data-driven FASI Algorithm is developed based on the utilization of the R-value. This algorithm, which can be deployed with user-specified learning algorithms (e.g. random forest, neural networks), is intuitively appealing and easy to interpret. Thirdly, rigorous theoretical justifications are provided for the FASI algorithm. The finite-sample theory on FSR control is established with minimal assumptions, accommodating scores generated by black-box algorithms. Finally, the empirical performance of FASI is investigated through extensive experimentation using simulated and real-world data sets, demonstrating the effectiveness and practical utility of the proposed approach.

The rest of the paper is structured as follows. In Section 2 we define the FSR and describe the problem formulation. Section 3 introduces the R-value and FASI algorithm. The numerical results for simulated and real data are presented in Sections 4 and 5, respectively. Section 6 concludes the main article with a discussion of related works and possible extensions. The Online Supplementary Material provides additional technical details about the methodology, proof of theorems, and supplementary numerical results.

2 Problem Formulation

Suppose we observe a data set $\{(X_i, A_i, Y_i) : i \in \mathcal{D}\}$, where $\mathcal{D} = [n] \equiv \{1, ..., n\}$ is an index set, $X_i \in \mathbb{R}^p$ is a *p*-dimensional vector of features, $A_i \in \mathcal{A}$ is an additional feature representing the protected or sensitive attribute, and Y_i is a class label taking values in $\mathcal{C} = \{1, ..., C\}$. The goal is to predict the classes for *m* new individuals indexed by $\mathcal{D}^{test} = \{n+1, ..., n+m\}$, with observed features $\{(X_j, A_j) : j \in \mathcal{D}^{test}\}$. The predicted values for their class labels $\{Y_j : j \in \mathcal{D}^{test}\}$ are denoted by $\{\hat{Y}_j : j \in \mathcal{D}^{test}\}$.

2.1 Background: predictive parity in classification

We focus on scenarios where an individual's membership to a particular protected group is known. Group-fairness approaches, which explicitly enforce fairness across groups, have been widely applied across various disciplines, ranging from medicine to the criminal justice system. To provide context for our fairness notion, we start with the widely used predictive parity or sufficiency principle in classification, as discussed in Crisp (2003), Barocas et al. (2017) and Chouldechova (2017). According to this principle, the probability of misclassifying an individual to class c should be equal across all protected groups:

$$\mathbb{P}(Y \neq \hat{Y} | \hat{Y} = c, A = a) \text{ are the same for all } a \in \mathcal{A}.$$
(1)

We highlight three primary issues related to machine learning methods developed under the sufficiency principle (Zeng et al. 2022, Pleiss et al. 2017, Zafar et al. 2017). First, the calibration by group method (Barocas et al. 2017), a popular approach for ensuring fair outcomes for subgroups, does not offer a theoretical guarantee on controlling the misclassification rate at a user-specified level. This lack of a guarantee can be particularly problematic in high-stakes decision-making situations. Second, current classification methods only focus on the accuracy of individual classifications, neglecting the complexities that arise when multiple individuals are classified simultaneously. This oversight regarding multiplicity can lead to severe inflation of misclassification errors. Finally, concurrent state-of-the-art classifiers typically exhibit high complexity and analytical intractability, making it difficult to quantify the uncertainties around their predictions. Even such theoretical analyses are feasible, they often involve strong assumptions about the underlying model and the accuracy of its outputs, which may not hold in practice. In response to these challenges, we propose a comprehensive approach comprising a selective classification framework (Section 2.2), a modified error rate criterion (Section 2.3), and a novel class of model-free algorithms with finite-sample theoretical guarantees (Sections 3.1-3.2). Together, these components provide a highly effective solution to the identified issues.

2.2 A selective inference framework for binary classification

This article mainly focuses on binary classification problems. The extension to the general multi-class setting is discussed briefly in Section 6.

Consider an application scenario that involves the prediction of mortgage default, where Y = 2 indicates default and Y = 1 otherwise. A common practice is to use risk assessment software to produce a *confidence score*, which is used to classify an individual into "high", "medium" or "low" risk classes. Let S(x, a) denote a score that maps an instance (x, a) to a real value, where a higher value indicates a higher risk of default. When the score represents the estimated class probability, we have $S(x, a) \in [0, 1]$. Suppose we observe a new instance

 $(X^*, A^*) = (x, a)$. Consider a class of decision rules of the form:

$$\hat{Y} = \mathbb{I}\{S(x, a) < t_l\} + 2\mathbb{I}\{S(x, a) > t_u\},\$$

where $\mathbb{I}(\cdot)$ is the indicator function, and $t_l < t_u$ are data-driven thresholds. The predicted label \hat{Y} takes three possible values in the action space $\Lambda = \{1, 2, 0\}$, respectively indicating that an individual has low $(\hat{Y} = 1)$, high $(\hat{Y} = 2)$ and medium $(\hat{Y} = 0)$ risks of default. The value "0", referred to as an "indecision" or "reject option" in classification (Herbei & Wegkamp 2006, Sun & Wei 2011, Lei 2014, Lee et al. 2021), is used to express "doubt," indicating insufficient confidence to make a definitive decision. For example, an individual with $\hat{Y} = 1$ will be approved for a mortgage, an individual with $\hat{Y} = 2$ will be rejected, while an individual with $\hat{Y} = 0$ will be asked to provide additional information and resubmit the application.

Now, we turn to a simultaneous inference task involving m individuals, with their confidence scores $\{S_j : j \in \mathcal{D}^{test}\}$. Consider the following decision rule,

$$\hat{Y}_j = \mathbb{I}(S_j < t_l) + 2\mathbb{I}(S_j > t_u), \quad \text{for } j \in \mathcal{D}^{test}.$$
(2)

We can view (2) as a selective inference procedure, which selects individuals with extreme scores into the high and low risk classes, while returning an indecision on the remainder. The selective inference view provides a flexible framework that allows for various types of classification rules. For example, if it is only of interest to select high-risk individuals, then the action space is $\Lambda = \{0, 2\}$, and one can use the following rule:

$$\hat{Y}_j = 2 \cdot \mathbb{I}(S_j > t_u), \text{ for } j \in \mathcal{D}^{test}.$$
 (3)

2.3 False selection rate and the fairness issue

In practice, it is desirable to avoid erroneous selections, which often have negative social or economic impacts. In the mortgage example, approving an individual who will truly default (i.e., $\hat{Y} = 1$ but Y = 2) would increase the financial burden of the lender, while rejecting an individual who will not default (i.e., $\hat{Y} = 2$ but Y = 1) would lead to a loss of profit. In situations where m is large, controlling the inflation of selection errors is a crucial task for policy makers. A practically useful notion is the false selection rate (FSR), which is defined as the expected fraction of erroneous decisions among all definitive decisions. We use the notation FSR^{C'}, where $\mathcal{C}' \subset \mathcal{C} = \{1, 2\}$ is the set of class labels that we are interested in selecting. To illustrate the definition, we consider two scenarios. In the first, we select individuals from both classes using rule (2). Denote $\mathcal{S} = \{j \in \mathcal{D}^{test} : \hat{Y}_j \neq 0\}$ the index set of the selected cases and $|\mathcal{S}|$ its cardinality. Then we have

$$\operatorname{FSR}^{\{1,2\}} = \mathbb{E}\left[\frac{\sum_{j \in \mathcal{D}^{test}} \sum_{c=1}^{2} \mathbb{I}(\hat{Y}_{j} = c, Y_{j} \neq c)}{|\mathcal{S}| \vee 1}\right],\tag{4}$$

where $x \lor y = \max\{x, y\}$, and the expectation \mathbb{E} is taken over both the observed data $\{(X_i, A_i, Y_i) : i \in \mathcal{D}\}$ and future test data $\{(X_j, A_j, Y_j) : j \in \mathcal{D}^{test}\}$.

In the second scenario, our goal is to select individuals in class c = 2 using rule (3). Denote $\mathcal{S} = \{j \in \mathcal{D}^{test} : \hat{Y}_j = 2\}$ the index set of the selected cases. Then

$$\operatorname{FSR}^{\{2\}} = \mathbb{E}\left[\frac{\sum_{j \in \mathcal{D}^{test}} \mathbb{I}(\hat{Y}_j = 2, Y_j \neq 2)}{|\mathcal{S}| \vee 1}\right].$$
(5)

 $FSR^{\{1\}}$ can be defined similarly. Incorporating indecisions into the decision-making process enables the derivation of a decision rule that is capable of controlling the FSRs at a user-specified level. However, achieving this goal is difficult under the standard classification setup that mandates definitive decisions for all individuals. For instance, if the minimum condition on the classification boundary (Meinshausen & Rice 2006, Cai & Sun 2017) is not met, simultaneously achieving small FSR^1 and FSR^2 will be impossible without an indecision option.

The FSR framework, which defines the error rate by focusing only on subjects that are selected or classified into class c, is inspired by the widely used and powerful idea of false discovery rate (FDR; Benjamini & Hochberg 1995) in multiple testing. The FSR is a general concept for selective inference that encompasses important special cases such as the misclassification rate, FDR and beyond. If we set both the state space and action space to be $\{1, 2\}$, so there are no indecisions, then the FSR defined by (4) reduces to the misclassification rate $m^{-1}\mathbb{E}\left\{\sum_{j\in\mathcal{D}^{test}} (\hat{Y}_j \neq Y_j)\right\}$. To see the connection between the FSR and FDR, consider a multiple testing problem with

$$H_{j0}: Y_j = 1$$
 vs. $H_{j1}: Y_j = 2$, for $j \in \mathcal{D}^{test}$.

The state space is $C = \{1, 2\}$. A multiple testing procedure $\hat{Y}_j \in \{0, 2\}, j \in \mathcal{D}^{test}$, corresponds to the selection rule (3). The action space $\Lambda = \{0, 2\}$ differs from the state space $C = \{1, 2\}$, with $\hat{Y}_j = 2$ indicating that H_{j0} is rejected, and $\hat{Y}_j = 0$ indicating that there is not enough evidence to reject H_{j0} . Then FSR^{2}, as defined in (5), precisely yields the widely used FDR, the expected fraction of false rejections among all rejections.

In practical scenarios, minimizing the number of indecisions is highly desirable. To quantify this concept, we introduce the expected proportion of indecisions:

$$\operatorname{EPI} = \frac{1}{m} \mathbb{E} \left\{ \sum_{j \in \mathcal{D}^{test}} \mathbb{I} \left(\hat{Y}_j = 0 \right) \right\} = 1 - \frac{\mathbb{E}(|\mathcal{S}|)}{m}.$$
 (6)

Under the same FSR level, a smaller EPI corresponds to greater statistical power.

Next we turn to the important fairness issue in selective inference. A major concern is that the rate of erroneous decisions might be unequally shared between the protected groups, as illustrated in the COMPAS example. To address this issue, it is desirable to control the FSR for each protected attribute in A. Therefore, we aim to find a selective classification rule obeying the following constraint on group-wise FSRs:

$$\operatorname{FSR}_{a}^{\{c\}} = \mathbb{E}\left[\frac{\sum_{j \in \mathcal{D}^{test}} \mathbb{I}(\hat{Y}_{j} = c, Y_{j} \neq c, A_{j} = a)}{\left\{\sum_{j \in \mathcal{D}^{test}} \mathbb{I}(\hat{Y}_{j} = c, A_{j} = a)\right\} \vee 1}\right] \leq \alpha_{c}, \quad \text{for all } a \in \mathcal{A},$$
(7)

where $\alpha_c, c \in \{1, 2\}$, is a user-specified tolerance level for class c. The fairness-adjusted error rate constraint (7) equally bounds the fraction of erroneous decisions among protected groups. We aim to develop a selective classification rule that solves the following constrained optimization problem: for $c \in \{1, 2\}$,

minimize the EPI, subject to
$$\text{FSR}_a^{\{c\}} \leq \alpha_c$$
, for all $a \in \mathcal{A}$. (8)

Although our problem formulation (8) only sets upper bounds for group-wise FSR levels, minimizing the EPI enforces the exhaustion of allowable FSR levels for each group. This leads to algorithms that closely align all group-wise FSR levels with the designated nominal level,



Figure 2: For FCC and RCC, the degree of unfairness increases as $\pi_{2|M}$ and $\pi_{2|F}$ become more disparate. FASI ensures that the group-wise FSRs are effectively controlled and approximately equalized.

ensuring comparability of error rates across all protected groups.

2.4 The construction of fair classifiers: issues and roadmap

We investigate the important issue of what makes a "fair" classifier. In most classification tasks, the standard operation is to first construct a confidence score, and then secondly to turn this score into a decision by setting a threshold. Let $\hat{S}_j^c \in [0, 1]$ denote a confidence score obtained from a machine learning model that predicts the conditional probability of individual $j \in \mathcal{D}^{test}$ being from class $c \in \mathcal{C}$, where for discussion below $\mathcal{C} = \{1, 2\}$. Consider the following thresholding rule:

$$\hat{Y}_j = \sum_{c=1}^2 c \cdot \mathbb{I}(\hat{S}_j^c > t_c), \quad \text{for } j \in \mathcal{D}^{test}.$$
(9)

We consider only the threshold $t_c \ge 0.5$ in (9) to prevent overlapping selections. A systematic two-stage strategy for dealing with overlapping selection is presented in Section A.4 of the Supplement.

We now present two approaches for constructing the score \hat{S}_j^c in the ideal setting where we assume that an oracle possesses knowledge of the true data generating model. The first method, referred to as the "full covariate classifier" (FCC), involves thresholding the following score:

$$S_j^{c,FCC}(x,a) = \mathbb{P}\left(Y_i = c | X_j = x, A_j = a\right), \quad \text{for } j \in \mathcal{D}^{test}.$$
 (10)

Here, $S_j^{c,FCC}(x, a)$ is used to denote the oracle probability of an individual being in class c based on all covariates. The second approach, referred to as the "reduced covariate classifier" (RCC), involves applying (9) by thresholding the following similar but less informative score:

$$S_j^{c,RCC}(x) = \mathbb{P}\left\{Y_j = c | X_j = x\right\}.$$
(11)

In this case, $S_j^{c,RCC}(x)$ is used to assess the same probability by eliminating the sensitive attribute from the covariate list. However, as we shall illustrate next, both the FCC and RCC approaches can be inadequate for addressing the issue of fairness in our context.

Consider the mortgage example where we simulate a data set that contains a sensitive attribute "gender". The goal is to select individuals into the high risk class with FSR control at 10%; the simulation setup is detailed in Section 4. We highlight here that the proportions of individuals with label "2" are different across the protected groups: for the male group, the proportion of individuals with label "2", denoted as $\pi_{2|M}$, is fixed at 50%, whereas for the female group the proportion $\pi_{2|F}$ varies from 15% to 85%.

We apply the FCC approach and plot the overall FSR and group-wise FSRs as functions of $\pi_{2|F}$ on the left panel of Figure 2. We can see that FCC controls the overall FSR but not the group-wise FSRs. Hence thresholding rules based on (10) are harmful in the sense that the burden of erroneous decisions is not shared equally among the two gender groups. The RCC approach can be harmful as well, as illustrated in the middle panel of Figure 2. While the overall FSR is still controlled at 10%, the issue of unfairness is in fact aggravated rather than mitigated, with widened gaps in the group-wise FSRs. In addition, the RCC approach has two further drawbacks. Firstly, disregarding a sensitive attribute can result in significant power loss. Secondly, if feature X is highly predictive of sensitive attribute A, then the RCC approach can still lead to unfair decisions due to the issue of surrogate encoding (Kusner et al. 2017, Long & Albert 2021). Concretely, in fairness research surrogate encoding pertains to the circumstance where the sensitive attribute A is absent from the list of predictors, but its information is encoded or concealed within other predictors, causing A to still influence the outcome Y.

It is essential to emphasize that the patterns presented in Figure 2 are not tied to a particular classification algorithm. Instead, they represent a systematic bias. In the above study, we consider the oracle setting, where the user has access to perfect scores (i.e., oracle posterior probabilities). However, even in this ideal scenario, the unfairness in Figure 2 still persists. In contrast, our proposed FASI algorithm, illustrated in the right panel of Figure 2, effectively controls the FSR and approximately equalizes the FSRs across all protected groups.

To achieve fairness across protected groups, there are two strategies to consider. Strategy (a) involves modifying the current confidence scores to create new scores that are comparable across groups. Conversely, Strategy (b) entails maintaining the original confidence scores and setting group-adjusted thresholds. Strategy (a), which employs a universal threshold for all individuals, is appealing because the decision-making process would be straightforward for practitioners, given that the new scores are comparable across groups. In contrast, although equally effective for addressing the fairness issue, Strategy (b) can be less intuitive and more difficult to implement. Practitioners without a full understanding of the underlying algorithm may find Strategy (b) confusing and even controversial, as varied thresholds are used for different protected groups, causing concern about possible discrimination at another level. Therefore, we will develop methodology by adopting Strategy (a) in the following sections.

3 Methodology

This section develops a fairness-adjusted selective inference (FASI) procedure for binary classification (with state space $\mathcal{C} = \{1, 2\}$). We focus on the goal of controlling the FSR^{c}_a defined in (5). The methodologies for controlling the FSR of the form (4) and the case of multinomial classification will be briefly discussed in Section 6 and Section A.4 of the Supplement.

A major challenge in our methodological development is that most state-of-the-art learning algorithms are complex and do not offer performance guarantees on their outputs. This limitation makes uncertainty quantification and error rate control challenging and even intractable. To overcome this challenge, we develop a model-free framework that is applicable to any black-box algorithm and relies only on the exchangeability between observed and test data.

3.1 The R-value and FASI algorithm

We first introduce a significance index, called the R-value, for ranking individuals and then discuss how the R-values can be converted to a selective classification rule via thresholding.

The R-value is computed via the FASI algorithm, which consists of three steps: training, calibrating and thresholding. The observed data set $\{(X_i, A_i, Y_i) : i \in \mathcal{D}\}$ is randomly divided into a training set and a calibration set: $\mathcal{D} = \mathcal{D}^{train} \cup \mathcal{D}^{cal}$. In the first step, we train a score function, denoted $\hat{S}^c(x, a)$, using data $\{(X_i, A_i, Y_i) : i \in \mathcal{D}^{train}\}$. By convention a larger score indicates a higher probability of belonging to class c. The scores, often representing estimated class probabilities, can be generated from any user-specified classifier satisfying $\hat{S}^1(x, a) + \hat{S}^2(x, a) = 1$. We make no assumptions on the accuracy of these scores.

In the second step, we first calculate scores $\hat{S}_i^c \coloneqq \hat{S}^c(X_i, A_i)$ for $i \in \mathcal{D}^{cal} \cup \mathcal{D}^{test}$, then define the following ratio

$$Q_k^c = \sum_{a \in \mathcal{A}} \mathbb{I}(A_k = a) \cdot \frac{\frac{1}{n_a^{cal} + 1} \left\{ \sum_{i \in \mathcal{D}^{cal}} \mathbb{I}\left(A_i = a, \hat{S}_i^c \ge \hat{S}_k^c, Y_i \ne c\right) + 1 \right\}}{\frac{1}{m_a + n_a^{cal} + 1} \left\{ \sum_{i \in \mathcal{D}^{cal} \cup \mathcal{D}^{test}} \mathbb{I}\left(A_i = a, \hat{S}_i^c \ge \hat{S}_k^c\right) + 1 \right\}} \land 1,$$
(12)

for $k \in \mathcal{D}^{cal} \cup \mathcal{D}^{test}$, where $m_a = \sum_{j \in \mathcal{D}^{test}} \mathbb{I}(A_j = a)$, $n_a^{cal} = \sum_{i \in \mathcal{D}^{cal}} \mathbb{I}(A_i = a)$. The operation \land indicates that Q_k^c is set to 1 if it exceeds 1. This makes sense because, as we will see in Section 3.2, Q_k^c can be interpreted as a proportion between 0 and 1. We included both the test and calibration data in the denominator to compute the ratio, which increases the stability of Q_k^c .

In practical scenarios, higher-scoring individuals may not consistently correspond to smaller Q_k^c values. To ensure an intuitive and sensible R-value definition, we propose the incorporation of the following monotonicity adjustment:

$$R_j^c \equiv \min_{\{k \in \mathcal{D}^{cal} \cup \mathcal{D}^{test}: \hat{S}_k^c \le \hat{S}_j^c\}} Q_k^c, \quad j \in \mathcal{D}^{test}.$$
(13)

Note that the computation of Q_k^c as an intermediate quantity is performed for $j \in \mathcal{D}^{cal} \cup \mathcal{D}^{test}$. However, the calculation of R_j^c , referred to as the R-value and utilized for decision making purposes, is only performed for $j \in \mathcal{D}^{test}$.

In the third step, we compare the R-value (13) with α_c . To provide a clear understanding, we consider the scenario where only individuals belonging to one specific class "c" are to be selected. Our decision rule is given by:

$$\hat{Y}_j = c \cdot \mathbb{I}(R_j^c \le \alpha_c), \quad j \in \mathcal{D}^{test}.$$
(14)

The FASI algorithm, summarized in Algorithm 1, possesses several attractive properties. First, as we will explain shortly, the R-value provides an estimate of a proportion, making it easily interpretable and comparable across groups. A smaller value of R_j^c indicates stronger evidence of belonging to class c. Second, the FSR analysis based on R-values is straightforward: practitioners can directly make decisions using the R-values by comparing them with the userspecified FSR level. Third, fairness consideration has been incorporated into the R-value. As revealed in Lemma 1 in Section D.1 of the Supplement, the decision rule (14) entails searching for the smallest group-wise threshold for \hat{S}^c subject to $\text{FSR}_a^{\{c\}} \leq \alpha_c$, thereby approximately aligning the group-wise FSR levels with the nominal level. Finally, FASI is model-free and provides a powerful framework for FSR control, as discussed in the next subsection.

3.2 Why FASI works?

We start by explaining why the R-value provides a sensible estimate of the FSR. To simplify the discussion, we ignore the sensitive attribute A for the moment and consider a thresholding rule

Algorithm 1 The FASI Algorithm

Input: $\{(X_i, A_i, Y_i) : i \in \mathcal{D}\}, \{(X_j, A_j) : j \in \mathcal{D}^{test}\}, \text{ a specific class } c, \text{FSR level } \alpha_c.$ **Output**: a selective classification rule $\{\hat{Y}_j \in \{0, c\} : j \in \mathcal{D}^{test}\}.$

- 1: Randomly split \mathcal{D} into \mathcal{D}^{train} and \mathcal{D}^{cal} .
- 2: Train a machine learning model on $\{(X_i, A_i, Y_i) : i \in \mathcal{D}^{train}\}$.
- 3: Predict confidence scores \hat{S}_i^c for $i \in \mathcal{D}^{cal} \cup \mathcal{D}^{test}$.
- 4: Compute the R-values $\{R_j^c: j \in \mathcal{D}^{test}\}$ according to Equations (12) and (13).
- 5: Return $\hat{Y}_j = c$ if $R_j^c \leq \alpha_c$ and $\hat{Y}_j = 0$ if $R_j^c > \alpha_c$, for $j \in \mathcal{D}^{test}$.

of the form $\{\mathbb{I}(\hat{S}_j^c \ge t) : j \in \mathcal{D}^{test}\}$. Consider the false selection proportion (FSP) process:

$$FSP(t) = \frac{\sum_{j \in \mathcal{D}^{test}} \mathbb{I}\left(\hat{S}_j^c \ge t, Y_j \neq c\right)}{\left\{\sum_{j \in \mathcal{D}^{test}} \mathbb{I}(\hat{S}_j^c \ge t)\right\} \lor 1},$$
(15)

with FSP(t) = 0 if no individual is selected. The FSP cannot be computed from data because we do not observe the true states $\{Y_j : j \in \mathcal{D}^{test}\}$. Then the effectiveness of the FASI algorithm relies on the following exchangeability condition:

Assumption 1. The triples $\{(X_i, A_i, Y_i) : i \in \mathcal{D}^{cal} \cup \mathcal{D}^{test}\}$ are exchangeable.

Specifically, under Assumption 1, the unobserved process $\sum_{j \in \mathcal{D}^{test}} \mathbb{I}\left(\hat{S}_j^c \geq t, Y_j \neq c\right)$ is strongly resembled by its "mirror process" in the calibration data $\sum_{i \in \mathcal{D}^{cal}} \mathbb{I}\left(\hat{S}_i^c \geq t, Y_i \neq c\right)$. Constructing a mirror process and exploiting the symmetry property to make inference is a powerful idea that has been explored in recent works (cf. Barber & Candès 2015, Weinstein et al. 2017, Lei & Fithian 2018, Leung & Sun 2022, Du et al. 2023). Meanwhile, we leverage both \mathcal{D}^{cal} and \mathcal{D}^{test} to enhance the stability of the denominator of (15), which does not require the knowledge of $Y_j, j \in \mathcal{D}^{cal} \cup \mathcal{D}^{test}$. To account for the unequal sample sizes between \mathcal{D}^{cal} and $\mathcal{D}^{cal} \cup \mathcal{D}^{test}$, we derive the mirror FSP process as follows:

$$Q^{c}(t) = \frac{\frac{1}{n^{cal}+1} \left\{ \sum_{i \in \mathcal{D}^{cal}} \mathbb{I}(\hat{S}_{i}^{c} \ge t, Y_{i} \neq c) + 1 \right\}}{\frac{1}{m+n^{cal}+1} \left\{ \sum_{j \in \mathcal{D}^{cal} \cup \mathcal{D}^{test}} \mathbb{I}(\hat{S}_{j}^{c} \ge t) + 1 \right\}},$$
(16)

where m and n^{cal} represent the cardinalities of \mathcal{D}^{test} and \mathcal{D}^{cal} , respectively. By restricting the $Q^c(t)$ process to a specific group $a \in \mathcal{A}$, substituting a candidate confidence score \hat{S}_k^c , for $k \in \mathcal{D}^{test}$, in place of t in (16), and implementing the monotonicity adjustment, we recover the fairness-adjusted R-value as presented in (13).

Remark 1. When inspecting (15) and (16), we observe that a "+1" has been added to the count of false selections on \mathcal{D}^{cal} in both the numerator and the denominator. This technical adjustment only has a minor impact on the empirical performance of FASI when the number of selections is small. However, it guarantees that (16) corresponds to a martingale, which is crucial for proving the theory. Therefore, it is reasonable to incorporate the same "+1" adjustment to n^{cal} in both the numerator and the denominator, slightly enhancing the power of the algorithm.

The FSP process (15) and its mirror process (16) together offer an intuitive interpretation of the R-value. Roughly speaking, the R-value represents the smallest estimated FSP at which

the i^{th} individual is just selected. Put differently, if we set the threshold at R = r, selecting all individuals with R-values less than or equal to r into class c, then we anticipate that, for every group $a \in \mathcal{A}$, roughly 100r% of the selections will be incorrect decisions. The fairness notion is inherently integrated into the R-value, enabling the calibration of a universal threshold to align all group-wise FSRs with the nominal level. Moreover, our interpretation resembles the q-value (Storey 2003) in FDR analysis; further details are provided in Section C in the Supplement. We highlight that while Storey's q-value relies on the empirical distribution of p-values, our R-value is derived from calibration data through a carefully designed mirror process.

We now present a theorem that establishes the finite-sample theory of FASI. Our theory significantly differs from existing works in that we do not make any assumptions about the accuracy of \hat{S}_i^c . Instead, the accuracy of the scores only impacts the power rather than the validity for FSR control. Practical guidelines on how to construct more informative confidence scores (and hence effective R-values) are discussed in Section B.4.

Theorem 1. Define $\gamma_{c,a} = \mathbb{E}\left(p_{c,null}^{test,a}/p_{c,null}^{cal,a}\right)$, where $p_{c,null}^{test,a}$ and $p_{c,null}^{cal,a}$ are the empirical proportions of individuals in group a that do not belong to class c in the test and calibration data, respectively. Then under Assumption 1, we have, for all $a \in \mathcal{A}$, the FASI algorithm with R-value (13) satisfies $FSR_a^{\{c\},*} \leq \gamma_{c,a}\alpha_c$, where

$$FSR_a^{\{c\},*} = \mathbb{E}\left[\frac{\sum_{j\in\mathcal{D}^{test}}\mathbb{I}(\hat{Y}_j=c,Y_j\neq c,A_j=a)}{\sum_{j\in\mathcal{D}^{test}}\mathbb{I}(\hat{Y}_j=c,A_j=a)+1}\right].$$
(17)

Assumption (1) on exchangeability implies that $\gamma_{c,a}$ is typically very close to 1, resulting in nearly exact control; this has been corroborated in our numerical studies (Section G.2 of the Supplement). Moreover, in Section A.3, we present a corollary demonstrating that a conservative version of the R-value guarantees FSR control below α , eliminating $\gamma_{c,a}$ from the bound. However, the conservative version tends to result in a higher proportion of indecisions (or decreased power). As a result, we recommend the R-value presented in (12) and (13), which are more powerful while providing almost exact control in practical situations.

Remark 2. In the modified FSR definition (17), the "+1" adjustment is utilized to account for the extra uncertainty in the approximation of the number of rejections. A similar modification was employed in Theorem 1 of Barber & Candès (2015), but for different reasons. In Section A.2 of the Supplement, we demonstrate that a simpler version of the R-value, which only incorporates \mathcal{D}^{test} in the denominator, can effectively control the unadjusted $\text{FSR}_a^{\{c\}}$ (7). We recommend the use of (12) in practice for two reasons: (a) the difference between $\text{FSR}_a^{\{c\},*}$ and $\text{FSR}_a^{\{c\}}$ is usually negligible, and (b) incorporating both \mathcal{D}^{cal} and \mathcal{D}^{test} in the denominator increases the stability of the R-value (cf. Section G.1 of the Supplement).

Finally we explain the challenges and ideas for proving Theorem 1. The discussion focuses on the R-value process (16), but can be easily extended to the group-adjusted case (12). Three major challenges in the theoretical analysis include (a) how to handle the dependence between the scores \hat{S}_i^c [as the same training data have been used to compute (16)], (b) how to evaluate the FSR without knowledge about quality of the scores, and (c) how to develop non-asymptotic guarantees on the performance of FASI. Inspired by the elegant ideas in the FDR literature (Storey et al. 2004, Barber & Candès 2015), we have carefully designed the R-values so that the corresponding FSP process (16) is stochastically dominated by a super-martingale. We then apply the optional stopping theorem and leverage the exchangeability assumption to establish the upper bound for the FSR. We stress that our theory utilizes no assumptions on the underlying models or quality of scores, and the algorithm is provably valid in finite samples.

3.3 Connections to existing work

This section explores the connections and distinctions between FASI and existing methods developed under the sufficiency principle in fairness research. Additionally, we provide insights about recent developments in conformal inference that bear relevance to FASI.

Our formulation (7) is closely related to the sufficiency principle (1) in fairness literature, but it overcomes several of its limitations. Firstly, (7) operates within a selective classification framework, offering an indecision option for individuals needing further review, thereby enabling effective error rate control at user-specified levels. Secondly, we define the FSR notion to aggregate decision errors on m new individuals, overcoming the sufficiency principle's limitation, which only pertains to the error rate of an individual decision. Lastly, many algorithms developed under the sufficiency principle are complex and computationally intensive, lacking finite sample guarantees over outputs from blackbox models. In contrast, the FASI algorithm effectively controls the FSR in finite samples without assumptions about the underlying model, classification algorithm, or score accuracy. A detailed comparison with related works, including Zeng et al. (2022) and Lee et al. (2021), is provided in Section F of the Supplement.

The R-value has a compelling interpretation under the *conformal inference* (Vovk et al. 2005, Lei & Wasserman 2014) framework. In Section C of the Supplement, we show that a variation of our R-value corresponds to the Benjamini-Hochberg (BH) adjusted q-value (Storey 2003) of the *conformal p-values* (Bates et al. 2023, Mary & Roquain 2022) under the *one-class classification* setting (Moya & Hush 1996, Khan & Madden 2009, Kemmler et al. 2013). The connection to conformal inference and the BH method provides valuable insights into why the FASI algorithm is model-free and offers effective FSR control in finite samples, as claimed in Theorem 1.

The theory presented in Bates et al. (2023) encounters a complication similar to ours as the conformal p-values are dependent. To address this, Bates et al. (2023) first show that the conformal p-values satisfy the PRDS condition and then apply the theory in Benjamini & Yekutieli (2001) to establish the validity of FDR control. While we conjecture that the PRDS approach may be relevant, its extension to our specific context is non-trivial as our R-values do not explicitly utilize conformal p-values under the binary classification setup. Therefore, our theory via martingales appears to be a suitable and equally effective alternative. Moreover, incorporating conformal p-values, which rely on one-class classifiers, directly into our binary classification problem would entail discarding labeled outliers and consequently lead to information loss; this issue has been explored in a recent study by Liang et al. (2024).

3.4 Theoretical R-value and optimality theory

We briefly mention the theoretical R-value and the optimality theory, which follows from the works of Sun & Cai (2007) and Cai et al. (2019). Details are deferred to Section B of the Supplement due to page constraints. Although developed under a highly idealized setup, the theory provides insights for practitioners regarding how to train score functions to construct informative R-values. We highlight two essential messages.

Firstly, our analysis shows that the choice of optimal score function indicates that, during the training stage, we should emulate the full covariate classifier (10), which utilizes *all features*, including the sensitive attribute A, to best capture individual level information. Scores trained without the sensitive attribute [e.g. (11)] are usually suboptimal. The fairness adjustments for decisions should not be made during the training stage but in the calibration stage, where the fully informative scores can be converted to R-values to adjust the disparity in error rates across groups. This strategy shares the same spirit with the *learn-then-test* or *selection-by-prediction* framework advocated by Angelopoulos et al. (2022) and Jin & Candes (2023).

Secondly, we find that the optimal selection rule equalizes the group-wise error rates. The intuition is that in order to minimize the EPI, the pre-specified mFSRs must be *exhausted in all separate groups*; hence the group-wise mFSRs are all equal to the nominal level (thereby automatically equalized). This implies that handling the fairness issue through a constrained optimization problem (7) leads to the asymptotic equality of our FASI algorithm. Our numerical studies corroborate this claim, though a full analysis is complicated due to the dependence between the scores. We leave this for future research.

4 Simulations

This section presents two simulations under the binary classification setup. The objective is to compare the performance of FASI against the Full Covariate Classifier (FCC). We did not include the Restricted Covariate Classifier (RCC) in these simulations, as RCC has consistently demonstrated larger deviations from the target group-wise FSR levels. We demonstrate that both the oracle and data-driven versions of FASI can control the group-wise FSRs, while RCC fails to do so. The oracle versions of FASI and FCC use the exact posterior probabilities for S_j^c , defined in Equation 10, while the data-driven procedures estimate S_j^c via the GAM method (Hastie et al. 2009, James et al. 2023, Chen & Guestrin 2016).

In all simulations, we set $|\mathcal{D}^{train}| = 1,500$, $|\mathcal{D}^{cal}| = 1,000$ and $|\mathcal{D}^{test}| = 1,000$. Gender is our protected attribute taking two values A = F (females) and A = M (males). The feature vectors $X \in \mathbb{R}^3$ are simulated according to the following model:

$$F(\cdot) = \pi_M \{ \pi_{1|M} F_{1,M}(\cdot) + \pi_{2|M} F_{2,M}(\cdot) \} + \pi_F \{ \pi_{1|F} F_{1,F}(\cdot) + \pi_{2|F} F_{2,F}(\cdot) \},\$$

where $\pi_a = \mathbb{P}(A = a)$, $\pi_{c|a} = \mathbb{P}(Y = c|A = a)$ and $F_{c,a}$ is the conditional distribution of X given Y = c and A = a. Let $\pi_M = \pi_F = 0.5$. We will consider two scenarios in our simulation. Although only the GAM method is employed in our simulation, we report that our findings remain consistent regardless of the specific learning algorithms utilized. For a comparison of different machine learning algorithms, please refer to Section G.3 of the Supplement.

In the first scenario, the conditional distributions of X given class Y are assumed to be multivariate normal and are identical for males and females:

$$F_{1,M} = F_{1,F} = \mathcal{N}(\boldsymbol{\mu}_1, 2 \cdot \mathbf{I}_3), \quad F_{2,M} = F_{2,F} = \mathcal{N}(\boldsymbol{\mu}_2, 2 \cdot \mathbf{I}_3),$$

where \mathbf{I}_3 is a 3×3 identity matrix, $\boldsymbol{\mu}_1 = (0, 1, 6)^{\top}$ and $\boldsymbol{\mu}_2 = (2, 3, 7)^{\top}$. The only difference between the group-wise distributions lies in the conditional proportions: we fix $\pi_{2|M} = \mathbb{P}(Y = 2|A = M) = 0.5$, while varying $\pi_{2|F} = \mathbb{P}(Y = 2|A = F)$ from 0.15 to 0.85. We shall see that in the asymmetric situation (i.e., when $\pi_{2|F}$ is very large or small), the unadjusted FCC rule leads to unfair policies (i.e. we observe disparate FSRs across the male and female groups).

We simulate 1,000 data sets and apply both the FCC and FASI methods at FSR level 0.1 to the simulated data sets. The FASI method is implemented with R-values defined in (13). For the FCC method, the protected attributes are ignored when computing the fractions in (12):

$$Q_k^{c,\text{FCC}} = \frac{\frac{1}{n_a^{cal}+1} \left\{ \sum_{i \in \mathcal{D}^{cal}} \mathbb{I} \left(\hat{S}_i^c \ge \hat{S}_j^c, Y_i \ne c \right) + 1 \right\}}{\frac{1}{m_a + n_a^{cal}+1} \left\{ \sum_{i \in \mathcal{D}^{cal} \cup \mathcal{D}^{test}} \mathbb{I} \left(\hat{S}_i^c \ge \hat{S}_j^c \right) + 1 \right\}};$$

then these $Q_k^{c,\text{FCC}}$ values are adjusted according to (13). The corresponding selection rule is $\left\{ \hat{Y}_j = c \cdot \mathbb{I}(R_j^{c,\text{FCC}} \leq 0.1) : j \in \mathcal{D}^{test} \right\}$.



Figure 3: Simulation 1. Top row: the oracle procedure. Bottom row: the data-driven procedure using GAM. Left and middle columns: FSR_a^1 and FSR_a^2 levels for both females and males. Right column: the EPI levels.

The FSR levels are computed by averaging the respective false discovery proportions (FSPs) from 1,000 replications. The simulation results are summarized in Figure 3. The first and second rows respectively correspond to the oracle and data-driven versions of each method. The first two columns respectively plot the group-wise FSRs for class 1 and class 2 as functions of $\pi_{2|F}$. The final column plots the EPI (6), obtained by averaging the results from 1,000 replications. The following patterns can be observed.

- FCC fails to control the group-wise FSRs. As $\pi_{2|F}$ moves away from $\pi_{2|M} = 0.5$, the gap between the FSR control for Females and Males dramatically widens due to the asymmetry in the proportions of the signals (true class 2 observations) in the male and female groups.
- Both the oracle and data-driven FASI procedures consistently control the FDR at the nominal level. However, when $\pi_{2|F}$ is high, the number of selections decreases, resulting in a reduced total number of selections from both groups. Consequently, both methods exhibit increased conservativeness. This pattern can be attributed to the conservative nature of the R-value (12), which includes a "+1" adjustment and functions as an estimate of the true false selection proportion: the level of conservativeness becomes more pronounced as the proportion $\pi_{2|F}$ become close to either 0 or 1.
- Both oracle and data-driven FASI algorithms are able to roughly equalize the group-wise FSRs between the Female and Male groups. The data-driven FASI is able to closely mirror the behavior of the oracle method.
- The parity in FSR control is achieved at the price of slightly higher EPI levels.

Our second simulation considers the setting where $F_{c,M} \neq F_{c,F}$. Denoting the mean for class c and protected attribute a as $\boldsymbol{\mu}_{c,a}$, the data is generated from $F_{c,a} = \mathcal{N}(\boldsymbol{\mu}_{c,a}, 2 \cdot \mathbf{I}_3)$, with components $\boldsymbol{\mu}_{1,M} = (0, 1, 6)^{\top}$, $\boldsymbol{\mu}_{2,M} = (2, 3, 7)^{\top}$, $\boldsymbol{\mu}_{1,F} = (1, 2, 7)^{\top}$ and $\boldsymbol{\mu}_{2,F} = (3, 4, 8)^{\top}$. In all other respects Simulations 1 and 2 are identical. The results for the second simulation scenario are provided in Figure 4. We notice very similar patterns to our first simulation setup. FASI



Figure 4: Simulation 2. Comparable setup to Simulation 1 except that the female and male distributions now differ from each other.

controls the group-wise FSRs for all values of $\pi_{2|F}$ while the FCC fails to do so. The data-driven FASI closely emulates the oracle procedure, for both the FSR and EPI levels.

Finally, we examine the variability of the false discovery proportions (FSP), which can fluctuate across replications. Specifically, the FSR is derived as the average of the FSPs. While our theory ensures that the FSR can be controlled under the nominal level α , it is important to note that the FSP may deviate significantly from α . To investigate this variability, we focus on the same experimental setting in Simulation 1 used to generate Figure 3, and present the 90% quantiles of the group-wise FSPs. The summarized results are depicted in Figure 5.

The group-wise FSRs, represented by solid blue and dot-dashed red lines, are effectively controlled at the desired 10% level. The quantiles are visually depicted by the blue or red shaded regions, corresponding to the male and female groups, respectively. For the male group, where $\pi_{2|M}$ remains constant, the 90% quantiles range between 5% and 15%. In contrast, the FSP variability for the female group is more pronounced, with greater variability observed when $\pi_{2|F}$ assumes larger values. This larger FSP variability is closely associated with a smaller number of selections made from the female group.

5 Real Data Examples

This section demonstrates the application of FASI on two real data sets. Sections 5.1 and 5.2 respectively analyze the COMPAS data (Angwin et al. 2016) and US census data (Dua & Graff 2017). For the COMPAS and census data, we have employed GAM and Adaboost models, respectively, to construct confidence scores. It is important to note that users have the flexibility to choose the best model for their specific application by utilizing their own training data. To facilitate the implementation of FASI with user-specified models, the R package fasi has been developed and is readily available on CRAN.



Figure 5: 90% quantiles of the FSPs. Red region: female group; Blue region: male group.

5.1 COMPAS data analysis

In 2016, ProPublica's investigative journalists curated a data set of 6,172 individuals, where 3,175 were Black and the remaining 2,997 belonged to other racial categories, who had been arrested in Broward County, Florida. These racial categories, Black and Other, serve as our protected attributes in this study. Within the data set, the "Black" group consisted of 1,773 individuals who were identified as having recidivated within the 2-year time frame considered in the study, while the "Other" group consisted of 1,217 individuals who also recidivated during this period. This 2-year window was chosen as a proxy for the true label of identifying recidivists.

All individuals were assigned a risk score by the COMPAS algorithm (a whole number between 1 and 10) developed by NorthPointe Inc. This score was used to inform the judge of each person's risk of recidivating during their bail hearing. The data set contains demographic information about each person including their race, age, number of previous offenses, sex, number of prior offenses, and their assigned COMPAS risk score.

In this analysis, our objective is to utilize FASI to address potential disparities in FSRs among different racial groups. The literature has extensively examined various fairness notions, such as disparate treatment (Zafar et al. 2017), as well as studies specifically related to the COMPAS data set (Angwin et al. 2016, Dieterich et al. 2016). While our approach is highly relevant and sensible in this context where high-stake consequences are involved, it is crucial to carefully evaluate and scrutinize its implementation, bearing in mind the societal trade-offs associated with different definitions of fairness.

We performed 100 random splits of the data set, where for each protected group and class label Y (our proxy for recidivism), 90% of the data was assigned to \mathcal{D} and the remaining 10% to \mathcal{D}^{test} . Furthermore, we evenly split \mathcal{D} into \mathcal{D}^{train} and \mathcal{D}^{cal} . To assess the performance, we present the results across a range of α values from 0.15 to 0.30. The first two columns in Figure 6 illustrate the difference between the true and target FSRs for the FCC and FASI algorithms, respectively. The last column of the figure plots the EPI levels.

While the FCC approach effectively controls the overall FSR, it falls short in controlling the FSRs across different racial groups. In the left panel of Figure 6, we can observe that the race-wise FSRs deviate substantially from the nominal level, and the FSR levels for the Black



Figure 6: COMPAS data analysis for predicting recidivists. Left and Middle: False Selection Rate minus the desired control level for varying levels of α for the FCC and FASI method respectively. Right: The EPI for both the FCC and FASI method.

group are significantly lower compared to those of the Other group. This discrepancy persists consistently across all values of α . In contrast, the middle panel of Figure 6 demonstrates that by employing the FASI algorithm, the race-wise FSR levels are effectively controlled below the nominal level and are approximately equalized across the sensitive groups. Moreover, the right panel illustrates that FASI achieves a nearly identical EPI level as the FCC approach.

5.2 1994 census income data analysis

The US census is a primary source of information for generating data concerning the American population. Consequently, the data they collect plays a direct role in informing future policy decisions, such as allocating resources for programs that offer economic assistance to vulnerable populations. These resources encompass necessities such as food, healthcare, job training, housing, and other forms of economic aid, which rely on accurate estimates of income levels within the population. The potential consequences of making unfair decisions when predicting income levels can be significant, as these predictions contribute to determining how hundreds of billions of dollars in federal funding will be allocated over the next decade. In this case study, we utilize the 1994 US Census Data set from the UCI Machine Learning Repository to predict whether an individual earns above or below \$50,000 per year, with Class 1 representing individuals earning less than \$50,000 and Class 2 representing those earning more than \$50,000. To avoid overlapping selections, we utilize the two-stage procedure described in Section A.4 of the Supplement.

The data set in this study comprises 32,561 observations on 14 variables, predominantly demographic factors such as education level, age, and hours worked per week, among others. The protected attributes under consideration are Female and Male. Specifically, the Female group consists of a total of 10,771 observations, with 1,179 individuals earning over \$50,000 per year. Similarly, the Male attribute encompasses the remaining 21,790 observations, with 6,662 individuals earning over \$50,000 per year.

We applied the FCC and FASI algorithms at different FSR levels, ranging from 0.05% to 10%. We performed 100 random splits of the dataset, where for each gender and class label, 70% of the data was randomly assigned to \mathcal{D} , and the remaining 30% was assigned to \mathcal{D}^{test} . Furthermore, \mathcal{D} was evenly divided into \mathcal{D}^{train} and \mathcal{D}^{cal} . The left and middle panels of Figure 7 respectively show the FSR levels for both the FCC and FASI.

From the left column, we can observe that the group-wise FSR levels of FCC consistently deviate from the nominal level α , resulting in unfair decisions for the sensitive groups. This



Figure 7: Census income prediction. Class 1 (top row) comprises individuals earning *less than* \$50,000 per year, while Class 2 (bottom row) includes individuals earning *more than* \$50,000 per year. Left and Middle: False Selection Rate minus varying levels of α . Right: The EPI levels.

pattern is observed in both Class 1 and Class 2, although in opposite directions. The disparity in group-wise FSR levels becomes more pronounced as α increases. In contrast, the middle column demonstrates that for Class 1, the group-wise FSR levels of FASI remain close to α . For Class 2, the group-wise FSR levels of FASI exhibit conservativeness but are roughly equalized across the two sensitive groups. The conservativeness can be attributed to the R-value, which provides a conservative estimate of the true FSP. Furthermore, the right column highlights that FASI effectively achieves approximate parity, ensuring that the burden is roughly equally shared across the two genders, with only a slight increase in the EPI level.

6 Discussions

This section concludes the article by discussing additional fairness notions and extensions of the FSR concept, highlighting limitations in existing research and suggesting possible future directions for further exploration.

6.1 Related fairness notions

Fairness in machine learning presents a complex challenge. Multiple studies focus on addressing representation or sampling bias, which arises when data are collected in a non-representative fashion (Mehrabi et al. 2021). By contrast, algorithmic bias emerges when the model itself introduces bias beyond the inherent biases in the input data. This article addresses the issue of algorithmic bias, with the objective of ensuring an equitable distribution of erroneous decisions across different groups. FASI is model-free, allowing for deployment with any user-specified model. It achieves fairness by aligning group-wise FSRs to the same designated level, requiring only mild conditions on data exchangeability.

In addition to the sufficiency principle, the *separation principle* (Barocas et al. 2017), a second notion of fairness, has been widely used. It requires that

$$P(Y \neq \hat{Y}|Y = c, A = a)$$
 are the same for all $a \in \mathcal{A}$. (18)

This principle differs from the sufficiency principle (1), whereby \hat{Y} and Y interchange positions in the conditional probability expression. A third notion on fairness, in the context of prediction intervals, has been considered in Romano, Barber, Sabatti & Candès (2020). Rather than conditioning on either Y or \hat{Y} , this fairness criterion is concerned with the joint probabilities of (\hat{Y}, Y) , requiring that the misclassification rates are equalized across all protected groups:

$$P(Y \neq \hat{Y}|A = a)$$
 are the same for all $a \in \mathcal{A}$. (19)

The fourth notion, known as *demographic parity* (Jiang et al. 2020) requires that

$$P(\hat{Y} \neq c|A = a)$$
 are the same for all $a \in \mathcal{A}$. (20)

Other popular fairness notions include *equalized odds* (Hardt et al. 2016, Romano, Bates & Candès 2020) and *equalized risks* (Corbett-Davies & Goel 2018).

Zafar et al. (2017) proposed the use of cost-sensitive classifiers with group-specific costs (Menon & Williamson 2018) to address a fairness issue comparable to our work. However, their technique forces a decision to be made on all individuals, whereas our approach is a selective inference procedure that only makes confident judgments on a subset of subjects. Given human intervention, FASI can achieve higher accuracy than cost-sensitive classifiers, as practitioners are aware of the undecided cases that merit additional scrutiny, ultimately reducing erroneous decisions with potentially extensive societal costs.

Our fairness criterion, as described in Equation 7, constitutes a group fairness notion that presupposes full knowledge of the protected groups. This approach is widely adopted in the literature and finds applications across diverse domains, including medicine and the criminal justice system (Manrai et al. 2016, Angwin et al. 2016), often facilitated by specialized software tools (Bellamy et al. 2018, Saleiro et al. 2018). However, situations may arise where the protected groups lack clear delineation, such as when the sensitive attribute pertains to age or income. New ideas, such as individual fairness and counterfactual fairness, provide useful alternatives. Specifically, individual fairness aims to ensure that comparable individuals receive commensurate outcomes (Mukherjee et al. 2020), while counterfactual fairness posits that fairness should not be exclusively contingent on observable attributes but should also consider potential counterfactual factors. Given the substantial complexities associated with individual and counterfactual fairness algorithms, we leave exploration of this promising avenue in future research.

A highly contentious matter is that disparate fairness criteria often yield distinct algorithms and different decisions in practice. For instance, the sufficiency and separation principles can be incompatible with one another (Kleinberg et al. 2016, Friedler et al. 2021), and classification parity or group calibration can potentially harm the very groups that these algorithms are intended to protect (Corbett-Davies & Goel 2018). Despite growing awareness of fairness concerns in decision-making, a consensus is yet to be reached on the best approaches for achieving fairness in machine learning. While we do not claim that FASI is ubiquitously superior to competing approaches, adjusting group-wise FSRs appears to be an effective and suitable fairness criterion for high-stake applications, overcoming several limitations of the widely used sufficiency principle. Much research is still needed for understanding the trade-offs and applicability of different fairness notions across diverse contexts and applications.

6.2 FSR concepts in multinomial classification

The selective inference framework and FSR concepts can be extended beyond the binary classification setting. Denote the collection of all class labels by $C = \{1, \ldots, C\}$. The case with C = 1corresponds to the one-class classification problem, as recently explored in Guan & Tibshirani (2021) and Bates et al. (2023), which can be encompassed within our general framework.

For situations with $C \geq 2$, denote the set of classes to be selected by \mathcal{C}' , and assume $\mathcal{C}' \subset \mathcal{C}$. With indecisions being allowed, the action space is given by $\Lambda = \{0, \mathcal{C}'\}$. Denote the selection rule $\{\hat{Y}_i : i \in \mathcal{D}^{test}\} \in \Lambda^m$. Then the FSR with respect to subset \mathcal{C}' is defined as the expected fraction of erroneous selections among all selections:

$$FSR^{\mathcal{C}'} = \mathbb{E}\left[\frac{\sum_{j \in \mathcal{D}^{test}} \mathbb{I}(\hat{Y}_j \in \mathcal{C}', \hat{Y}_j \neq Y_j)}{\left\{\sum_{i \in \mathcal{D}^{test}} \mathbb{I}(\hat{Y}_j \in \mathcal{C}')\right\} \lor 1}\right].$$

The group-wise FSRs taking into account the protected attribute A can be defined analogously to (7) by restricting the selections to specific groups. The EPI definition remains the same. The extension of the R-value in this setup is briefly discussed in Section A.4 of the Supplement, with some remaining issues left for future research.

Acknowledgement

We are grateful to the referees whose constructive feedback has much improved the clarity and presentation of our manuscript. The valuable suggestions from Matteo Sesia and Zinan Zhao have been instrumental in advancing the methodology and theory of our work.

References

- Angelopoulos, A. N., Bates, S., Candès, E. J., Jordan, M. I. & Lei, L. (2022), 'Learn then test: Calibrating predictive algorithms to achieve risk control', arXiv preprint arXiv:2110.01052.
- Angwin, J., Larson, J., Mattu, S. & Kirchner, L. (2016), 'Machine bias: There's software used across the country to predict future criminals', And it's biased against blacks. ProPublica 23, 77–91.
- Barber, R. F. & Candès, E. J. (2015), 'Controlling the false discovery rate via knockoffs', The Annals of Statistics 43(5), 2055–2085.
- Barocas, S., Hardt, M. & Narayanan, A. (2017), 'Fairness in machine learning', Nips tutorial 1, 2.
- Bates, S., Candès, E., Lei, L., Romano, Y. & Sesia, M. (2023), 'Testing for outliers with conformal p-values', The Annals of Statistics 51(1), 149 – 178. URL: https://doi.org/10.1214/22-AOS2244
- Bellamy, R. K. E., Dey, K., Hind, M., Hoffman, S. C., Houde, S., Kannan, K., Lohia, P., Martino, J., Mehta, S., Mojsilovic, A., Nagar, S., Ramamurthy, K. N., Richards, J., Saha, D., Sattigeri, P., Singh, M., Varshney, K. R. & Zhang, Y. (2018), 'Ai fairness 360: An extensible toolkit for detecting, understanding, and mitigating unwanted algorithmic bias'. URL: https://arxiv.org/abs/1810.01943
- Benjamini, Y. (2010), 'Simultaneous and selective inference: Current successes and future challenges', Biometrical Journal 52(6), 708–721.
- Benjamini, Y. & Hochberg, Y. (1995), 'Controlling the false discovery rate: a practical and powerful approach to multiple testing', J. Roy. Statist. Soc. B 57, 289–300.
- Benjamini, Y. & Hochberg, Y. (2000), 'On the adaptive control of the false discovery rate in multiple testing with independent statistics', Journal of Educational and Behavioral Statistics 25, 60–83.

- Benjamini, Y. & Yekutieli, D. (2001), 'The control of the false discovery rate in multiple testing under dependency', Ann. Statist. 29(4), 1165–1188.
- Cai, T. & Sun, W. (2017), 'Optimal screening and discovery of sparse signals with applications to multistage high-throughput studies', Journal of the Royal Statistical Society: Series B (Statistical Methodology) 79(1), 197.
- Cai, T. T. & Sun, W. (2009), 'Simultaneous testing of grouped hypotheses: Finding needles in multiple haystacks', J. Amer. Statist. Assoc. 104, 1467–1481.
- Cai, T. T., Sun, W. & Wang, W. (2019), 'CARS: Covariate assisted ranking and screening for large-scale two-sample inference (with discussion)', J. Roy. Statist. Soc. B 81, 187–234.
- Chen, T. & Guestrin, C. (2016), XGBoost: A scalable tree boosting system, in 'Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining', KDD '16, ACM, New York, NY, USA, pp. 785–794. URL: http://doi.acm.org/10.1145/2939672.2939785
- Chouldechova, A. (2017), 'Fair prediction with disparate impact: A study of bias in recidivism prediction instruments', *Big data* 5(2), 153–163.
- Corbett-Davies, S. & Goel, S. (2018), 'The measure and mismeasure of fairness: A critical review of fair machine learning', arXiv preprint arXiv:1808.00023.
- Crisp, R. (2003), 'Equality, priority, and compassion', *Ethics* **113**(4), 745–763.
- Dieterich, W., Mendoza, C. & Brennan, T. (2016), 'Compas risk scales: Demonstrating accuracy equity and predictive parity', *Northpointe Inc*.
- Du, L., Guo, X., Sun, W. & Zou, C. (2023), 'False discovery rate control under general dependence by symmetrized data aggregation', *Journal of the American Statistical Association* **118**(541), 607–621.
- Dua, D. & Graff, C. (2017), 'UCI machine learning repository'. URL: http://archive.ics.uci.edu/ml
- Friedler, S. A., Scheidegger, C. & Venkatasubramanian, S. (2021), 'The (im) possibility of fairness: different value systems require different mechanisms for fair decision making', *Communications of* the ACM 64(4), 136–143.
- Guan, L. & Tibshirani, R. (2021), 'Prediction and outlier detection in classification problems', Journal of the Royal Statistical Society: Series B (Statistical Methodology) to appear, arXiv:1905.04396.
- Hardt, M., Price, E. & Srebro, N. (2016), 'Equality of opportunity in supervised learning', Advances in neural information processing systems 29, 3315–3323.
- Hastie, T., Tibshirani, R. & Friedman, J. (2009), The Elements of Statistical Learning: Data Mining, Inference, and Prediction, Springer series in statistics, Springer.
- Herbei, R. & Wegkamp, M. H. (2006), 'Classification with reject option', The Canadian Journal of Statistics/La Revue Canadienne de Statistique pp. 709–721.
- James, G., Witten, D., Hastie, T. & Tibshirani, R. (2023), An Introduction to Statistical Learning: with Applications in R, Vol. 2, Springer.
- Jiang, R., Pacchiano, A., Stepleton, T., Jiang, H. & Chiappa, S. (2020), Wasserstein fair classification, in R. P. Adams & V. Gogate, eds, 'Proceedings of The 35th Uncertainty in Artificial Intelligence Conference', Vol. 115 of *Proceedings of Machine Learning Research*, PMLR, pp. 862–872.
- Jin, Y. & Candes, E. J. (2023), 'Selection by prediction with conformal p-values', Journal of Machine Learning Research 24(244), 1–41.

- Kemmler, M., Rodner, E., Wacker, E.-S. & Denzler, J. (2013), 'One-class classification with gaussian processes', *Pattern recognition* 46(12), 3507–3518.
- Khan, S. S. & Madden, M. G. (2009), A survey of recent trends in one class classification, in 'Irish conference on artificial intelligence and cognitive science', Springer, pp. 188–197.
- Kleinberg, J., Mullainathan, S. & Raghavan, M. (2016), 'Inherent trade-offs in the fair determination of risk scores', arXiv preprint arXiv:1609.05807.
- Kusner, M. J., Loftus, J., Russell, C. & Silva, R. (2017), Counterfactual fairness, in I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan & R. Garnett, eds, 'Advances in Neural Information Processing Systems', Vol. 30, Curran Associates, Inc.
- Lee, J. K., Bu, Y., Rajan, D., Sattigeri, P., Panda, R., Das, S. & Wornell, G. W. (2021), Fair selective classification via sufficiency, in 'Proceedings of the 38th International Conference on Machine Learning', Vol. 139 of Proceedings of Machine Learning Research, PMLR, pp. 6076–6086.
- Lei, J. (2014), 'Classification with confidence', *Biometrika* **101**(4), 755–769.
- Lei, J. & Wasserman, L. (2014), 'Distribution-free Prediction Bands for Non-parametric Regression', Journal of the Royal Statistical Society Series B: Statistical Methodology 76(1), 71–96.
- Lei, L. & Fithian, W. (2018), 'AdaPT: An Interactive Procedure for Multiple Testing with Side Information', Journal of the Royal Statistical Society Series B: Statistical Methodology 80(4), 649–679.
- Leung, D. & Sun, W. (2022), 'ZAP: Z-Value Adaptive Procedures for False Discovery Rate Control with Side Information', Journal of the Royal Statistical Society Series B: Statistical Methodology 84(5), 1886–1946.
- Liang, Z., Sesia, M. & Sun, W. (2024), 'Integrative conformal p-values for powerful out-of-distribution testing with labeled outliers', *Journal of the Royal Statistical Society Series B: Statistical Methodology* p. to appear. arXiv preprint arXiv:2208.11111.
- Long, K. D. & Albert, S. M. (2021), Use of zip code based aggregate indicators to assess race disparities in covid-19, Vol. 31, Ethnicity & disease.
- Manrai, A. K., Funke, B. H., Rehm, H. L., Olesen, M. S., Maron, B. A., Szolovits, P., Margulies, D. M., Loscalzo, J. & Kohane, I. S. (2016), 'Genetic misdiagnoses and the potential for health disparities', *New England Journal of Medicine* **375**(7), 655–665. PMID: 27532831.
- Mary, D. & Roquain, E. (2022), 'Semi-supervised multiple testing', *Electronic Journal of Statistics* **16**(2), 4926 4981.
- Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K. & Galstyan, A. (2021), 'A survey on bias and fairness in machine learning', ACM Comput. Surv. 54(6).
- Meinshausen, N. & Rice, J. (2006), 'Estimating the proportion of false null hypotheses among a large number of independently tested hypotheses.', Ann. Statist. 34, 373–393.
- Menon, A. K. & Williamson, R. C. (2018), The cost of fairness in binary classification, in S. A. Friedler & C. Wilson, eds, 'Proceedings of the 1st Conference on Fairness, Accountability and Transparency', Vol. 81 of Proceedings of Machine Learning Research, PMLR, pp. 107–118.
- Moya, M. M. & Hush, D. R. (1996), 'Network constraints and multi-objective optimization for one-class classification', *Neural networks* **9**(3), 463–474.
- Mukherjee, D., Yurochkin, M., Banerjee, M. & Sun, Y. (2020), Two simple ways to learn individual fairness metrics from data, *in* 'Proceedings of the 37th International Conference on Machine Learning', ICML'20, JMLR.org.
- Pleiss, G., Raghavan, M., Wu, F., Kleinberg, J. & Weinberger, K. Q. (2017), 'On fairness and calibration', arXiv preprint arXiv:1709.02012.

- Romano, Y., Barber, R. F., Sabatti, C. & Candès, E. (2020), 'With malice toward none: Assessing uncertainty via equalized coverage'. https://hdsr.mitpress.mit.edu/pub/qedrwcz3.
- Romano, Y., Bates, S. & Candès, E. J. (2020), Achieving equalized odds by resampling sensitive attributes, in 'Advances in Neural Information Processing Systems 33 (NIPS 2020)', Curran Associates, Inc. To appear.
- Saleiro, P., Kuester, B., Hinkson, L., London, J., Stevens, A., Anisfeld, A., Rodolfa, K. T. & Ghani, R. (2018), 'Aequitas: A bias and fairness audit toolkit'. URL: https://arxiv.org/abs/1811.05577
- Silverman, B. W. (1986), Density estimation for statistics and data analysis / B.W. Silverman, Chapman and Hall London; New York.
- Storey, J. D. (2002), 'A direct approach to false discovery rates', J. Roy. Statist. Soc. B 64, 479–498.
- Storey, J. D. (2003), 'The positive false discovery rate: a Bayesian interpretation and the q-value', Ann. Statist. 31, 2013–2035.
- Storey, J. D., Taylor, J. E. & Siegmund, D. (2004), 'Strong control, conservative point estimation and simultaneous conservative consistency of false discovery rates: a unified approach', J. Roy. Statist. Soc. B 66(1), 187–205.
- Sun, W. & Cai, T. T. (2007), 'Oracle and adaptive compound decision rules for false discovery rate control', J. Amer. Statist. Assoc. 102, 901–912.
- Sun, W. & Wei, Z. (2011), 'Large-scale multiple testing for pattern identification, with applications to time-course microarray experiments', J. Amer. Statist. Assoc. 106, 73–88.
- Vovk, V., Gammerman, A. & Shafer, G. (2005), Algorithmic learning in a random world, Vol. 29, Springer.
- Weinstein, A., Barber, R. & Candes, E. (2017), 'A power and prediction analysis for knockoffs with lasso statistics'. arXiv preprint arXiv:1712.06465.
- Zafar, M. B., Valera, I., Gomez Rodriguez, M. & Gummadi, K. P. (2017), Fairness beyond disparate treatment & disparate impact: Learning classification without disparate mistreatment, WWW '17, International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, CHE, p. 1171–1180. URL: https://doi.org/10.1145/3038912.3052660
- Zeng, X., Dobriban, E. & Cheng, G. (2022), Fair bayes-optimal classifiers under predictive parity, *in* 'Advances in Neural Information Processing Systems', Vol. 35, Curran Associates, Inc., pp. 27692–27705.

Supplementary Material

This supplement provides alternative R-value notions (Section A), additional technical details of the methodology (Sections B-C), technical proofs (Sections D-E), discussion of related fairness algorithms (Section \mathbf{F}), and supplementary numerical results (Section \mathbf{G}).

A Alternative Versions of the R-value

A.1 General considerations

In this paper, we propose the utilization of the R-value presented in (13), which is adjusted from (12), as the basic operational unit of our FASI algorithm. This section discusses some limitations of this definition and offers alternative versions of the R-value that may be of either practical or theoretical interest.

Firstly, one advantage of the R-value (12) lies in its inclusion of both \mathcal{D}^{cal} and \mathcal{D}^{test} in the denominator, which enhances the algorithm's stability. This is particularly valuable when m, the size of the test set, is small. However, the inclusion of both \mathcal{D}^{cal} and \mathcal{D}^{test} introduces additional complexity to the theory. Specifically, the resulting FASI algorithm can only control a modified version of the FSR, which includes a "+1" adjustment in the denominator compared to the original FSR defined in (7). Although this is a minor concern, we provide a simplified version (referred to as the "vanilla version") of the R-value in Section A.2. This version of the R-value, which only includes \mathcal{D}^{test} in the denominator, will be particularly relevant for readers who prefer a theory that eliminates the +1 adjustment in the FSR definition.

Secondly, one consideration, pertaining to the multiplicative factor $\gamma_{c,a} = \mathbb{E}\left(p_{c,null}^{test,a}/p_{c,null}^{cal,a}\right)$ in the theorem, is that FASI fails to provide precise FSR control due to the possible fluctuations in $\gamma_{c,a}$. While this issue is also minor (as under Assumption 1, this constant is approximately 1 and numerically negligible, cf. Section G.2 of this Supplement), we present a conservative version of the R-value in Section A.3. We demonstrate that this multiplicative factor $\gamma_{c,a}$ can be eliminated from the theory when the conservative version is employed. It is important to note that the conservative R-value is primarily of theoretical interest, as it leads to a substantial loss in power in many practical scenarios.

Finally, in situations where multiple classes are being selected, the R-value rule may result in overlapping selections. To mitigate this issue, we introduce a modified definition of the R-value in Section A.4 to address such undesirable scenarios.

A.2 The vanilla version

We modify the quantity Q_k^c in Equation 12 by considering only \mathcal{D}^{test} in the denominator:

$$Q_k^{c,\chi} = \sum_{a \in \mathcal{A}} \mathbb{I}(A_k = a) \cdot \frac{\frac{1}{n_a^{cal} + 1} \left\{ \sum_{i \in \mathcal{D}^{cal}} \mathbb{I}\left(A_i = a, \hat{S}_i^c \ge \hat{S}_k^c, Y_i \neq c\right) + 1 \right\}}{\frac{1}{m_a} \left[\sum_{j \in \mathcal{D}^{test}} \mathbb{I}\left(A_j = a, \hat{S}_j^c \ge \hat{S}_k^c\right) \right]} \land 1.$$
(A.1)

This calculation is performed for all $k \in \mathcal{D}^{cal} \cup \mathcal{D}^{test}$. We further adjust $Q_k^{c,\chi}$, as done in Equation 13, to enforce the monotonicity property:

$$R_j^{c,\chi} \equiv \min_{\{k \in \mathcal{D}^{cal} \cup \mathcal{D}^{test}: \hat{S}_k^c \le \hat{S}_j^c\}} Q_k^{c,\chi}, \quad j \in \mathcal{D}^{test}.$$
 (A.2)

This leads to our original proposal of the R-value (the vanilla version).

The next theorem shows that the FASI algorithm, implemented using the vanilla version of R-value, is valid for FSR control.

Theorem 2. Define $\gamma_{c,a} = \mathbb{E}\left(p_{c,null}^{test,a}/p_{c,null}^{cal,a}\right)$, where $p_{c,null}^{test,a}$ and $p_{c,null}^{cal,a}$ are the proportions of individuals in group a that do not belong to class c in the test and calibration data, respectively. Then under Assumption 1, we have, for all $a \in \mathcal{A}$, the FASI algorithm with R-value (A.2) satisfies $FSR_a^{\{c\}} \leq \gamma_{c,a}\alpha_c$.

Theorem 2 closely resembles Theorem 1 discussed in the main text. However, it is important to note that in Theorem 2, the denominator of the R-value only incorporates \mathcal{D}^{test} . As a result, we achieve FSR control on the desired quantity (7), rather than the modified quantity (17). In practical scenarios, the disparity between these two FSRs is negligible, which is why we have chosen to solely present the more robust and stable R-value (12) in the main text.

A.3 The conservative version

The two versions of the R-value, presented in (A.2) and (13) [adjusted from (A.1) and (12)] respectively, only provide FSR control up to a multiplicative factor $\gamma_{c,a} = \mathbb{E}\left(p_{c,null}^{test,a}/p_{c,null}^{cal,a}\right)$. While this factor is typically close to 1 under Assumption 1, the subsequent corollary provides a conservative version of the R-value for practitioners who wish to exclude the stochastic variation in $\gamma_{c,a}$ from their analysis. The proof of the corollary follows directly from the proofs of Theorems 1 and 2 and is therefore omitted.

Corollary 1. Suppose we apply the FASI algorithm with the conservative R-values:

$$\hat{R}_{j}^{c,\chi} = \frac{n_{a}^{cal} + 1}{n_{a,null}^{cal,c} + 1} R_{j}^{c,\chi}$$
(A.3)

for $j \in \mathcal{D}^{test}$, where $n_a^{cal} = \sum_{i \in \mathcal{D}^{cal}} \mathbb{I}(A_i = a)$ and $n_{a,null}^{cal,c} = \sum_{i \in \mathcal{D}^{cal}} \mathbb{I}(A_i = a, Y_i \neq c)$. Further define $n_a^{test} = \sum_{j \in \mathcal{D}^{test}} \mathbb{I}(A_j = a)$ and $n_{a,null}^{test,c} = \sum_{j \in \mathcal{D}^{test}} \mathbb{I}(A_j = a, Y_j \neq c)$. Then we have (a)

$$\textit{FSR}_a^{\{c\}} \leq \mathbb{E}\left(n_{a,null}^{\textit{test},c}/n_a^{\textit{test}}\right) \alpha \leq \alpha$$

for all $a \in \mathcal{A}$ when $\hat{R}_i^{c,\chi}$ are used in the FASI algorithm.

Corollary 1 shows that all group-wise FSRs are strictly controlled at less than or equal to the nominal level α . The ratio $n_{a,null}^{test,c}/n_a^{test}$, referred to as the null proportion in multiple testing, also appears in the classical Benjamini-Hochberg (BH) procedure for FDR control. In Section C.2 of this supplement, we will elaborate the connection between the FASI algorithm implemented with conservative R-values (A.3) and the BH algorithm implemented with conformal p-values.

It is expected that the FASI algorithm with conservative R-values (A.3) can be enhanced by incorporating the unknown ratio $n_{a,null}^{test}/n_a^{test}$ into the analysis. This approach has been successfully adopted in various works, such as Benjamini & Hochberg (2000) and Storey (2002), to boost the power of the conservative BH algorithm in the context of FDR control. The FASI algorithm with adjusted R-values in (13), derived from (12), can be regarded as such an approach. Specifically, the unknown ratio $n_{a,null}^{test,c}/n_a^{test}$ is initially estimated as $(n_{a,null}^{cal,c} + 1)/(n_a^{cal} + 1)$. This estimated ratio is then incorporated into the FASI algorithm by utilizing the conservative version of FASI at the modified level of $(n_a^{cal} + 1)/(n_{a,null}^{cal,c} + 1)\alpha$. This practice leads to improved power at the expense of the additional factor $\gamma_{c,a}$ in Theorems 1 and 2.

Finally, we note that our focus in this article is not on the conservative R-values since they tend to result in a higher proportion of indecisions (or decreased power). Our recommended R-values, as presented in (12) and adjusted to (13), on the other hand, are intuitive, powerful, and provide almost exact control in practical situations.

A.4 The multi-class version

When considering the selection of multiple classes, additional complexities arise, necessitating further modifications to our previous R-values. Let $\mathcal{C}' \subset \mathcal{C}$ represent the set of classes to be selected. With

indecisions being allowed, the action space is $\Lambda = \{0, \mathcal{C}'\}$. We denote the selection rule for *m* individuals in the test set as $\{\hat{Y}_j : j \in \mathcal{D}^{test}\} \in \Lambda^m$.

The FSR can be defined in two ways with respect to the subset C'. The first definition evaluates the fraction of incorrect selections for each individual class separately:

$$\operatorname{FSR}_{a}^{\{c\}} = \mathbb{E}\left[\frac{\sum_{j \in \mathcal{D}^{test}} \mathbb{I}(\hat{Y}_{j} = c, Y_{j} \neq c, A_{j} = a)}{\left\{\sum_{j \in \mathcal{D}^{test}} \mathbb{I}(\hat{Y}_{j} = c, A_{j} = a)\right\} \lor 1}\right], \text{ for all } a \in \mathcal{A} \text{ and } c \in \mathcal{C}'.$$

By contrast, the second definition calculates an overall error rate by combining selections from all classes in C':

$$\operatorname{FSR}_{a}^{\mathcal{C}'} = \mathbb{E}\left[\frac{\sum_{j \in \mathcal{D}^{test}} \mathbb{I}(\hat{Y}_{j} \in \mathcal{C}', \hat{Y}_{j} \neq Y_{j}, A_{j} = a)}{\left\{\sum_{j \in \mathcal{D}^{test}} \mathbb{I}(\hat{Y}_{j} \in \mathcal{C}', A_{j} = a)\right\} \lor 1}\right], \text{ for all } a \in \mathcal{A}$$

We will focus on the binary classification scenario and adopt the first definition of FSR. A potential issue arises when controlling for $\text{FSR}_a^{\{c\}}$ for both c = 1 and c = 2, where simply comparing the R-value with α_c may result in overlapping selections. To simplify the discussion, we assume that the selection criterion remains consistent across both classes, denoted as $\alpha_c = \alpha$ for c = 1, 2. Our proposed procedure consists of two steps:

- 1. Step 1: Define $R_j^* = \min(R_j^1, R_j^2)$ and $c_j^* = \arg\min_{c \in \{1,2\}} R_j^c$.
- 2. Step 2: The decision rule is given by $\left\{ \hat{Y}_j = c_j^* \cdot \mathbb{I}(R_j^* \leq \alpha), j \in \mathcal{D}^{test} \right\}$.

This two-step procedure effectively avoids assigning an individual to multiple classes. Specifically, the procedure first identifies the most promising class based on the minimum R-value, and then makes a decision on whether the individual should be selected into this most promising class.

The second definition of FSR introduces several complicated issues. Firstly, it requires the employment of a new score function to achieve optimality under the oracle setting. Secondly, substantial adjustments must be made to the mirror process described in Section 3.2. Thirdly, the development of martingale theories becomes notably more intricate. Finally, when dealing with scenarios involving more than two classes, an additional layer of complexity arises. These various issues offer intriguing and crucial avenues for future exploration and research.

A.5 Summary

The aforementioned modifications are provided for addressing various considerations and limitations regarding the R-value within the FASI framework. To ensure conciseness and clarity, we have chosen to focus on only one version of the R-value (13), adjusted from (12), in the main text. This version strikes a balance between simplicity, stability, and power, and we recommend its practical utilization. However, we acknowledge that the alternative versions of R-values possess their own merits, be it from a practical or theoretical standpoint. All these alternative versions, when utilized within our FASI algorithm, are provably valid for FSR control.

B Theoretical R-value and Optimality Theory

In this section, we introduce the theoretical R-value and derive the optimal score function under a simplified setup. Our theory provides practical insights for practitioners on how to train score functions to construct informative R-values.

B.1 The mixture model under an oracle setting

Our subsequent discussions are purely theoretical, where we assume an oracle with access to all distributional information and make several simplifying assumptions. Our primary goal is to develop a theoretical version of the R-value and an optimality theory for FSR control. This theoretical framework serves as a foundation for our practical algorithm and provides valuable insights into the properties of the R-value.

Our discussions are centered around the following random mixture model:

$$F(x) = \sum_{a \in \mathcal{A}} \left\{ \pi_{1|a} F_{1,a}(x) + \pi_{2|a} F_{2,a}(x) \right\},$$
(B.4)

where $F_{1,a}(x)$ and $F_{2,a}(x)$ are the conditional CDFs of X from classes 1 and 2, respectively. Let $f_{c,a}(x)$ be the corresponding density functions of $F_{c,a}$. The probabilities $\pi_{c|a} = \mathbb{P}(Y = c, A = a)$ represent the joint probabilities of Y = c and A = a for c = 1, 2. Denote $\pi_a = P(A = a), \pi_{c|a} = P(Y = c|A = a)$. For our analysis, we consider a selection rule of the form $\hat{Y}(t) = c \cdot \mathbb{I}(S^c \ge t)$, where t is a threshold, and $\mathbb{I}(\cdot)$ is the indicator function. We assume that an oracle has knowledge of the conditional probabilities and conditional CDFs defined above.

B.2 The conversion algorithm

In this section, we present a systematic approach for converting an arbitrary score $S^{c}(x, a)$ into a fair score $R^{c}(S^{c})$, which we refer to as the theoretical R-value. Although the discussion is theoretical in nature, it highlights the existence of a fair score that corresponds to every confidence score. This algorithm can be regarded as a method of *calibration by group*, a widely used technique in the fairness literature (see Barocas et al. (2017) for an example).

The conversion algorithm consists of three steps. In Step 1, we find the distributional information with respect to the score S^c . Let $G^c(s) = \sum_{a \in \mathcal{A}} \pi_a G^c_a(s)$ be the CDF of S^c , where

$$G_a^c(s) = \pi_{1|a} G_{1,a}^c(s) + \pi_{2|a} G_{2,a}^c(s),$$

with $G_{1,a}^c(s)$ and $G_{2,a}^c(s)$ denoting the conditional CDFs of S^c given A = a and Y, and $\pi_{c|a} = \mathbb{P}(Y_i = c|A_i = a)$ being the conditional probabilities, for c = 1, 2.

Suppose an oracle knows the conditional probabilities and conditional CDFs defined above. In Step 2, we compute the conditional error probability when the threshold is t:

$$Q_a^c(t) \coloneqq \mathbb{P}(Y \neq c | S^c \ge t, A = a) = \begin{cases} \frac{\pi_{2|a} \left\{ 1 - G_{2,a}^1(t) \right\}}{1 - G_a^1(t)}, & c = 1; \\ \frac{\pi_{1|a} \left\{ 1 - G_{1,a}^2(t) \right\}}{1 - G_a^2(t)}, & c = 2. \end{cases}$$

Finally, in Step 3 we compute a fair score, referred to the theoretical R-value, for an individual from group a with observed score $S^c = s$:

$$R^{c}(s) = \inf_{t \leq s} \left\{ Q^{c}_{a}(t) \coloneqq \mathbb{P}(Y \neq c | \hat{Y}(t) = c, A = a) \right\},\tag{B.5}$$

where if the threshold is t, $Q_a^c(t)$ corresponds to the conditional error probability. If the confidence score satisfies the monotone likelihood ratio condition (MLRC, Sun & Cai 2007) then the infimum is achieved at s exactly; see Section E.1 for related discussions.

B.3 Theoretical R-value and fairness

Consider random mixture model (B.4). Suppose an oracle knows the score function $S^c(x, a) = \mathbb{P}(Y = c | X = x, A = a)$. The goal is to assign labels "0", "1" and "2" to new instances $\{(X_j, A_j) : j \in \mathcal{D}^{test}\}$. We assume that the instances (X_j, A_j) are independent draws from an underlying distribution F(x, a).

Define the marginal FSR

$$\mathrm{mFSR}_{a}^{\{c\}} = \frac{\mathbb{E}\left\{\sum_{j\in\mathcal{D}^{test}}\mathbb{I}(\hat{Y}_{j}=c,Y_{j}\neq c,:A_{j}=a)\right\}}{\mathbb{E}\left\{\sum_{j\in\mathcal{D}^{test}}\mathbb{I}(\hat{Y}_{j}=c,A_{j}=a)\right\}}.$$

Under the random mixture model (B.4), it can be shown that, following arguments in Storey (2003) for FDR analysis,

$$\mathrm{mFSR}_{a}^{\{c\}} = \mathbb{P}(Y \neq c | \hat{Y} = c, A = a), \tag{B.6}$$

which is the conditional probability required in the sufficiency principle.

Based on the work of Cai et al. (2019), we can similarly show that under mild conditions,

$$\operatorname{FSR}_{a}^{\{c\}} = \operatorname{mFSR}_{a}^{\{c\}} + o(1), \text{ when } m_a \coloneqq |\{j \in \mathcal{D}^{test} : A_j = a\}| \to \infty.$$
(B.7)

The next proposition, which follows directly from (B.5), shows that thresholding the theoretical R-value leads to a fair selective inference procedure.

Proposition 1. Consider a classifier that claims $\hat{Y} = c$ if $R^c \leq \alpha$. Then

$$\mathbb{P}(Y \neq c | \hat{Y} = c, A = a) \le \alpha \text{ for all } a \in \mathcal{A}.$$
(B.8)

We would like to make two important remarks. Firstly, the theoretical R-value, which may be viewed as the counterpart of the data-driven R-value, represents the minimum conditional probability required to ensure that an individual with score $S^c = s$ is selected into class c. Secondly, the theoretical R-value is a fundamental quantity that is closely linked to the sufficiency principle in the fairness literature. Proposition 1 highlights that by setting thresholds for the theoretical R-values, the thresholding procedure fulfills the sufficiency principle and controls the group-wise error rates.

B.4 A sketch of the optimality theory

We present and prove an intuitive result that shows $S^c(x, a) = \mathbb{P}(Y = c | X = x, A = a)$ is the optimal choice of score function for calibrating the theoretical R-value. To simplify the arguments, we develop our optimality theory based on the mFSR, an asymptotically equivalent variation of the FSR. The relationship between the mFSR and FSR has been established in Equation (B.7).

We aim to construct a selection rule under the binary classification setting that solves the following constrained optimization problem:

Minimize the EPI, subject to
$$mFSR_a^c \le \alpha_c, c = 1, 2$$
 for all $a \in \mathcal{A}$. (B.9)

Here, we denote $S_j^c = \mathbb{P}(Y_j = c | X_j = x_j, A_j = a_j)$, and the scores can be transformed to theoretical R-values, denoted R_j^1 and R_j^2 . The process of conversion follows the general strategy outlined in Section B.2, and is described in more detail in the proof of Theorem 3 below.

Define the oracle procedure

$$\boldsymbol{\delta}^{OR} = \{\delta_j^{OR} : j \in \mathcal{D}^{test}\}, \text{ where } \delta_j^{OR} = \mathbb{I}(R_j^1 \le \alpha_1) + 2\mathbb{I}(R_j^2 \le \alpha_2).$$
(B.10)

The optimality of the oracle procedure is established in the next theorem.

Theorem 3. Consider random mixture model (B.4). Assume that α_1 and α_2 have been properly chosen such that (B.10) does not have overlapping selections. Let $\mathcal{D}_{\alpha_1,\alpha_2}$ denote the collection of selection rules that satisfy $mFSR_a^c \leq \alpha_c$ for c = 1, 2 and all $a \in \mathcal{A}$. Let $EPI_{\boldsymbol{\delta}}$ denote the EPI of an arbitrary decision rule $\boldsymbol{\delta}$. Then the oracle procedure (B.10) is optimal in the sense that $EPI_{\boldsymbol{\delta}_{OR}} \leq EPI_{\boldsymbol{\delta}}$ for any $\boldsymbol{\delta} \in \mathcal{D}_{\alpha_1,\alpha_2}$.

The optimality theory indicates that, during the training stage, we should utilize *all features*, including the sensitive attribute A, to best capture individual level information.

B.5 Connections to Storey's q-value

The theoretical R-value is closely connected to the q-value, a useful tool in large-scale testing due to its intuitive interpretation and ease of use, as described in Storey (2003).

To test hypotheses $\{H_j : j \in \mathcal{D}^{test}\}$ with associated p-values $\{p_j : j \in \mathcal{D}^{test}\}$, let π be the proportion of non-nulls and G(t) the alternative distribution of p-values. The q-value for hypothesis H_j is defined as

$$\inf_{t \ge p_j} \left\{ \text{pFDR}(t) \coloneqq \frac{(1-\pi)t}{(1-\pi)t + \pi G(t)} \right\},\,$$

which roughly measures the fraction of false discoveries when H_i is rejected.

The q-value and R-value algorithms operate in the same manner. Conducting an FDR analysis at a given level α entails obtaining the q-value for hypothesis j and rejecting it if the q-value is less than or equal to α . Likewise, conducting an FSR analysis at level α involves obtaining the R-value for individual j and selecting it if the R-value is less than or equal to α .

C R-value and Conformal P-value

In this section, we adopt a multiple testing perspective to gain further insights into the R-value. Despite its distinct motivation, we demonstrate that the R-value can be derived as the (BH) q-value of the conformal p-values (Bates et al. 2023) in the context of one-class classification scenario. For comparability purposes, we exclude the sensitive attribute A and concentrate on the vanilla version of the R-value defined by (A.1) and (A.2) in Section A.2.

C.1 A brief review of conformal p-values

The problem of one-class classification, also known as outlier detection or out-of-distribution testing in conformal inference, can be formulated within the framework of selective inference. Consider the observed data $\{(X_i, Y_i) : i \in \mathcal{D}\}$ originating from two classes, $Y_i = 1$ and $Y_i = 2$. We divide the set \mathcal{D} into two subsets, $\mathcal{D}^c = \{i : Y_i = c\}, c = 1, 2$, where \mathcal{D}^1 and \mathcal{D}^2 represent the index sets of inliers and outliers, respectively.

In the context of outlier detection, the objective is to accurately identify outliers (individuals with label Y = 2) in a set of unlabeled test data $\{X_j : j \in \mathcal{D}^{test}\}$, while maintaining strict control over the error rate. By considering individuals in class "1" as the null cases, we can formulate an equivalent multiple testing problem as follows:

$$H_{j0}: Y_j = 1$$
 vs. $H_{j1}: Y_j \neq 1$ (i.e., $Y_j = 2$), for $j \in \mathcal{D}^{test}$.

As explained in Section 2.3, the false selection rate $FSR^{\{2\}}$, as defined in (5), precisely yields the widely used FDR, the expected fraction of false rejections among all rejections.

Remark 3. In the context of outlier detection, the standard practice is to only consider the labeled inliers \mathcal{D}^1 when computing conformal p-values. It is worth noting that recent research by Liang et al. (2024) has revealed that this approach may result in potential efficiency loss. Nevertheless, for the sake of comparability, we adhere to the conventional practice and exclude \mathcal{D}^2 in our investigation.

The construction of split-conformal p-values (Bates et al. 2023) involves partitioning \mathcal{D}^1 into two subsets: \mathcal{D}^{train} for training a score function and \mathcal{D}^{cal} for calibrating a significance index. Treating $\hat{S}^2(\cdot)$ as a conformity score function that indicates the likelihood of belonging to class 2, the conformal p-value for testing H_{i0} can be expressed using our notation as:

$$\hat{u}_j \equiv \hat{u}(X_j) = \frac{\sum_{i \in \mathcal{D}^{cal}} \mathbb{I}\{\hat{S}^2(X_i) \ge \hat{S}^2(X_j)\} + 1}{n^{cal} + 1}.$$
(C.11)

Remark 4. To avoid confusion, it's worth noting that in our setup, a higher score indicates a higher likelihood of being an outlier. This interpretation is contrary to the convention in Bates et al. (2023). In order to align the definitions, we have swapped the expression " $S \leq t$ " from the conformal p-value definition in Bates et al. (2023) with " $S \geq t$ " in our formula (C.11). This adjustment ensures equivalence between the two formulations.

C.2 R-value is the BH q-value of conformal p-values (for outlier detection)

To see the connection of our R-value to the conformal p-value (C.11), recall the definition of Storey's q-value

$$\hat{q}^{ST}\{\hat{u}(t)\} = (1-\pi)\hat{u}(t)/G\{\hat{u}(t)\},\$$

where π is the proportion of non-null cases in \mathcal{D}^{test} and $G(\cdot)$ is the cumulative distribution function (CDF) of the p-values. Now recall $m = |\mathcal{D}^{test}|$, let $\hat{G}(t)$ denote the empirical process of the scores $\{\hat{S}_{j}^{2} : j \in \mathcal{D}^{test}\}$:

$$\hat{G}(t) = \frac{1}{m} \sum_{j \in \mathcal{D}^{test}} \mathbb{I}\left\{\hat{u}(\hat{S}_j^2) \le \hat{u}(t)\right\} = \frac{1}{m} \sum_{j \in \mathcal{D}^{test}} \mathbb{I}\left(\hat{S}_j^2 \ge t\right),\tag{C.12}$$

where the last equality holds because, by (C.11), a larger score corresponds to a smaller conformal p-value. Next we consider a modification of Storey's q-value, referred to as the BH q-value, which ignores the $(1 - \pi)$ term and substitutes \hat{G} in place of G in Storey's q-value:

$$\hat{q}_i^{BH} = \frac{\hat{u}(\hat{S}_i^2)}{\hat{G}(\hat{S}_i^2)} \quad \text{for } i \in \mathcal{D}^{cal} \cup \mathcal{D}^{test}.$$
(C.13)

Combining (C.11) - (C.13), we have

$$\hat{q}^{BH}(t) = \frac{m}{n^{cal}+1} \cdot \frac{\sum_{i \in \mathcal{D}^{cal}} \mathbb{I}\left(\hat{S}_i^2 \ge t\right) + 1}{\sum_{j \in \mathcal{D}^{test}} \mathbb{I}\left(\hat{S}_j^2 \ge t\right)} = \frac{m}{n^{cal}+1} \cdot \frac{\sum_{i \in \mathcal{D}^{cal}} \mathbb{I}\left(\hat{S}_i^2 \ge t, Y_i = 1\right) + 1}{\sum_{j \in \mathcal{D}^{test}} \mathbb{I}\left(\hat{S}_j^2 \ge t\right)}.$$
 (C.14)

The last equality holds because under the one-class classification setup, \mathcal{D}^{cal} is a "pure" training set in which all observations are from the null class "1". Let t take values in $\{S_k : k \in \mathcal{D}^{cal} \cup \mathcal{D}^{test}\}$ and denote the corresponding values $\{\hat{q}_k^{BH} : k \in \mathcal{D}^{cal} \cup \mathcal{D}^{test}\}$.

We also need to apply a monotonicity adjustment to ensure that the q-value function is nondecreasing in the conformity score. Following (13) of Section 3.1, we let

$$q_j^{BH} = \min_{k \in \mathcal{D}^{cal} \cup \mathcal{D}^{test}: \hat{S}_k^2 \le \hat{S}_j^2} \hat{q}_k^{BH}, \quad \text{for } j \in \mathcal{D}^{test}.$$
(C.15)

This precisely recovers the vanilla version of our R-value.

C.3 Discussion

We emphasize that the fundamental connection between the R-value and conformal q-values only holds under the one-class classification setup. The BH q-value (C.14) will be different from the R-value (16) under the binary classification setup that we have considered in this article. Specifically, the cardinalities of the calibration sets will be different under the two setups, and the equality (C.14) does not hold. Our R-value does not explicitly utilize conformal p-values under the binary classification setup.

The conformal p-value approach by Bates et al. (2023) remains applicable for selective inference in the binary classification setup, specifically for the selection of cases from class 2. Nevertheless, it is noteworthy that the conformal p-value method utilizes a smaller data set, as the data set \mathcal{D}^2 is discarded, in comparison to our R-value approach. Consequently, this may lead to suboptimal information utilization and a reduction in statistical power. In addition, it is worth noting that the FASI algorithm may not be well-suited for the outlier detection problem, as it presumes that the test data and calibration data are exchangeable, which is unlikely to hold in practical scenarios. Therefore, both the conformal p-value and FASI approaches would require modification to address the outlier detection problem with labeled outliers. Related issues have gone beyond the scope of this study and will be pursued in future research.

D Proof of Theorems 1 and 2

We begin by presenting the proof of Theorem 2 that involves the vanilla version of the R-value, followed by the more intricate proof of Theorem 1 with the R-value in the main text.

D.1 Proof of Theorem 2

D.1.1 An empirical process description of the FASI algorithm

Denote $\hat{S}_i^c := \hat{S}^c(X_i = x_i, A_i = a_i)$, for $i \in \mathcal{D}^{cal} \cup \mathcal{D}^{test}$. Suppose we select subjects into class c if the confidence score is greater than or equal to a threshold t. The estimated false discovery proportion (FSP), as a function of t, in group a can be described as the following empirical process:

$$\hat{Q}_{c}^{a}(t) = \frac{\frac{1}{n_{a}^{cal}+1} \left\{ \sum_{i \in \mathcal{D}^{cal}} \mathbb{I}(\hat{S}_{i}^{c} \ge t, Y_{i} \neq c, A_{i} = a) + 1 \right\}}{\frac{1}{m_{a}} \left\{ \sum_{j \in \mathcal{D}^{test}} \mathbb{I}(\hat{S}_{j}^{c} \ge t, A_{j} = a) \lor 1 \right\}}.$$
(D.16)

Let $\mathbf{S} = \{\hat{S}_i^c : i \in \mathcal{D}^{cal} \cup \mathcal{D}^{test}\}.$

Consider a selection procedure represented by the empirical process given in Equation (D.16). This selection procedure aims to find the smallest threshold, denoted as τ , for which the estimated FSP is less than or equal to α .

$$\tau = \hat{Q}_c^{a,-1}(\alpha) = \min\left\{t \in \mathbf{S} : \hat{Q}_c^a(t) \le \alpha\right\}.$$
(D.17)

The proof in this section centers on the vanilla version of the R-value defined through (A.1) and (A.2). For conciseness in notation, we omit the χ superscript in subsequent discussions. It is easy to see that

$$R_j^c \equiv \min_{\{t \in \mathbf{S}: t \le \hat{S}_j^c\}} \hat{Q}_c^a(t), \quad j \in \mathcal{D}^{test}.$$
(D.18)

Now we show that the decision rule based on thresholding the R-value (D.18) can be equivalently represented using a decision rule based on thresholding the confidence scores $\{S_j^c, j \in \mathcal{D}^{test}\}$. This conclusion is formally stated in the following lemma. The same arguments are equally applicable to the R-value defined by (13).

Lemma 1. The FASI algorithm can be represented in two equivalent ways:

$$\delta_j = \mathbb{I}(R_j^c \le \alpha) \quad \Longleftrightarrow \quad \delta'_j = \mathbb{I}(\hat{S}_j^c \ge \tau), \quad j \in \mathcal{D}^{test}, \tag{D.19}$$

where τ is defined in (D.17).

Proof of Lemma 1. First, suppose that $\mathbb{I}\{\hat{S}_j^c \geq \tau\} = 1$ holds for some $j \in \mathcal{D}^{test}$. By the definition of τ , we have that $\tau \in \{t \in \mathbf{S} : t \leq \hat{S}_j^c\}$ and

$$R_j^c = \min_{\{t \in \mathbf{S}: t \le \hat{S}_j^c\}} \hat{Q}_c^a(t) \le \hat{Q}_c^a(\tau) \le \alpha$$

It follows that $\mathbb{I}\{R_j^c \leq \alpha\} = 1.$

Next, suppose that $\mathbb{I}\{R_j^c \leq \alpha\} = 1$. By the definition of R_j^c , we have $\min_{\{t \in \mathbf{S}: t \leq \hat{S}_i^c\}} \hat{Q}_c^a(t) \leq \alpha$. That is, there exists a threshold $t \leq \hat{S}_j^c$ in **S** such that $\hat{Q}_c^a(t) \leq \alpha$. It follows that

$$\tau = \min\left\{t \in \mathbf{S} : \hat{Q}_c^a(t) \le \alpha\right\} \le \min\left\{t \in \mathbf{S}, t \le \hat{S}_j^c : \hat{Q}_c^a(t) \le \alpha\right\} \le \hat{S}_j^c,$$

implying that $\mathbb{I}\{\hat{S}_j^c \geq \tau\} = 1$. Combining the two arguments above, we have

$$\mathbb{I}(R_j^c \le \alpha) = 1 \Longleftrightarrow \mathbb{I}(\hat{S}_j^c \ge \tau) = 1.$$

We can similarly show that $\mathbb{I}(R_j^c \leq \alpha) = 0 \iff \mathbb{I}(\hat{S}_j^c \geq \tau) = 0$, establishing (D.19).

We will now proceed to describe the true FSP process of the FASI algorithm, as outlined in (D.16) and the second representation in (D.19) through S^c . Let

$$V^{t}(t) = \sum_{j \in \mathcal{D}^{test}} \mathbb{I}(\hat{S}_{j}^{c} \ge t, Y_{j} \ne c, A_{j} = a),$$

$$W^{t}(t) = \sum_{j \in \mathcal{D}^{test}} \mathbb{I}(\hat{S}_{j}^{c} \ge t, Y_{j} = c, A_{j} = a),$$

$$R^{t}(t) = V^{t}(t) + W^{t}(t) = \sum_{j \in \mathcal{D}^{test}} \mathbb{I}(\hat{S}_{j}^{c} \ge t, A_{j} = a)$$

denote the counts of false selections, correct selections and total selections, respectively, in the test set \mathcal{D}^{test} when the threshold is t. The sensitive attribute "a" has been omitted from the expressions of V^t and R^t for notation simplicity. Additionally, define

$$\begin{split} V^c(t) &= \sum_{i \in \mathcal{D}^{cal}} \mathbb{I}(\hat{S}_i^c \ge t, Y_i \ne c, A_i = a), \\ W^c(t) &= \sum_{i \in \mathcal{D}^{cal}} \mathbb{I}(\hat{S}_i^c \ge t, Y_i = c, A_i = a), \\ R^c(t) &= V^c(t) + W^c(t) = \sum_{i \in \mathcal{D}^{cal}} \mathbb{I}(\hat{S}_i^c \ge t, A_i = a) \end{split}$$

as the corresponding counts in the calibration set \mathcal{D}^{cal} . Consider the data-driven threshold τ defined in (D.17). Then the FSP of the proposed FASI algorithm, as defined by the second representation in (D.19), can be computed as

$$\operatorname{FSP}_{a}^{\{c\}}(\tau) = \frac{V^{t}(\tau)}{R^{t}(\tau) \vee 1}.$$

In what follows, we always consider the quantities in group "A = a" and omit the superscript "a" in the remaining expressions for notational simplicity. Moreover, instead of using m_a and n_a^{cal} , we use notations $|\mathcal{D}^{cal}|$ and $|\mathcal{D}^{test}|$ (dropping "a") to reflect that the cardinality of the sets is random. The operation of the FASI algorithm implies that

$$FSP_a^{\{c\}}(\tau) = \frac{V^t(\tau)}{V^c(\tau) + 1} \cdot \frac{V^c(\tau) + 1}{R^t(\tau) \vee 1}$$
$$= \hat{Q}_a^c(\tau) \cdot \frac{|\mathcal{D}^{cal}| + 1}{|\mathcal{D}^{test}|} \cdot \frac{V^t(\tau)}{V^c(\tau) + 1}$$
$$\leq \alpha \cdot \frac{|\mathcal{D}^{cal}| + 1}{|\mathcal{D}^{test}|} \cdot \frac{V^t(\tau)}{V^c(\tau) + 1},$$
(D.20)

where the last two steps utilize definitions (D.16) and (D.17), respectively.

D.1.2 Martingale arguments

A crucial step in establishing FSR control, i.e., $\mathbb{E}\{FSP_a^{\{c\}}(\tau)\} \leq \alpha$, is to demonstrate that the ratio appearing in expression (D.20),

$$\frac{V^t(t)}{V^c(t)+1},\tag{D.21}$$

is a martingale. We start with the following continuous-time filtration:

$$\mathcal{F}_t = \sigma\{V^t(s), V^c(s), W^t(s), W^c(s) : t_l \le s \le t\},\$$

where t_l represents the lower limit of the threshold. That is, if t_l is employed, all subjects are classified into class c. It is equivalent to write $\mathcal{F}_t = \sigma\{V^t(s), V^c(s), R^t(s), R^c(s) : t_l \leq s \leq t\},\$

However, in our proof, it is sufficient to consider a discrete-time filtration since the FASI algorithm only selects thresholds from $\mathbf{S} = \{\hat{S}_i^c : i \in \mathcal{D}^{test} \cup \mathcal{D}^{cal}\}$. Let $m^* = |\mathcal{D}^{cal}| + |\mathcal{D}^{test}|$ denote the total number of selections in both \mathcal{D}^{cal} and \mathcal{D}^{test} when the threshold is t_l . We consider a σ -field that contains all information of the entire selection process. Consider a sequence of thresholds (times) $\{s_k : k = m^*, \dots, 1\}$, where s_k is the threshold when exactly k subjects, including those from both \mathcal{D}^{cal} and \mathcal{D}^{test} , are selected into class c, and k takes values in the order of $m^*, m^* - 1, \dots, 1$ (backward in time). This leads to the following discrete-time filtration:

$$\mathcal{F}_{k} = \sigma \left\{ V^{t}(s_{j}), V^{c}(s_{j}), W^{t}(s_{j}), W^{c}(s_{j}) : j = m^{*}, m^{*} - 1, \dots, k \right\}.$$
 (D.22)

We can see that \mathcal{F}_k is a backward-running filtration in the sense that for $k_1 < k_2$, $\mathcal{F}_{k_2} \subset \mathcal{F}_{k_1}$.

Note that at time s_k , only one of the four following events is possible:

$$\begin{split} A_1 &= \{ V^t(s_{k-1}) = V^t(s_k) - 1, V^c(s_{k-1}) = V^c(s_k), W^t(s_{k-1}) = W^t(s_k), W^c(s_{k-1}) = W^c(s_k) \}, \\ A_2 &= \{ V^t(s_{k-1}) = V^t(s_k), V^c(s_{k-1}) = V^c(s_k) - 1, W^t(s_{k-1}) = W^t(s_k), W^c(s_{k-1}) = W^c(s_k) \}, \\ A_3 &= \{ V^t(s_{k-1}) = V^t(s_k), V^c(s_{k-1}) = V^c(s_k), W^t(s_{k-1}) = W^t(s_k) - 1, W^c(s_{k-1}) = W^c(s_k) \}, \\ A_4 &= \{ V^t(s_{k-1}) = V^t(s_k), V^c(s_{k-1}) = V^c(s_k), W^t(s_{k-1}) = W^t(s_k), W^c(s_{k-1}) = W^c(s_k) - 1 \}. \end{split}$$

According to Assumption 1 which claims that the data points in \mathcal{D}^{cal} and \mathcal{D}^{test} are exchangeable, and the fact that FASI uses same fitted model to compute the scores, we have

$$\frac{\mathbb{P}(A_1|\mathcal{F}_k)}{\mathbb{P}(A_2|\mathcal{F}_k)} = \frac{V^t(s_k)}{V^c(s_k)}.$$
(D.23)

Moreover, we must have

$$\mathbb{P}(A_1|\mathcal{F}_k) + \mathbb{P}(A_2|\mathcal{F}_k) + \mathbb{P}(A_3|\mathcal{F}_k) + \mathbb{P}(A_4|\mathcal{F}_k) = 1.$$
(D.24)

It follows from (D.23) and (D.24) that there exists a γ_k , such that

$$\mathbb{P}(A_1|\mathcal{F}_k) = \gamma_k \cdot \frac{V^t(s_k)}{V^t(s_k) + V^c(s_k)},$$
$$\mathbb{P}(A_2|\mathcal{F}_k) = \gamma_k \cdot \frac{V^c(s_k)}{V^t(s_k) + V^c(s_k)},$$

and $\mathbb{P}(A_3|\mathcal{F}_k) + \mathbb{P}(A_4|\mathcal{F}_k) = 1 - \gamma_k$. It will soon become evident that the value of γ_k is inconsequential in the theory, as it will be canceled out in the calculations.

We show that the ratio defined in (D.21) is a discrete-time (backward) martingale with respect to

the filtration \mathcal{F}_k . This desired result can be easily established by noting that

$$\mathbb{E}\left\{\frac{V^{t}(s_{k-1})}{V^{c}(s_{k-1})+1}|\mathcal{F}_{k}\right\}$$

$$= \frac{V^{t}(s_{k})-1}{V^{c}(s_{k})+1}\cdot\gamma_{k}\cdot\frac{V^{t}(s_{k})}{V^{t}(s_{k})+V^{c}(s_{k})} + \frac{V^{t}(s_{k})}{V^{c}s_{k}}\cdot\gamma_{k}\cdot\frac{V^{c}(s_{k})}{V^{t}(s_{k})+V^{c}(s_{k})} + \frac{V^{t}(s_{k})}{V^{c}(s_{k})+1}(1-\gamma_{k})$$

$$= \frac{V^{t}(s_{k})}{V^{c}(s_{k})+1}\cdot(\gamma_{k}+1-\gamma_{k})$$

$$= \frac{V^{t}(s_{k})}{V^{c}(s_{k})+1}.$$

D.1.3 FSR Control

The threshold τ defined by (D.17) is a stopping time with respect to the filtration \mathcal{F}_k since $\{\tau \leq s_k\} \in \mathcal{F}_k$. In other words, the event whether the *k*th selection occurs completely depends on the information prior to time s_k (including s_k).

Let $\mathcal{D}^{test,0}$ and $\mathcal{D}^{cal,0}$ be the index sets for subjects in the testing and calibration data that do not belong to class c, respectively. In the final step of our proof, we shall apply the optional stopping theorem to the filtration $\{\mathcal{F}_k\}$. Recall that t_l is lower limit of the threshold, and m^* is the total number of misclassifications in both \mathcal{D}^{cal} and \mathcal{D}^{test} when the threshold is t_l . The group-wise FSR of the FASI algorithm is

$$\begin{aligned} \operatorname{FSR}_{a}^{\{c\}} &= \mathbb{E}\{\operatorname{FSP}_{a}^{\{c\}}(\tau)\} \end{aligned} \qquad (D.25) \\ &\leq \alpha \cdot \mathbb{E}\left[\cdot \mathbb{E}\left\{\frac{|\mathcal{D}^{cal}| + 1}{|\mathcal{D}^{test}|} \frac{V^{t}(\tau)}{V^{c}(\tau) + 1} | \mathcal{F}_{m^{*}}\right\}\right] \\ &= \alpha \cdot \mathbb{E}\left[\frac{|\mathcal{D}^{cal}| + 1}{|\mathcal{D}^{test}|} \cdot \frac{V^{t}(t_{l})}{V^{c}(t_{l}) + 1}\right] \\ &= \alpha \cdot \mathbb{E}\left\{\frac{|\mathcal{D}^{cal}| + 1}{|\mathcal{D}^{test}|} \cdot \frac{|\mathcal{D}^{test,0}|}{|\mathcal{D}^{cal,0}| + 1}\right\} \end{aligned} \qquad (D.26) \\ &\leq \alpha \cdot \mathbb{E}\left\{\frac{p_{c}^{test,0}}{p_{c}^{cal,0}}\right\} \\ &\coloneqq \gamma_{c}\alpha, \end{aligned}$$

where $p_c^{test,0}$ and $p_c^{cal,0}$ are the proportions of individuals that do not belong to class c in the test and calibration data, respectively. To get Equation (D.26) we have used the fact that when t_l is used then all subjects are classified to class c. This completes the proof.

Remark 5. We provide a remark to explain the \mathbb{E} operator in (D.25). As indicated by (D.19), Algorithm 1 is equivalent to a thresholding rule based on \hat{S}_i^c . The data-driven threshold (or stopping time), τ , is a random variable that varies across different realizations or data sets. The FSP, denoted as $FSP_a^{\{c\}}(\tau)$, is a random variable that differs across data sets. The FSR, defined as the expectation of the FSP, integrates the randomness across the observed data $\{(X_i, A_i, Y_i) : i \in \mathcal{D}\}$ and the test data $\{(X_j, A_j, Y_j) : j \in \mathcal{D}^{test}\}$.

D.2 Proof of Theorem 1

The theory can be established by adjusting some steps in the previous proof. The estimated FSP in group a for a given threshold t is:

$$\hat{Q}_c^a(t) = \frac{|\mathcal{D}^{cal}| + |\mathcal{D}^{test}| + 1}{|\mathcal{D}^{cal}| + 1} \cdot \frac{\sum_{i \in \mathcal{D}^{cal}} \mathbb{I}(\hat{S}_i^c \ge t, Y_i \neq c, A_i = a) + 1}{\sum_{i \in \mathcal{D}^{test} \cup \mathcal{D}^{cal}} \mathbb{I}(\hat{S}_i^c \ge t, A_i = a) + 1}.$$

Similar to the previous proof, define

$$\tau = \hat{Q}_c^{a,-1}(\alpha) = \min\left\{t \in \mathbf{S} : \hat{Q}_c^a(t) \le \alpha\right\}.$$

Then following the arguments in Lemma 1, the FASI algorithm is equivalent to the following thresholding rule: $\{\mathbb{I}(\hat{S}_{j}^{c} \leq \tau) : j \in \mathcal{D}^{test}\}.$

Define $V^t(t)$, $R^t(t)$, $V^c(t)$, $R^c(t)$, $W^c(t)$ and $W^c(t)$ as before. The operation of the FASI algorithm implies that

$$\frac{1+R^c(\tau)}{1+R^c(\tau)+R^t(\tau)} \le \frac{|\mathcal{D}^{cal}|+1}{|\mathcal{D}^{cal}|+|\mathcal{D}^{test}|+1}\alpha.$$

It follows that

$$\frac{R^t(\tau) \vee 1}{1 + R^c(\tau)} \ge \frac{1}{\alpha} \left(\frac{|\mathcal{D}^{test}|}{|\mathcal{D}^{cal}| + 1} + 1 - \alpha \right).$$

Some simple algebra provides an upper bound of the following ratio at the stopping time τ :

$$\frac{1+R^c(\tau)}{R^t(\tau)\vee 1} \le \left\{\frac{|\mathcal{D}^{test}|}{|\mathcal{D}^{cal}|+1} + 1 - \alpha\right\}^{-1} \alpha \le \frac{|\mathcal{D}^{cal}|+1}{|\mathcal{D}^{test}|} \alpha.$$
(D.27)

Let $\mathcal{D}^{test,0}$ and $\mathcal{D}^{cal,0}$ denote the index sets for subjects in the testing and calibration data that do not belong to class c, respectively. Conditional on the filtration $\{\mathcal{F}_k\}$, and combining the above results, we have

$$\operatorname{FSR}_{a}^{\{c\},*} \leq \mathbb{E}\left(\mathbb{E}\left[\frac{V^{t}(\tau)}{1+V^{c}(\tau)} \cdot \frac{1+V^{c}(\tau)}{1+R^{c}(\tau)} \cdot \frac{1+R^{c}(\tau)}{R^{t}(\tau) \vee 1} | \mathcal{F}_{m^{*}}\right]\right)$$

$$\leq \mathbb{E}\left(1 \cdot \frac{|\mathcal{D}^{cal}|+1}{|\mathcal{D}^{test}|} \alpha \cdot \mathbb{E}\left[\frac{V^{t}(\tau)}{V^{c}(\tau)+1} | \mathcal{F}_{m^{*}}\right]\right)$$

$$= \alpha \mathbb{E}\left\{\frac{|\mathcal{D}^{cal}|+1}{|\mathcal{D}^{test}|} \cdot \frac{|\mathcal{D}^{test,0}|}{|\mathcal{D}^{cal,0}|+1}\right\}$$

$$\leq \gamma_{c}\alpha,$$

establishing the desired result. Here γ_c is defined as $\gamma_c = \mathbb{E}\left\{\frac{p_c^{test,0}}{p_c^{cal,0}}\right\}$, with $p_c^{test,0}$ and $p_c^{cal,0}$ being the proportions of individuals that do not belong to class c in the test and calibration data, respectively. In the second line of the above equation, we have utilized the trivial fact $V^c(\tau) \leq R^c(\tau)$ and the inequality (D.27). In the third line of the above equation, we have used the fact that $\frac{V^t(t)}{V^c(t)+1}$ is a martingale (as established in Section D.1.2). We apply the optional stopping theorem to the filtration \mathcal{F}_{m^*} . Note that when t_l is used then all subjects are classified to class c. The proof is complete.

E Proof of Theorem 3

The theorem implies that the optimal confidence score for constructing R-values should be $S^c(x,a) = \mathbb{P}(Y = c | X = x, A = a)$. A similar optimality theory has been developed in the context of multiple testing with groups (Cai & Sun 2009). However, the proof for the binary classification setup with the indecision option is much more complicated; we provide the proof here for completeness. We first establish an essential monotonicity property in Section E.1, then prove the optimality theory in Section E.2.

E.1 A monotonicity property

Suppose we use $S_j^c(x, a) = \mathbb{P}(Y_j = c | X_j = x, A_j = a)$, for $j \in \mathcal{D}^{test}$, as the confidence score. The corresponding theoretical R-values can be obtained via the conversion algorithm in Appendix B.5. Under Model B.4, the mFSR level with threshold t is mFSR_a^{c}(t) = $\mathbb{P}(Y \neq c | S^c \geq t, A = a)$. The

theoretical R-values is defined as $R^c(s^c) = \inf_{t \leq s^c} \left\{ \operatorname{mFSR}_a^{\{c\}}(t) \right\}$. Let $Q_a^c(t)$ be the mFSR level when the threshold is t. The next proposition characterizes the monotonic relationship between $Q_a^c(t)$ and t.

Proposition 2. $Q_a^c(t)$ is monotonically decreasing in t.

The proposition is essential for expressing the oracle procedure as a thresholding rule based on S^c . Specifically, denote $Q_a^{c,-1}(\cdot)$ the inverse of $Q_a^c(\cdot)$. The monotonicity of $Q_a^c(t)$ and the definition of the theoretical R-value together imply that $S_j^c(x,a) = Q_a^{c,-1}(R_j^c)$ for $a \in \mathcal{A}$. For notational convenience, let $T_j(x,a) = \mathbb{P}(Y_j = 2|X_j = x, A_j = a)$. Then $S_j^1 = 1 - T_j$ and $S_j^2 = T_j$. Therefore the oracle rule

$$\delta_j^{OR} = \mathbb{I}(R_j^1 \le \alpha_1) + 2\mathbb{I}(R_j^2 \le \alpha_2)$$

can be equivalently written as

$$\delta_{j}^{OR} = \mathbb{I}\left\{S_{j}^{1} \ge Q_{1,a}^{-1}(\alpha_{1})\right\} + 2\mathbb{I}\left\{S_{j}^{2} \ge Q_{2,a}^{-1}(\alpha_{2})\right\}$$
$$= \mathbb{I}\left\{T_{j} \le 1 - Q_{1,a}^{-1}(\alpha_{1})\right\} + 2\mathbb{I}\left\{T_{j} \ge Q_{2,a}^{-1}(\alpha_{2})\right\},$$
(E.28)

for $j \in \mathcal{D}^{test}$. This provides a key technical tool in Section E.2.

Proof of Proposition 2.

Define $\tilde{Q}_a^c(t) = 1 - Q_a^c(t)$. We only need to show that $Q_a^c(t)$ is monotonically increasing in t. Let $\mathcal{M}_a = \{j \in \mathcal{D}^{test} : A_j = a\}$. According to the definition of the mFSR and the definition of S_j^c , we have

$$\mathbb{E}\left\{\sum_{j\in\mathcal{M}_a}\left\{S_j^c - \tilde{Q}_a^c(t)\right\}\mathbb{I}(S_j^c > t)\right\} = 0,\tag{E.29}$$

where the expectation is taken over both \mathcal{D}^{test} . It is important to note that the oracle procedure, which assumes that all distributional information is known, does not utilize \mathcal{D}^{train} and \mathcal{D}^{cal} . It is easy to see from Equation (E.29) that $\tilde{Q}_a^c(t) > t$ otherwise the summation on the LHS must be positive, leading to a contradiction.

Next we show that $t_1 < t_2$ implies $\tilde{Q}_a^c(t_1) \leq \tilde{Q}_a^c(t_2)$. Assume instead that $\tilde{Q}_a^c(t_1) > \tilde{Q}_a^c(t_2)$. We focus on group a, then

$$\begin{split} &\sum_{j \in \mathcal{M}_a} \{S_j^c - \tilde{Q}_a^c(t_1)\} \mathbb{I}(S_j^c > t_1) \\ = &\sum_{j \in \mathcal{M}_a} \{S_j^c - \tilde{Q}_a^c(t_2) + \tilde{Q}_a^c(t_2) - \tilde{Q}_a^c(t_1)\} \mathbb{I}(S_j^c > t_1) \\ = &\sum_{i \in \mathcal{M}_a} \{S_j^c - \tilde{Q}_a^c(t_2)\} \mathbb{I}(S_j^c > t_2) + \sum_{j \in \mathcal{M}_a} \{S_j^c - \tilde{Q}_a^c(t_2)\} \mathbb{I}(t_1 \le S_j^c \le t_2) \\ &+ \sum_{j \in \mathcal{M}_a} \left\{ \tilde{Q}_a^c(t_2) - \tilde{Q}_a^c(t_1) \right\} \mathbb{I}(S_j^c > t_1) \\ = &I + II + III. \end{split}$$

Taking expectations on both sides, it is easy to see that the LHS is zero. However, the RHS is strictly greater than zero. For term I, we have $\mathbb{E}(I) = 0$ according to the definition of mFSR. For term II, we have $\mathbb{E}(II) < 0$ as we always have $\tilde{Q}_a^c(t) > t$. For term III, we have $\mathbb{E}(III) < 0$ since we assume $\tilde{Q}_a^c(t_1) > \tilde{Q}_a^c(t_2)$. It follows that the assumption $\tilde{Q}_a^c(t_1) > \tilde{Q}_a^c(t_2)$ cannot be true, and the proposition is proved.

E.2Proof of the theorem

Define the expected number of true selections $\text{ETS} = \sum_{j \in \mathcal{D}^{test}} \mathbb{I}(Y_j = c, \hat{Y}_j = c)$. Then it can be shown that minimizing the EPI subject to the FSR constraint is equivalent to maximizing the ETS subject to the same constraint.

According to Proposition 2, the oracle rule can be written as

$$\delta_j^{OR} = \mathbb{I}\left\{T_j \le 1 - Q_{1,a}^{-1}(\alpha_1)\right\} + 2\mathbb{I}\left\{T_j \ge Q_{2,a}^{-1}(\alpha_2)\right\}.$$

The mFSR constraints for the oracle rule imply that

$$\mathbb{E}\left\{\sum_{j\in\mathcal{M}_a} (T_j - \alpha_1)\mathbb{I}(\delta_j^{OR} = 1)\right\} = 0, \ \mathbb{E}\left\{\sum_{j\in\mathcal{M}_a} (1 - T_j - \alpha_2)\mathbb{I}(\delta_j^{OR} = 2)\right\} = 0.$$
(E.30)

Let $\boldsymbol{\delta} \in \{0, 1, 2\}^m$ be a general selection rule in $\mathcal{D}_{\alpha_1, \alpha_2}$. Then the mFSR constraints for $\boldsymbol{\delta}$ implies that

$$\mathbb{E}\left\{\sum_{j\in\mathcal{M}_a} (T_j - \alpha_1)\mathbb{I}(\delta_j = 1)\right\} \le 0, \ \mathbb{E}\left\{\sum_{j\in\mathcal{M}_a} (1 - T_j - \alpha_2)\mathbb{I}(\delta_j = 2)\right\} \le 0.$$
(E.31)

The ETS of $\boldsymbol{\delta} = \{\delta_j : j \in \mathcal{D}^{test}\}$ is given by

$$\operatorname{ETS}_{\boldsymbol{\delta}} = \mathbb{E} \left[\sum_{a \in \mathcal{A}} \sum_{j \in \mathcal{M}_a} \{ \mathbb{I}(\delta_j = 1)(1 - T_j) + \mathbb{I}(\delta_j = 2)T_j \} \right]$$
$$= \sum_{a \in \mathcal{A}} \operatorname{ETS}_{\boldsymbol{\delta}}^{1,a} + \operatorname{ETS}_{\boldsymbol{\delta}}^{2,a}.$$

The goal is to show that $\operatorname{ETS}(\boldsymbol{\delta}^{OR}) \geq \operatorname{ETS}(\boldsymbol{\delta})$. We only need to show $\operatorname{ETS}_{\boldsymbol{\delta}^{OR}}^{c,a} \geq \operatorname{ETS}_{\boldsymbol{\delta}}^{c,a}$ for all c and a. We will show $\operatorname{ETS}_{\boldsymbol{\delta}^{OR}}^{1,a} \geq \operatorname{ETS}_{\boldsymbol{\delta}}^{1,a}$ for a given a. The remaining inequalities follow similar arguments. According to (E.30) and (E.31), we have

$$\mathbb{E}\left[\sum_{j\in\mathcal{M}_a} (T_j - \alpha_1) \left\{ \mathbb{I}(\delta_j^{OR} = 1) - \mathbb{I}(\delta_j = 1) \right\} \right] \ge 0.$$
(E.32)

Let $\lambda_{1,a} = (1 - Q_{1,a}^{-1}(\alpha_1) - \alpha_1)/Q_{1,a}^{-1}(\alpha_1)$. It can be shown that $\lambda_{1,a} > 0$. For $i \in \mathcal{M}_a$, we claim that the oracle rule can be equivalently written as

$$\delta_j^{OR} = \mathbb{I}\left\{\frac{T_j - \alpha_1}{1 - T_j} < \lambda_{1,a}\right\}.$$

Using the previous expression and techniques similar to the Neyman-Pearson lemma, we claim that the following result holds for all $j \in \mathcal{M}_a$:

$$\left\{ \mathbb{I}(\delta_j^{OR} = 1) - \mathbb{I}(\delta_j = 1) \right\} \left\{ T_j - \alpha_1 - \lambda_{1,a}(1 - T_j) \right\} \le 0.$$

It follows that

$$\mathbb{E}\left[\sum_{j\in\mathcal{M}_a} \left\{ \mathbb{I}(\delta_j^{OR}=1) - \mathbb{I}(\delta_j=1) \right\} \{T_j - \alpha_1 - \lambda_{1,a}(1-T_j)\} \right] \le 0.$$
(E.33)

According to (E.32) and (E.33), we have

$$\lambda_{OR} \mathbb{E} \sum_{j \in \mathcal{M}_a} (1 - T_j) \left\{ \mathbb{I}(\delta_j^{OR} = 1) - \mathbb{I}(\delta_j = 1) \right\} = \lambda_{OR} \left(\mathrm{ETS}_{\boldsymbol{\delta}^{OR}}^{1,a} - \mathrm{ETS}_{\boldsymbol{\delta}}^{1,a} \right) \ge 0$$

Note that $\lambda_{OR} > 0$, the desired result follows.

\mathbf{F} **Related Fairness Algorithms**

We discuss two closely related works developed based on the sufficiency principle, in order to emphasize the advantages of FASI.

Zeng et al. (2022) presents a group-wise thresholding rule that maximizes the classifier's power subject to the constraints imposed by the sufficiency principle. However, this method does not allow for indecisions, thereby rendering it impossible to control the error rate at user-specified levels. In contrast, Lee et al. (2021) proposes a selective classification procedure that satisfies the sufficiency principle and allows for indecisions. This enables fair decision-making with error rate control. However, the approach by Lee et al. (2021) relies on complex fitting algorithms and imposes stringent assumptions for theoretical development, which lacks reliable theoretical guarantees regarding output reliability in practical scenarios. Moreover, both methods fail to address the issue of inflated decision errors that arise when classifying multiple individuals simultaneously.

We emphasize that the choice of fairness definition should be contextual and informed by the specific automated decision-making scenario. FASI offers several advantages that make it a more practical choice for practitioners. Firstly, in high-stakes scenarios, the proposed selective inference framework with an indecision option effectively handles situations where the consequences of incorrect decisions are significant. This approach provides practitioners with guidance on which observations require further attention, rather than automatically making decisions when the accuracy may not be sufficient. Secondly, when multiple individuals need to be classified simultaneously, it is crucial to employ a suitable error criterion that can aggregate cumulative errors and control for multiplicity. FASI addresses this concern by providing the FSR, which generalizes the powerful and practical FDR criterion in large-scale testing problems. Lastly, in scenarios where complex or blackbox machine learning models are utilized, having a model-free algorithm like FASI becomes essential. This allows for the deployment of userspecified blackbox models while simultaneously ensuring provable validity in controlling the associated risks without imposing strong model assumptions.

	User Specified Error Rate	Finite Sample Theory	No Assumptions on Model Accuracy
Zeng et al. (2022):	No	No	No
Lee et al. (2021):			
Fair Selective Classification	No	No	No
Via Sufficiency			
FASI	Yes	Yes	Yes

Table 1: Comparison of algorithms developed to fulfill the sufficiency principle.

G Additional Numerical Results

G.1Comparing the R^{χ} -value and R-value

In this section, we demonstrate through simulation that the R-value (12) is more stable than the R^{χ} value (A.1) when $|\mathcal{D}^{test}|$ is small. To do this, we will look at the distributions of R-value and R^{χ} -value for a fixed confidence score of s(x, a) = 0.9.



Figure 8: The comparison between the R^{χ} -value and R-value for varying sizes of the test data set. The left column shows the histograms of the R^{χ} -value (orange) and the right column shows the histograms of the R-value (green). The R^{χ} -values and R-values are computed for a fixed confidence score of s(x, a) = 0.9 based on 1,000 randomly generated data sets.

We consider the setting described in Section 4 with $F_{1,M} = F_{1,F} = \mathcal{N}(\boldsymbol{\mu}_1, 2 \cdot \mathbf{I}_3)$ and $F_{2,M} = F_{2,F} = \mathcal{N}(\boldsymbol{\mu}_2, 2 \cdot \mathbf{I}_3)$. We set $\pi_{2|F} = \pi_{2|M} = 0.8$, $\boldsymbol{\mu}_1 = (1, 1, 1)^{\top}$ and $\boldsymbol{\mu}_2 = (2, 2, 2)^{\top}$. The confidence scores are constructed as the oracle class probabilities P(Y = c|X, A).

In Figure 9, we compute 1,000 R^{χ} -values and R-values for a fixed score of s = 0.9 based on randomly generated \mathcal{D}^{cal} and \mathcal{D}^{test} . The size of the calibration set is fixed at $|\mathcal{D}^{cal}| = 1,000$ and the test set has sizes $|\mathcal{D}^{test}| \in \{5, 50, 200\}$. The columns of Figure 9 show the histograms the R^{χ} -values (left) and R-values (right) with \mathcal{D}^{test} increasing from 5 (first row) to 200 (last row).

When $|\mathcal{D}^{test}| = 5$, we notice that the R^{χ} -value has much more variability than the R-value. This is because the denominator of the R^{χ} -value only utilizes 5 observations when computing the total number of selections. By contrast, the R-value uses 1,005 observations since it has access to data from both \mathcal{D}^{cal} and \mathcal{D}^{test} . Moving further down the rows of Figure 9, the advantage of the R-value slowly disappears as $|\mathcal{D}^{test}|$ increases. This causes the variability of both R^{χ} -value and R-value to become almost identical. We conclude from this simulation that the R-value is more desirable in settings where $|\mathcal{D}^{test}|$ is small since it can use more data to decrease its variability. However, while the R^{χ} -value has more variability for small $|\mathcal{D}^{test}|$, this disadvantage can be quickly overcome through the introduction of a reasonably sized test set.

G.2 Numerical investigations of the factor $\gamma_{c,a}$

In Theorem 1, we show that the FASI algorithm can control the FSR at level $\gamma_{c,a}\alpha_c$. This section investigates the deviations of $\gamma_{c,a}$ from 1. For simplicity, we only focus on $\gamma_{1,a}$. The setup of the simulations is identical to that in Section 4.

Figure 8 shows the estimates of $\gamma_{1,a}$ for both the Female (green solid line) and Male (orange dashed line) groups. We vary $\pi_{2|F}$ from 0.15 to 0.85 while fixing $\pi_{2|M} = 0.5$. The y-axis plots the estimate of $\gamma_{1,a}$ averaged over 1,000 independent simulation runs. In both settings, $\gamma_{1,a}$ is nearly 1 across both the Female and Male groups. In the most extreme setting ($\pi_{1|F} = 0.85$), $\gamma_{1,a}$ deviates away from 1 by 0.01.



Figure 9: Estimates of $\gamma_{1,a}$ from the simulations in Section 4. The solid (green) line represents the estimate of $\gamma_{1,F}$ for the Female protected group and similarly the orange (long-dashed) line for the Male protected group.

G.3 FASI deployed with other machine learning models

One of the attractive guarantees of our proposed selective inference framework is that we can have the guarantees of Theorem 1, regardless of the machine learning algorithm that is used to generate the confidence scores. In this section, Figure 10 replicates the results of Simulation 1 in Section 4, for a variety of machine learning models where the data has two protected groups, Female and Male. In this section we use, logistic regression, GAM, Nonparametric Naive Bayes, and XGBoost (James et al. 2023, Hastie et al. 2009, Silverman 1986, Chen & Guestrin 2016) to estimate the confidence scores that will be converted to the R-values for our FASI framework.

The left column of Figure 10 plots the FSR for classification group 2 against a varying proportion of signal $\pi_{2|F}$ from the Female protected group i.e. the true proportion of Females that belong to class 2. The right column shows the corresponding EPI for each ML model. The goal is to control FSR at the 10% level.

As we go down the rows, we notice that every model is able to effectively control the False Selection Rate (similar to Simulation 1), however each model has a different EPI. Here, it seems that Logistic Regression, GAM and Nonparametric Naive Bayes have a similar EPI that gets close to 20% in the most extreme case. However, XGBoost has a slightly higher EPI that gets closer to 30% in the worst case. This is a consequence of the accuracy that each ML model has when estimating the true posterior probability P(Y = 2|X, A) for use in our FASI algorithm. However while some models are more or less accurate than others, they are all able to control the FSR at the desired level.



Figure 10: FSR control for the high risk classification. Left column: The resulting FSR from multiple different ML models that are used to estimate the confidence scores used to calculate the R-value. Right column: The corresponding EPI from different confidence scores. The overall FSR (green / solid) as well as both the Female (blue / dashed) and Male (orange / dot-dashed) protected group FSR's are controlled at the desired 10% level, for all ML algorithms. The x-axis varies the amount of true proportion of high risk observations from the Female protected group, while fixing the true proportion from the male group at 50%.