
RISE: Robust Individualized Decision Learning with Sensitive Variables

Xiaoqing Tan
University of Pittsburgh
xit31@pitt.edu

Zhengling Qi
George Washington University
qizhengling@email.gwu.edu

Christopher W. Seymour
University of Pittsburgh
seymourc@pitt.edu

Lu Tang*
University of Pittsburgh
lutang@pitt.edu

Abstract

This paper introduces RISE, a robust individualized decision learning framework with sensitive variables, where sensitive variables are collectible data and important to the intervention decision, but their inclusion in decision making is prohibited due to reasons such as delayed availability or fairness concerns. A naive baseline is to ignore these sensitive variables in learning decision rules, leading to significant uncertainty and bias. To address this, we propose a decision learning framework to incorporate sensitive variables during *offline* training but not include them in the input of the learned decision rule during model deployment. Specifically, from a causal perspective, the proposed framework intends to improve the worst-case outcomes of individuals caused by sensitive variables that are unavailable at the time of decision. Unlike most existing literature that uses mean-optimal objectives, we propose a robust learning framework by finding a newly defined quantile- or infimum-optimal decision rule. The reliable performance of the proposed method is demonstrated through synthetic experiments and three real-world applications.

1 Introduction

Recently, there has been a widespread interest in developing methodology for individualized decision rules (IDRs) based on observational data. When deriving IDRs, some collectible data are important to the intervention decision, while their inclusion in decision making is prohibited due to reasons such as delayed availability or fairness concerns. For example, sensitive characteristics of subjects regarding their income, sex, race and ethnicity may not be appropriate to be used directly for decision making due to fairness concerns. In the medical field especially for patients in severe life-threatening conditions such as sepsis, timely bedside intervention decisions have to be made before lab measurements are ordered, assayed and returned to the attending physicians. However, due to the delayed availability of lab results, most of the decisions are made with great uncertainty and bias due to partial information at hand. We define *sensitive variables* as variables whose inclusion into decision rules is prohibited. The formal definition of sensitive variables will be given in Section 3.

In this work, we propose RISE (**R**obust **I**ndividualized decision learning with **S**ensitive variables)¹, a robust IDR framework to improve the outcome of individuals when there are informative yet sensitive variables that are either not available or prohibited from using during IDR deployment. To achieve this, we propose to estimate the optimal IDR by optimizing a quantile- or infimum-based objective,

*Corresponding author.

¹Python code is available at <https://github.com/ellenxtan/rise>.

respectively, for continuous or discrete sensitive variables. This optimization problem is then shown to be equivalent to a weighted classification problem where *most existing machine learning classifiers can be readily applied*. Our idea falls along the lines of work that considers algorithmic fairness [15] while extending it to the setting of causal inference [56] in the sense that decisions are driven by causality rather than a general utility function. We show in our empirical analyses that this leads to fairer and safer real-life decisions with little sacrifice of the overall performance.

Assuming that a larger outcome value is preferable, optimal IDRs are traditionally derived through maximizing the mean outcome of the sample population. In this paper, we are interested in a specific yet broadly applicable setting of learning that involves sensitive variables. We consider offline learning where sensitive variables are collected and *can be used in training the IDRs*, but they *cannot be used as input in the resulting IDRs*. This is a setting commonly considered in the fairness and privacy-related literature for classification (e.g., [23]), but not from a causal standpoint. When there exist important variables that are simply left out from training, the estimated IDR will be biased. This bias can be removed if all important variables are used during training, which we will show in Section 3.1 a mean-optimal approach. The optimal action maximizes the mean outcome where the mean is taken over the sensitive variables, conditioning on other variables. This method, however, has no control of the disparity in sensitive variables. Subjects with different sensitive values may report large outcome differences, hence unfairly or unsafely treated. Therefore, objective functions with robustness guarantees for sensitive variables are preferred, since they offer protection to subjects in the lower tail of the outcome distribution with regards to the sensitive values.

For illustration, we consider a toy example with binary actions, $A \in \{-1, 1\}$. We remark that the decision can only be made based on the variable X whereas S is a sensitive variable. The setup is shown in Table 1 and the oracle values under the mean-optimal rule and RISE² are given in Table 2. The detailed setup can be found in Section 4.1 under Example 1. We consider *vulnerable subjects* as those with low outcome values, as highlighted in red in Table 1 (A full definition is given in Section 3.2). For $X \leq 0.5$, the mean-optimal rule would assign action $A = 1$ as it tries to achieve the largest average reward across $S = 0$ and $S = 1$. Recall that S is not available at the time of decision. However, this action results in great harm for subjects with $S = 1$ as they could get the worst expected outcome of 0. On the contrary, RISE improves the worst-case outcome by assigning $A = -1$, protecting the vulnerable subjects. Likewise, for $X > 0.5$, the mean-optimal rule assigns $A = -1$ while RISE assigns $A = 1$ protecting those with $S = 0$ that could have experienced an outcome of 5. Compared to the mean-optimal rule, the proposed rule achieves a larger reward among vulnerable subjects while maintaining a comparable overall reward.

Table 1: Toy example setup of $E(Y|X, S, A)$.

	$X \leq 0.5$		$X > 0.5$	
	$S = 0$	$S = 1$	$S = 0$	$S = 1$
$A = -1$	11	13	5	27
$A = 1$	30	0	15	13

Table 2: Toy example results.

	Average reward	
	Overall	Vulnerable
Mean-optimal rule	15.5	2.5
RISE	13	14

Main contributions. *Methodology-wise*, 1) we propose a novel framework, RISE, to handle sensitive variables in causality-driven decision making. Robustness is introduced to improve the worst-case outcome caused by sensitive variables, and as a result, it reduces the outcome variation across subjects. The latter is directly associated with fairness and safety in decision making. To the best of our knowledge, we are among the first to propose a robust-type fairness criterion under causal inference. 2) We introduce a classification-based optimization framework that can easily leverage most existing classification tools catered to different functional classes, including state-of-the-art random forest, boosting, or neural network models. *Application-wise*, 3) the consideration of sensitive variables in decision learning is important to applications in policy, education, healthcare, etc. Specifically, we illustrate the application of RISE using three real-world examples from fairness and safety perspectives where robust decision rules are needed, across which we have observed robust performance of the proposed approach. From a fairness perspective, we consider a job training program where age is considered as a sensitive variable. From a safety perspective, we consider two applications to healthcare where lab measurements are considered as sensitive variables.

²For the mean-optimal rule, overall average reward is calculated by $(30 + 0 + 5 + 27)/4 = 15.5$, reward among vulnerable subjects is calculated by $(0 + 5)/2 = 2.5$; for RISE, overall average reward is calculated by $(11 + 13 + 15 + 13)/4 = 13$, reward among vulnerable subjects is calculated by $(13 + 15)/2 = 14$.

2 Related Work

Our work focuses on individualized decision rules, which aim at assigning treatment decision based on subject characteristics. Existing methods for deriving IDRs include model-based methods such as Q-learning [75, 39, 38] and A-learning [53, 61], model-free policy search methods [77, 81, 80], and contextual bandit methods [6, 31]. In Appendix A, we provide additional literature review on general IDRs under causal settings. Fairness, safety and robustness are topics of interest that extend well beyond causal inference. In the following, we provide a review of these areas, with focus given to work related to causal inference and IDRs.

Fairness and safety in IDRs. The consideration of fairness and safety in machine learning has seen an explosion of interest in the past few years [15, 69, 4, 43, 20, 12, 37, 47, 65, 66], especially for solving classification and regression problems to help derive decisions that are not only accurate but also fair. In these work, sensitive variables are also referred to as sensitive, protected, or auxiliary attributes. We extend the definition of sensitive variables to include delayed information that is not available at deployment as it is also suitable for this framework.

Among earlier work, preprocessing [23, 17, 13, 58] and inprocess training approaches [5, 20, 29] consider disentangling the input X from a known or unknown sensitive variable S so that the transformed X does not contain any information related to S . Due to the causal nature of IDRs, effect of IDRs cannot be estimated consistently when an informative S is left out and the resulting rule is suboptimal. This follows from the classic argument that any unmeasured confounding (i.e., S), if not accounted for, would lead to bias. Similar issues persist in contextual bandits [22, 46]. [35] considers reducing the impact of auxiliary variables on prediction under distributional shift. Although it is motivated from a causal idea, its main focus is still on prediction. Inside the causal framework, [79, 42] extend fairness from prediction to policy learning using causal graphical models by incorporating fairness constraints. [10] considers counterfactual fairness that seeks to achieve conditional independence of the decisions via data preprocessing. Despite earlier efforts in bringing fairness into the causal framework, most of these approaches only ensure mean zero disparity in S but do not have robustness guarantees in the sense that the variance of the disparity in S is not controlled. Besides, most examples consider a single categorical sensitive variable, but not multiple or continuous ones.

Robustness in IDRs. Recently the statistical literature has witnessed a growing interest in developing robust methods for estimating IDRs. They introduce robustness into the objective function by using quantile-optimal treatment regimes or mean-optimal treatment regimes under certain constraints to improve the gain of individuals at the lower tail of the reward spectrum [72, 74, 51, 49, 16]. In particular, [72, 51] propose to estimate quantile- or tail-optimal treatment regimes. [74] studies the mean-optimal treatment regime under a constraint to control for the average potential risk. [49] proposes a decision-rule-based optimized covariates dependent equivalent for tasks of individualized decision making. [16] considers mean and quantile objectives simultaneously by maximizing the average value with the guarantee that its tail performance exceeds a prespecified threshold. Robustness, in their sense, pertains to the outcome distribution subject to the sampling error. When sensitive variables are present, we consider instead the robustness of the outcome distribution subject to the uncertainty due to sensitive variables, providing a more targeted way of ensuring robustness, which is directly related to fairness and safety. Compared to algorithms based on explicit fairness constraints (for example [76, 78] in classification and [79, 10] in causal inference) that seek to remove the disparity across different values of S , our method reduces the variance of disparity across S . In addition, constraint-based approaches typically require specialized optimization procedures whereas our approach presents an elegant and systematic way for optimization. To our knowledge, we are the first few to consider decision fairness via a robust objective under the causal framework.

3 Robust Decision Learning Framework with Sensitive Variables

3.1 Preliminaries

Notation. We let random variables be represented by upper-case letters, and their realizations be represented by lower-case letters. Suppose there are n independent subjects sampled from a given population. For subject i , let $A_i \in \{-1, 1\}$ denote a binary treatment assignment and Y_i denote the corresponding outcome. Without loss of generality, we assume a larger value of outcome is desirable.

Under the potential outcomes framework [55, 63], let $Y_i(a)$ be the potential outcome had the subject been assigned to $A = a$ for $a = 1$ or -1 . Let $X_i \in \mathbb{X}$ be the feature vector and, for now, S_i be a single sensitive variable. We consider $S \in \mathbb{S}$ where $\mathbb{S} = \{1, \dots, K\}$ if S is discrete and $\mathbb{S} = \mathbb{R}$ if S is continuous. The extension to multiple sensitive variables is presented in Section 3.4.

Definition of sensitive variables. We define sensitive variables S as variables that are important to the intervention decision, but their inclusion in decision making is prohibited. Formally, consider variables X and S that are both available during model training and are both determinants of conditional average treatment effect [54]. While S may be involved in training, the derived decision rule $d(\cdot)$ precludes the input of S due to sensitive concerns. Hence, the derived IDR is a function of the form $d(X) : \mathbb{X} \mapsto \mathbb{A}$. Following the above definition, we consider an offline learning framework where sensitive variables are collected and can be used in obtaining the IDRs, but they cannot be used in the resulting IDRs. A causal diagram and a decision diagram are provided in Figure 1. As it shows in Figure 1a, both X and S confound the effect of treatment A on outcome Y . The arrows represent the causal relationship between variables. Note that X and S can be correlated. This causal diagram is formalized in Assumption 1 below. On the other hand in Figure 1b, in the decision diagram under our setting, S is shown in a dotted circle as S may not be readily available at the time of decision making. We connect S and A with a dotted arrow to indicate that S is incorporated in the training of the decision rule, but it is not required at deployment.




Figure 1: (a) A causal diagram. (b) A decision diagram.

Assumption 1. Assume the following conditions hold:

1a Consistency: $Y = Y(-1)\mathbb{1}(A = -1) + Y(1)\mathbb{1}(A = 1)$.

1b Positivity: $0 < \Pr(A = 1|X, S) < 1$.

1c Unconfoundedness: $\{Y(-1), Y(1)\} \perp A | \{X, S\}$ and $\{Y(-1), Y(1)\} \not\perp A | X$.

Assumption 1a is the standard consistency assumption in causal inference and Assumption 1b states that every subject has a nonzero probability of getting the treatment. Assumption 1c states that given X and S , the potential outcomes are independent of the treatment assignments. Besides, unconfoundedness does not hold when only X is given, signifying the important role of S . Under causal settings, Assumption 1c implies that treatment effects cannot be non-parametrically identifiable without S [44, 54]. Approaches such as ones that disentangle X from S under supervised learning settings mentioned in Section 2 will introduce bias towards estimating IDRs.

Naive approaches that omit sensitive variables. When S is not available for future deployment, a naive approach is to maximize $E_X\{E(Y|X, A = d(X))\}$ over d using (X, A, Y) during the training procedure. This approach introduces bias in the estimation of conditional treatment effect and leads to a suboptimal IDR due to the unmeasured confounder S .

Mean-optimal approaches that use the sensitive variables. It is thus important that one includes S into the training procedure. For example, if we consider the value function framework (i.e., expected outcome) used by most existing works such as [36, 52], we can show that

$$\begin{aligned} E\{Y(d)\} &= E_{X,S}[E(Y(d)|X, S)] = E_X[E_{S|X}\{E(Y(d)|X, S)\}] \\ &= E_X[E_{S|X}\{E(Y|X, S, A = d(X))\}] \neq E_X[E(Y|X, A = d(X))], \end{aligned} \quad (1)$$

where the third equality in (1) holds by Assumption 1 and the last inequality also indicates the naive approaches without using S will in general fail. Then one valid approach is to maximize

$E_X [E_{S|X}\{E(Y|X, S, A = d(X))\}]$ over d using (X, S, A, Y) . The optimal IDR under this criterion is, for every $X \in \mathbb{X}$,

$$\tilde{d}(X) \in \text{sgn}(E_{S|X}\{E(Y|X, S, A = 1)\} - E_{S|X}\{E(Y|X, S, A = -1)\}),$$

which finds the treatment that maximizes the conditional expected outcome given X by averaging out the effect of the sensitive variable S . Mean-optimal approaches, however, fail to control the disparities across realizations of the sensitive variables due to the integration over S , which may lead to unsatisfactory decisions to certain subgroups, as illustrated in the toy example in Section 1.

3.2 Robust Optimality with Sensitive Variables

Driven by the limitation of existing approaches, our goal is to derive a robust decision rule that maximizes the worst-case scenarios of subjects when some sensitive information is not available at the time of deploying the decision rule. Specifically, our robust decision learning framework draws decisions based on individuals' available characteristics summarized in the vector X without the sensitive variable S , while improving the worst-case outcome of subjects in terms of the sensitive variable in the population. Formally, given a collection \mathbb{D} of all treatment decision rules depending only on X , the proposed RISE approach estimates the following IDR, which is defined as

$$d^* \in \arg \max_{d \in \mathbb{D}} E_X [G_{S|X}\{E(Y|X, S, A = d(X))\}], \quad (2)$$

where $G_{S|X}(\cdot)$ could be chosen as some risk measure for evaluating $E(Y|X, S, A = d(X))$ for each $S \in \mathbb{S}$. Examples include variance, conditional value at risk, quantiles, etc. In this paper, we consider $G_{S|X}$ as the conditional quantiles (for a continuous S) or the infimum (for a discrete S) over S .

Specifically, for a discrete S , $G_{S|X}$ is considered as an infimum operator of $E(Y|X, S, A = d(X))$ over S . We thus aim to find

$$d^* \in \arg \max_{d \in \mathbb{D}} E_X [\inf_{s \in \mathbb{S}} \{E(Y|X, S = s, A = d(X))\}],$$

where \inf is the infimum taken with respect to $E(Y|X, s, A = d(X))$ over $s \in \mathbb{S}$. This implies that for a given X , $d^*(X)$ assigns the treatment that yields the best worst-case scenario among all possible values of S for every $X \in \mathbb{X}$, or equivalently,

$$d^*(X) \in \text{sgn}(\inf_{s \in \mathbb{S}} \{E(Y|X, S = s, A = 1)\} - \inf_{s \in \mathbb{S}} \{E(Y|X, S = s, A = -1)\}).$$

For a continuous S , we consider $G_{S|X}\{E(Y|X, S, A = d(X))\}$ as $Q_{S|X}^\tau\{E(Y|X, S, A = d(X))\}$, which is the τ -th quantile of $E(Y|X, S, A = d(X))$ and $\tau \in (0, 1)$ is the quantile level of interest. Specifically, $Q_{S|X}^\tau\{E(Y|X, S, A = d(X))\} = \inf\{t : F(t) \geq \tau\}$ with F denoting the conditional distribution function of $E(Y|X, S, A = d(X))$ over \mathbb{S} given X and d . Note the randomness behind $E(Y|X, S, A = d(X))$ given X and d is fully determined by the sensitive variable S . Then optimal IDR under this criterion is defined as

$$d^* \in \arg \max_{\mathbb{D}} E_X [Q_{S|X}^\tau\{E(Y|X, S, A = d(X))\}].$$

This implies that for a given X , $d^*(X)$ assigns a treatment that yields the largest τ -th quantile of the outcome over the distribution related to S , or equivalently,

$$d^*(X) \in \text{sgn}(\{Q_{S|X}^\tau\{E(Y|X, S, A = 1)\} - Q_{S|X}^\tau\{E(Y|X, S, A = -1)\}).$$

We let $\tau = 0.25$ throughout the paper and suppress τ for simplicity. Results on varying the value of τ is provided in Appendix D; see Section 4.1 for details.

Identifying vulnerable subjects. Our RISE framework provides a natural way to define *vulnerable groups*. Specifically, for a discrete S , if $\inf_S \{E(Y|X, S, A = 1)\} > \inf_S \{E(Y|X, S, A = 0)\}$, then $\arg \inf_S \{E(Y|X, S, A = 0)\}$ is vulnerable given X , otherwise $\arg \inf_S \{E(Y|X, S, A = 1)\}$ is vulnerable. In other words, the vulnerable subjects are those in the worst-off group that needs protection. Similarly, for a continuous S , if $Q_S\{E(Y|X, S, A = 1)\} > Q_S\{E(Y|X, S, A = 0)\}$, then the set $\{S : E(Y|X, S, A = 0) \leq Q_S\{E(Y|X, S, A = 0)\}\}$ defines the vulnerable subjects given X , otherwise this group is defined as $\{S : E(Y|X, S, A = 1) \leq Q_S\{E(Y|X, S, A = 1)\}\}$.

3.3 Estimation and Algorithm

Here we provide a transformation of the proposed RISE from an optimization problem to a weighted classification problem. There are several advantages to this conversion: 1) The optimization problem defined in (2) involves a nonsmooth and nonconvex objective function that could lead to computational challenges. 2) With multiple powerful statistical and machine learning toolbox to choose from, a classification problem can be more readily solved in practice. Hyperparameter tuning and model selection could be conducted to further boost performance. 3) Compared to a direct optimization of (2), a classification-based optimizer allows the use of off-the-shelf software packages that can be tailored to different functional classes or incorporate different properties such as model sparsity.

Proposition 1. *Maximizing the objective function in (2) is equivalent to maximizing*

$$E_X \{ \mathbb{1}(d(X) = 1) [G_{S|X} \{E(Y|X, S, A = 1)\} - G_{S|X} \{E(Y|X, S, A = -1)\}] \}.$$

With Proposition 1 and a proper estimator of the outcome model $E(Y|X, S, A)$ using training data $\mathcal{D}_n = \{X_i, S_i, A_i, Y_i\}_{i=1}^n$, we replace the expectation of Y_i by its estimate \hat{Y}_i and solve the following,

$$\arg \max_{d \in \mathbb{D}} n^{-1} \sum_{i=1}^n [\mathbb{1}(d(x_i) = 1) \{g_1(x_i) - g_2(x_i)\}], \quad (3)$$

where $g_1(x_i) = G_{s|x} \{\hat{Y}_i(x_i, s, a_i = 1)\}$ and $g_2(x_i) = G_{s|x} \{\hat{Y}_i(x_i, s, a_i = -1)\}$. We have the following proposition to address noncontinuity in (3) and transform it into a classification problem. Define \mathbb{F} as a class of all measurable functions over \mathbb{X} .

Proposition 2. *Maximizing the empirical objective in (3) is equivalent to a weighted classification problem of minimizing over $f \in \mathbb{F}$,*

$$n^{-1} \sum_{i=1}^n \mathbb{1}[\text{sgn}\{g_1(x_i) - g_2(x_i)\} \cdot f(x_i) < 0] \cdot |g_1(x_i) - g_2(x_i)| \quad (4)$$

with features x_i , a label $\text{sgn}\{g_1(x_i) - g_2(x_i)\}$ and the sample weight $|g_1(x_i) - g_2(x_i)|$ for $1 \leq i \leq n$.

With Proposition 2, we have transformed the optimization problem (2) into a weighted classification problem (4) where for subject i with features x_i , the true label is $\text{sgn}\{g_1(x_i) - g_2(x_i)\}$ and the sample weight is $|g_1(x_i) - g_2(x_i)|$. The estimated optimal decision rule by (4) is then given by $\hat{d}(x) = \text{sgn}\{\hat{f}(x)\}$ ³. The proof of Proposition 1 and Proposition 2 is presented in Appendix B.

Algorithm 1 provides an algorithmic overview of RISE. The inner expectation $E(Y|X, S, A)$ can be modeled as $\hat{Y}(X, S, A)$ using a twin model separated by the treatment and control groups (i.e., a T-learner as in [27]). For a continuous S , $G(X, A) = Q_{S|X, A} \{E(Y|X, S, A)\}$ is estimated via a quantile regression of \hat{Y} on X but without S . For a discrete S , an estimate of $G(X, A) = \inf_S \{E(Y|X, S, A)\}$ is obtained by finding the minimum among $\{E(Y|X, S = 1, A), \dots, E(Y|X, S = K, A)\}$. The estimated decision rule can then be obtained from the weighted classification. In our implementation, neural networks are used to fit models in the training data sets. The details on modeling and hyperparameter tuning via cross-validations are given in Appendix C. A Python package `rise` based on neural networks is available on GitHub (<https://github.com/ellenxtan/rise>). Note that the model choices are flexible.

Algorithm 1: RISE (Robust individualized decision learning with sensitive variables)

Input: Training data $\mathcal{D}_n = \{Y_i, A_i, X_i, S_i\}_{i=1}^n$

Output: Estimated decision rule \hat{d}

for $i \leftarrow 1$ **to** n **do**

$\hat{Y}_i(x_i, s_i, a_i) \leftarrow$ Model $E(Y|X, S, A = a)$ using \mathcal{D}_n with $a = 1$ and $a = -1$, respectively;

if S is continuous **then**

$g_1(x_i) \leftarrow$ Model $Q_{S|X, A} \{E(Y|X, S, A = a)\}$ via quantile regressions of $\hat{Y}_i(x_i, s_i, a_i)$ on x_i , for \mathcal{D}_n with $a = 1$;

$g_2(x_i) \leftarrow$ Model $Q_{S|X, A} \{E(Y|X, S, A = a)\}$ via quantile regressions of $\hat{Y}_i(x_i, s_i, a_i)$ on x_i , for \mathcal{D}_n with $a = -1$;

else if S is discrete **then**

$g_1(x_i) \leftarrow$ Compute $\inf_{s \in \mathbb{S}} \{\hat{Y}_i(x_i, s, a_i = 1)\}$;

$g_2(x_i) \leftarrow$ Compute $\inf_{s \in \mathbb{S}} \{\hat{Y}_i(x_i, s, a_i = -1)\}$;

end

$\hat{d} \leftarrow$ Build a weighted classification model with features x_i , label $\text{sgn}\{g_1(x_i) - g_2(x_i)\}$, and sample weight $|g_1(x_i) - g_2(x_i)|$ for $1 \leq i \leq n$;

return \hat{d} .

³If $\hat{f}(x) = 0$, assign a random treatment.

3.4 Extension to Multiple Sensitive Variables

The extension from S being a single continuous variable to multiple continuous variables is straightforward in Algorithm 1. For multiple discrete sensitive variables, similar estimation procedure can be conducted as outlined in Section 3.3. Suppose there are L discrete sensitive variables, i.e., $\mathcal{S} = \{S_1, S_2, \dots, S_L\}$. The inner expectation $E(Y|X, S_1, \dots, S_L, A)$ can be obtained with a twin model of Y on X and all \mathcal{S} for each treatment level. The infimum over \mathbb{S} is obtained by finding the minimum iterating space of possible parameter values for each sensitive variable. See Section 4.2 for an example of using multiple discrete sensitive variables. We will discuss in Section 5 the challenges and future work related to the scenario with a mixture of continuous and discrete sensitive variables and the identification of vulnerable subjects under these cases.

4 Numerical Studies

In this section, we perform extensive numerical experiments to investigate the merit of robustness of the proposed framework via simulations and three real-data applications. The results demonstrate that the proposed rules achieve a robust objective with sensitive variables unavailable at the time of decision while maintaining comparable mean outcomes.

Compared approaches. For comparison, we consider the naive and mean-optimal approaches described in Section 3.1, which correspond to different choices of $G(\cdot)$ functions. The naive decision rule that simply disregard information of S , denoted as **Base**, can be formulated in our optimization framework of (2) by letting $G(X, A) = E(Y|X, A)$. The IDR can be estimated directly by fitting a model of Y on X in each treatment arm. The resulting IDR is not sensitive variables-aware and is biased due to confounding, as discussed. Another IDR that resembles traditional mean-optimal decision rules, denoted as **Exp**, can be formulated as $G(X, S, A) = E(Y|X, S, A)$. This can be obtained by training a classification model without S , i.e., only using X , after obtaining an outcome model for the inner expectation $E(Y|X, S, A)$. Note that this approach is not robust to extreme behaviors in S . The modeling approaches described in Appendix C apply to here. We also include the *double robust* [11, 68] versions of Base and Exp, respectively, by adapting Policytree (PT) [64, 3], the latest state-of-the-art policy learning method for maximizing the expected values. The double robust analogues of Base and Exp are termed **PT-Base** and **PT-Exp**, respectively.

Evaluation metrics. 1) *Objective*: the quantile objective is estimated and reported for a continuous S and the infimum objective is for a discrete S . The objective, when $\tau < 0.5$, (here $\tau = 0.25$) represents the value of the “low performers” among all possible value of S under a given d . 2) *Value*: the value function, or expected reward used by the most existing methods, such as [36, 52], is defined as $V(d) = E\{Y(d)\}$. It represents the “average performers”. For randomized trials, an unbiased estimator of $V(d)$ is given by $\hat{V}(d) = \{\sum_{i=1}^T Y_i \mathbb{1}(A_i = \hat{d}(X_i)) / \pi(A_i, X_i)\} / \{\sum_{i=1}^T \mathbb{1}(A_i = \hat{d}(X_i)) / \pi(A_i, X_i)\}$ [41], where T is the sample size of the test data and $\pi(A, X)$ is propensity score. For observational studies, the value is estimated with $\hat{V}(d) = T^{-1} \sum_{i=1}^T \hat{Y}_i(x_i, s_i, a_i = \hat{d})$. We report the metrics among all subjects and among the potential vulnerable subgroup, respectively. For simulation, we consider training data and testing data with sample sizes of 8,000 and 2,000, respectively. For real-data applications, we consider a 80-20 split of the dataset into a training data and a testing data. Continuous covariates are standardized before the estimation. All results are based on 100 replications. Experiments are performed on a 6-core Intel Xeon CPU E5-2620 v3 2.40GHz equipped with 64GB RAM.

4.1 Simulation Studies

Example 1. Here we provide details for the simulation of the motivating example introduced in Section 1. The outcome is generated according to: $Y_i = \mathbb{1}(X_i > 0.5)\{5 + 10\mathbb{1}(A_i = 1) + 22S_i - 24\mathbb{1}(A_i = 1)S_i\} + \mathbb{1}(X_i \leq 0.5)\{11 + 19\mathbb{1}(A_i = 1) + 2S_i - 32\mathbb{1}(A_i = 1)S_i\} + \epsilon_i$, where the covariate $X_i \sim \text{Unif}(0, 1)$, treatment assignment $A_i \sim \text{Bernoulli}(0.5)$, and the noise $\epsilon_i \sim N(0, 1)$. For a discrete type S , $S_i \sim \text{Bernoulli}(0.5)$. For a continuous type S , S_i is generated from a mixture of beta distributions, $\text{Beta}(4, 1)$ and $\text{Beta}(1, 4)$, with equal mixing proportions.

Example 2. We generate Y using the following model: $Y_i = \{0.5 + \mathbb{1}(A_i = 1) + \exp(S_i) - 2.5S_i \mathbb{1}(A_i = 1)\}\{1 + X_{i1} - X_{i2} + X_{i3}^2 + \exp(X_{i4})\} + \{1 + 2\mathbb{1}(A_i = 1) + 0.2 \exp(S_i) - 3.5S_i \mathbb{1}(A_i =$

1)}\{1+5X_{i1}-2X_{i2}+3X_{i3}+2\exp(X_{i4})\}+\epsilon, where $X_{ij} \sim Unif(0, 1)$, $j = 1, \dots, 6$, A satisfies $\log\{P(A_i = 1|X_i)/P(A = 0|X_i)\} = 0.6(-S_i + X_{i1} - X_{i2} + X_{i3} - X_{i4} + X_{i5} - X_{i6})$, and $\epsilon_i \sim N(0, 1)$. For a continuous type S , S_i is generated from a mixture of beta distributions, $Beta(4, 1)$ and $Beta(1, 4)$, with equal mixing proportions; for a discrete type S , we consider a binary S_i that satisfies $\log\{P(S_i = 1|X_i)/P(S_i = 0|X_i)\} = -2.5 + 0.8(X_{i1} + X_{i2} + X_{i3} + X_{i4} + X_{i5} + X_{i6})$.

Table 3 summarizes the performance of the proposed IDRs compared to the mean criterion for Example 1 and Example 2. The proposed RISE achieves the largest objectives and improves the value among vulnerable subjects, while maintaining comparative overall values. As for the objective, intuitively, the proposed rule is expected to achieve a larger objective than all other methods uniformly in \mathbb{X} . We also point out that there is no direct relationship between the objective among all subjects versus the objective among vulnerable subjects. For example, using the toy example with setup in Table 1, and limiting to subjects with $X \leq 0.5$ only, $S = 1$ is vulnerable and is assigned $A = -1$ by the proposed RISE. The objective among $S = 1$ is 13 but the objective among both $S = 0$ and $S = 1$ is $12 = (11 + 13)/2$, which is smaller than that among the vulnerable group. In other words, by protecting the vulnerable subjects, the proposed rule may lead to an increase in the outcome of the vulnerable group, and the gain may result in a higher outcome than the overall mean outcome. PT-Exp tends to show the best improvement in terms of the overall value, as the doubly robust-based estimators tend to reduce variance in value estimation. However, PT-Exp is shown to have minimal benefits for vulnerable subjects. RISE still shows the largest gain in the objective and value among vulnerable subjects among all compared methods.

Table 3: Simulation results for Example 1 and Example 2. Standard error in parenthesis.

Example	Type of S	IDR	Obj. (all)	Obj. (vulnerable)	Value (all)	Value (vulnerable)
1	Disc.	Base	7.03 (0.03)	7.01 (0.04)	14.3 (0.05)	7.92 (0.06)
		Exp	6.39 (0.03)	6.39 (0.04)	14.4 (0.05)	7.14 (0.06)
		PT-Base	2.66 (0.02)	2.65 (0.02)	15.4 (0.05)	2.58 (0.02)
		PT-Exp	2.62 (0.02)	2.62 (0.02)	15.5 (0.05)	2.55 (0.02)
		RISE	12.0 (0.01)	12.0 (0.01)	13.0 (0.01)	14.0 (0.01)
	Cont.	Base	9.12 (0.03)	9.14 (0.04)	14.5 (0.08)	8.25 (0.11)
		Exp	8.75 (0.03)	8.75 (0.04)	14.6 (0.08)	7.58 (0.06)
		PT-Base	6.71 (0.03)	6.72 (0.03)	15.3 (0.05)	4.52 (0.02)
		PT-Exp	6.68 (0.02)	6.67 (0.02)	15.4 (0.05)	4.47 (0.02)
		RISE	12.2 (0.02)	12.2 (0.03)	13.0 (0.01)	13.7 (0.01)
2	Disc.	Base	7.79 (0.02)	8.66 (0.03)	19.4 (0.04)	11.4 (0.06)
		Exp	9.12 (0.03)	10.1 (0.03)	19.5 (0.04)	14.4 (0.05)
		PT-Base	7.19 (0.03)	7.77 (0.03)	19.0 (0.05)	9.71 (0.05)
		PT-Exp	8.30 (0.02)	9.03 (0.03)	19.1 (0.04)	12.2 (0.05)
		RISE	13.5 (0.01)	14.0 (0.01)	17.4 (0.02)	22.1 (0.02)
	Cont.	Base	9.89 (0.02)	9.87 (0.03)	17.6 (0.02)	9.09 (0.04)
		Exp	11.1 (0.02)	11.1 (0.02)	17.8 (0.02)	12.2 (0.04)
		PT-Base	9.30 (0.02)	9.29 (0.03)	18.0 (0.03)	7.61 (0.04)
		PT-Exp	9.41 (0.02)	9.41 (0.02)	18.1 (0.02)	7.92 (0.04)
		RISE	14.1 (0.01)	14.2 (0.02)	17.0 (0.01)	20.3 (0.03)

In the appendix, we consider a continuous S for different quantile criteria $\tau = 0.1$ and 0.5 to test the robustness of RISE. Results show that when τ is small, there is more strength in the proposed method, as the algorithm aims to improve the worst-outcome scenarios. The proposed RISE has the largest gain in objective and value among vulnerable subjects when τ is 0.1, and has similar performance as the compared approaches when τ is 0.5. We also consider a scenario where S is not involved in the data generation of Y , i.e., Assumption 1c is simplified as $\{Y(-1), Y(1)\} \perp A|X$. The estimated objective and value function are similar across all compared approaches, which indicates the robustness of RISE. Finally, we study the performances of our method when Assumption 1b is nearly violated or Assumption 1c is violated. Similar patterns have been observed that the proposed RISE achieves the largest objectives and improves the value among vulnerable subjects, while maintaining comparable overall values. The details can be found in Appendix D.

4.2 Real-data Applications

We present three real-data examples to showcase the robust performance of RISE. These applications consider either fairness or safety in the context of policy [30] and healthcare [19, 59] where sensitive variables commonly exist.

Fairness in a job training program. To illustrate the implication of the proposed method from a fairness perspective, we consider the National Supported Work (NSW) program [30] for improving personalized recommendations of a job training program on increasing incomes. This program intended to provide a 6 to 18-month training for individuals in face of economical and social stress such as former drug addicts and juvenile delinquents. The original experimental dataset consists of 185 individuals who received the job training program ($A = 1$) and 260 individuals who did not ($A = -1$). The baseline covariates are age, years of schooling, race (1 = African Americans or Hispanics, 0 = others), married (1 = yes, 0 = no), high school diploma (1 = yes, 0 = no), earning in 1974, and earning in 1975. The outcome variable is the earning in 1978. In the exploratory analysis using causal forest [71], a random forest-based method for causal inference, we observe that age may play an important role in the causal effect of the job training program on the long-term post-market earning. In this application example we use age as the sensitive variable S and other baseline covariates as X . The earnings in years 1974, 1975, and 1978 are transformed by taking the logarithm of the earning plus one.

Improvement of HIV treatment. To illustrate the implication of the proposed method from a safety perspective when there is delayed information, we consider the ACTG175 dataset among HIV positive patients [19]. The original study considers a total of 2,139 patients who were randomly assigned into four treatment groups. In this data application, we focus on finding the optimal IDRs between two treatments: zidovudine combined with didanosine ($A = -1$) and zidovudine combined with zalcitabine ($A = 1$). The total number of patients receiving these two treatments is 1,046. The baseline covariates we consider are age, weight, CD4 T-cell amount at baseline, hemophilia (1 = yes, 0 = no), homosexual activity (1 = yes, 0 = no), Karnofsky score, history of intravenous drug use (1 = yes, 0 = no), gender (1 = male, 0 = female), CD8 T-cell amount at baseline, race (1 = non-Caucasian, 0 = Caucasian), number of days of previously received antiretroviral therapy, use of zidovudine in the 30 days prior to treatment initiation (1 = yes, 0 = no), and symptomatic indicator (1 = symptomatic, 0 = asymptomatic). The outcome variable is the CD4 T-cell amount at 96 ± 5 weeks from the baseline. We consider CD8 T-cell amount at baseline as the sensitive variable. The response of CD8 T-cell among HIV positive patients has not been fully understood [7]. Clinically, it is plausible that only CD4 is measured in clinical visits where treatments are based on, hence CD8 might not be measured and not used in decision making. As our exploratory analysis using causal forest shows, CD8 T-cell amount may play an important part in the treatment effect of the outcome.

Safe resuscitation for patients with sepsis. For this application, we apply the proposed method to treating sepsis, a life-threatening disease. This application intends to provide an example to apply our method with multiple categorical sensitive variables in the scenario where there is missing yet important information at the time of decision making. We apply the proposed method to a sepsis study from the University of Pittsburgh Medical Center (UPMC). The original study cohort includes 30,687 patients with Sepsis-3 [59] within 6 hours of hospital arrival from 14 UPMC hospitals between 2013 and 2017. For our data analysis, we consider X to be baseline patient characteristics 4 hours before sepsis onset, which includes patient demographics of age, gender (1 = male, 0 = female), race (1 = Caucasian, 0 = others), and weight, and vital signs of usage of mechanical ventilation (1 = yes, 0 = no), respiratory rate, temperature, intravenous fluids (1 = yes, 0 = no), Glasgow Coma Scale score, platelets, blood urea nitrogen, white blood cell counts, glucose, creatinine. We consider two sensitive variables, lactate and Sequential Organ Failure Assessment (SOFA) score 4 hours before sepsis onset. Lactate and SOFA score have been two important indicators of sepsis severity [21, 26, 60, 62]. Different from the baseline patient demographics or common vital signs that are typically obtained at the admission of patients, SOFA score combines performance of several organ systems in the body [59], which requires additional calculation and cannot be obtained directly. Lactate labs measures the level of lactic acid in the blood [2] and are less common in routine examination, which could be delayed in ordering. Hence, their measurements are obtained retrospectively after treatment decisions have been made and are not available at times of decision. We dichotomize lactate level at clinically meaningful value of 2 mmol/L [60], and SOFA score at value of 6 for analysis [70, 18]. The treatment option is whether the patient took any vasopressors during the first 24 hours after sepsis onset. The outcome is hospital survival ($Y = 1$) or death ($Y = 0$). The analysis cohort contains 6,539 patients

in total. We are interested in making decision about whether to treat patients with vasopressors in the first 24 hours after sepsis onset given the measurements of lactate and SOFA are not available at the time of decision making. Additional rationale and background on this example are provided in Appendix E.2.

Table 4: Estimated objective and value of different IDRs for the three data applications. Standard error in parenthesis. The outcome of each study is italicized.

Dataset	IDR	Obj. (all)	Obj. (vulnerable)	Value (all)	Value (vulnerable)
NSW <i>log(income+1)</i>	Base	5.26 (0.04)	5.28 (0.05)	6.32 (0.05)	6.33 (0.07)
	Exp	5.22 (0.04)	5.24 (0.05)	6.37 (0.05)	6.37 (0.07)
	PT-Base	4.97 (0.04)	5.08 (0.06)	6.40 (0.03)	6.38 (0.05)
	PT-Exp	5.03 (0.04)	5.11 (0.05)	6.43 (0.03)	6.40 (0.05)
	RISE	5.43 (0.04)	5.44 (0.04)	6.42 (0.04)	6.42 (0.06)
ACTG175 <i>CD4 T-cell amount</i>	Base	336.9 (1.65)	338.1 (2.23)	353.5 (1.86)	357.5 (2.24)
	Exp	337.5 (1.65)	338.9 (1.80)	355.9 (1.95)	359.1 (2.21)
	PT-Base	299.7 (1.01)	299.5 (1.91)	356.9 (1.72)	350.7 (2.54)
	PT-Exp	300.1 (0.99)	299.9 (1.83)	357.1 (1.55)	352.7 (2.61)
	RISE	351.5 (1.67)	351.2 (1.80)	351.8 (1.88)	363.1 (2.19)
Sepsis <i>survival rate</i>	Base	0.803 (0.001)	0.822 (0.001)	0.965 (0.001)	0.905 (0.002)
	Exp	0.803 (0.001)	0.822 (0.002)	0.966 (0.001)	0.908 (0.002)
	PT-Base	0.758 (0.001)	0.771 (0.002)	0.981 (0.001)	0.848 (0.003)
	PT-Exp	0.758 (0.001)	0.772 (0.002)	0.984 (0.001)	0.875 (0.003)
	RISE	0.836 (0.001)	0.833 (0.001)	0.972 (0.001)	0.923 (0.002)

Results. Table 4 presents the performance of various IDRs on the three applications. As expected, RISE has the largest objective as well as value among vulnerable subjects. The patterns are similar to that in Section 4.1. We apply the Shapley additive explanations (SHAP) approach [33] to help visualize and interpret the important covariates in the decision rules given by RISE and Exp, respectively, in Appendix E.3. The SHAP approach provides unified values to describe the correlation between each feature and the predicted decision rule, respectively [33]. Overall, the direction of correlations is similar for RISE and Exp, but their ranking of feature importance may be different.

5 Discussion

We have proposed RISE, a robust decision learning framework with a novel quantile- or infimum-optimal treatment objective intended to improve the worst-case scenarios of individuals when decisions with uncertainty need to be made, but with sensitive yet important information missing. Our approach can be applied to a broad range of applications, including but not limited to policy, education, healthcare, etc. For a mixture of continuous and discrete sensitive variables, the estimated rule can be obtained by first taking the infimum over the discrete ones as in Section 3.4, then obtaining the quantile over the continuous ones. However, challenges remain in finding the vulnerable subjects described in Section 3.2 under these settings as it may be computationally difficult to find a vulnerable set of S when it is multi-dimensional. Another future work includes the extension of the current binary treatment option to a multi-treatment option. It is also worth mentioning that our work can be naturally extended to the scenario where there exist unmeasured confounders. As long as the conditional average outcome given observed covariates can be identified (via instrumental variables such as [73] or negative control variables such as [50]), our method can be applied.

Acknowledgements

This research was supported in part by the Competitive Medical Research Fund of the UPMC Health System and the University of Pittsburgh Center for Research Computing through the resources provided. The authors thank Jason N. Kennedy for the sepsis data preparation, and Emily B. Brant, Chung-Chou H. Chang and Gong Tang for helpful discussions and feedback.

References

- [1] Martín Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, et al. {TensorFlow}: A system for {Large-Scale} machine learning. In *12th USENIX Symposium on Operating Systems Design and Implementation (OSDI 16)*, pages 265–283, 2016.
- [2] Lars W. Andersen, Julie Mackenhauer, Jonathan C. Roberts, Katherine M. Berg, Michael N. Cocchi, and Michael W. Donnino. Etiology and therapeutic approach to elevated lactate levels. *Mayo Clinic Proceedings*, 88(10):1127–1140, 2013.
- [3] Susan Athey and Stefan Wager. Policy learning with observational data. *Econometrica*, 89(1):133–161, 2021.
- [4] Solon Barocas, Moritz Hardt, and Arvind Narayanan. *Fairness and Machine Learning*. fairml-book.org, 2019. <http://www.fairmlbook.org>.
- [5] Alex Beutel, Jilin Chen, Zhe Zhao, and Ed H Chi. Data decisions and theoretical implications when adversarially learning fair representations. *arXiv preprint arXiv:1707.00075*, 2017.
- [6] Alberto Bietti, Alekh Agarwal, and John Langford. A contextual bandit bake-off. *Journal of Machine Learning Research*, 22(133):1–49, 2021.
- [7] Sushma Boppana and Paul Goepfert. Understanding the CD8 T-cell response in natural HIV control. *F1000Research*, 7, 2018.
- [8] Bibhas Chakraborty and Erica E Moodie. *Statistical Methods for Dynamic Treatment Regimes*. Springer, 2013.
- [9] Bibhas Chakraborty, Susan Murphy, and Victor Strehler. Inference for non-regular parameters in optimal dynamic treatment regimes. *Statistical Methods in Medical Research*, 19(3):317–343, 2010.
- [10] Haoyu Chen, Wenbin Lu, Rui Song, and Pulak Ghosh. On learning and testing of counterfactual fairness through data preprocessing. *arXiv preprint arXiv:2202.12440*, 2022.
- [11] Victor Chernozhukov, Denis Chetverikov, Mert Demirer, Esther Duflo, Christian Hansen, Whitney Newey, and James Robins. Double/debiased machine learning for treatment and structural parameters. *The Econometrics Journal*, 21(1):C1–C68, 01 2018.
- [12] Alexandra Chouldechova and Aaron Roth. A snapshot of the frontiers of fairness in machine learning. *Communications of the ACM*, 63(5):82–89, 2020.
- [13] Elliot Creager, David Madras, Jörn-Henrik Jacobsen, Marissa Weis, Kevin Swersky, Toniann Pitassi, and Richard Zemel. Flexibly fair representation learning by disentanglement. In *International Conference on Machine Learning*, pages 1436–1445. PMLR, 2019.
- [14] Jiewen Du, Xin Yan, Zhihong Liu, Lu Cui, Peng Ding, Xiaoqing Tan, Xiuming Li, Huihao Zhou, Qiong Gu, and Jun Xu. cBinderDB: a covalent binding agent database. *Bioinformatics*, 33(8):1258–1260, 2017.
- [15] Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard Zemel. Fairness through awareness. In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference*, pages 214–226, 2012.
- [16] Ethan X. Fang, Zhaoran Wang, and Lan Wang. Fairness-oriented learning for optimal individualized treatment rules. *Journal of the American Statistical Association*, 0(0):1–14, 2022.
- [17] Michael Feldman, Sorelle A Friedler, John Moeller, Carlos Scheidegger, and Suresh Venkatasubramanian. Certifying and removing disparate impact. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 259–268, 2015.

- [18] Flavio Lopes Ferreira, Daliana Peres Bota, Annette Bross, Christian Mélot, and Jean-Louis Vincent. Serial evaluation of the sofa score to predict outcome in critically ill patients. *The Journal of the American Medical Association*, 286(14):1754–1758, 2001.
- [19] Scott M Hammer, David A Katzenstein, Michael D Hughes, Holly Gundacker, Robert T Schooley, Richard H Haubrich, W Keith Henry, Michael M Lederman, John P Phair, Manette Niu, et al. A trial comparing nucleoside monotherapy with combination therapy in hiv-infected adults with cd4 cell counts from 200 to 500 per cubic millimeter. *New England Journal of Medicine*, 335(15):1081–1090, 1996.
- [20] Tatsunori Hashimoto, Megha Srivastava, Hongseok Namkoong, and Percy Liang. Fairness without demographics in repeated loss minimization. In *International Conference on Machine Learning*, pages 1929–1938. PMLR, 2018.
- [21] Michael D Howell, Michael Donnino, Peter Clardy, Daniel Talmor, and Nathan I Shapiro. Occult hypoperfusion and mortality in patients with suspected infection. *Intensive Care Medicine*, 33(11):1892–1899, 2007.
- [22] Matthew Joseph, Michael Kearns, Jamie H Morgenstern, and Aaron Roth. Fairness in learning: Classic and contextual bandits. *Advances in Neural Information Processing Systems*, 29, 2016.
- [23] Faisal Kamiran and Toon Calders. Data preprocessing techniques for classification without discrimination. *Knowledge and Information Systems*, 33(1):1–33, 2012.
- [24] Michael R Kosorok and Erica EM Moodie. *Adaptive treatment strategies in practice: planning trials and analyzing data for personalized medicine*. SIAM, 2015.
- [25] Filippos Koutroumpakis, Anna Evans Phillips, Dhiraj Yadav, Jorge D Machicado, Maaz Ahsan, Claudia Ramos Rivers, Xiaoqing Tan, Marc Schwartz, Siobhan Proksell, Elyse Johnston, et al. Serum igg4 subclass deficiency defines a distinct, commonly encountered, severe inflammatory bowel disease subtype. *Inflammatory Bowel Diseases*, 27(6):855–863, 2021.
- [26] Uma Krishna, Suresh P Joshi, and Mukesh Modh. An evaluation of serial blood lactate measurement as an early predictor of shock and its outcome in patients of trauma or sepsis. *Indian Journal of Critical Care Medicine*, 13(2):66, 2009.
- [27] Sören R Künnel, Jasjeet S Sekhon, Peter J Bickel, and Bin Yu. Metalearners for estimating heterogeneous treatment effects using machine learning. *Proceedings of the National Academy of Sciences*, 116(10):4156–4165, 2019.
- [28] Eric B Laber, Daniel J Lizotte, Min Qian, William E Pelham, and Susan A Murphy. Dynamic treatment regimes: Technical challenges and applications. *Electronic Journal of Statistics*, 8(1):1225–1272, 2014.
- [29] Preethi Lahoti, Alex Beutel, Jilin Chen, Kang Lee, Flavien Prost, Nithum Thain, Xuezhi Wang, and Ed Chi. Fairness without demographics through adversarially reweighted learning. *Advances in Neural Information Processing Systems*, 33:728–740, 2020.
- [30] Robert J LaLonde. Evaluating the econometric evaluations of training programs with experimental data. *The American Economic Review*, pages 604–620, 1986.
- [31] Lihong Li, Wei Chu, John Langford, and Xuanhui Wang. Unbiased offline evaluation of contextual-bandit-based news article recommendation algorithms. In *Proceedings of the fourth ACM international conference on Web search and data mining*, pages 297–306, 2011.
- [32] Gang Liu, Xiaoqing Tan, Chao Dang, Shuangquan Tan, Shihui Xing, Nianwei Huang, Kangqiang Peng, Chuanmiao Xie, Xiaoying Tang, and Jinsheng Zeng. Regional shape abnormalities in thalamus and verbal memory impairment after subcortical infarction. *Neurorehabilitation and Neural Repair*, 33(6):476–485, 2019.
- [33] Scott M Lundberg and Su-In Lee. A unified approach to interpreting model predictions. *Advances in Neural Information Processing Systems*, 30, 2017.

- [34] Jorge D Machicado, Amir Gougol, Xiaoqing Tan, Xiaotian Gao, Pedram Paragomi, Ioannis Pothoulakis, Rupjyoti Talukdar, Rakesh Kochhar, Mahesh K Goenka, Aiste Gulla, et al. Mortality in acute pancreatitis with persistent organ failure is determined by the number, type, and sequence of organ systems affected. *United European Gastroenterology Journal*, 9(2):139–149, 2021.
- [35] Maggie Makar, Ben Packer, Dan Moldovan, Davis Blalock, Yoni Halpern, and Alexander D’Amour. Causally motivated shortcut removal using auxiliary labels. In *International Conference on Artificial Intelligence and Statistics*, pages 739–766. PMLR, 2022.
- [36] Charles F Manski. Statistical treatment rules for heterogeneous populations. *Econometrica*, 72(4):1221–1246, 2004.
- [37] Ninareh Mehrabi, Fred Morstatter, Nripsuta Saxena, Kristina Lerman, and Aram Galstyan. A survey on bias and fairness in machine learning. *ACM Computing Surveys (CSUR)*, 54(6):1–35, 2021.
- [38] Erica EM Moodie, Thomas S Richardson, and David A Stephens. Demystifying optimal dynamic treatment regimes. *Biometrics*, 63(2):447–455, 2007.
- [39] Susan A Murphy. Optimal dynamic treatment regimes. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 65(2):331–355, 2003.
- [40] Susan A Murphy. An experimental design for the development of adaptive treatment strategies. *Statistics in Medicine*, 24(10):1455–1481, 2005.
- [41] Susan A Murphy, Mark J van der Laan, James M Robins, and Conduct Problems Prevention Research Group. Marginal mean models for dynamic regimes. *Journal of the American Statistical Association*, 96(456):1410–1423, 2001.
- [42] Razieh Nabi, Daniel Malinsky, and Ilya Shpitser. Learning optimal fair policies. In *International Conference on Machine Learning*, pages 4674–4682. PMLR, 2019.
- [43] Razieh Nabi and Ilya Shpitser. Fair inference on outcomes. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018.
- [44] Jerzy Neyman. Sur les applications de la thar des probabilités aux expériences agricoles: Essay des principe. Excerpts reprinted (1990) in English. *Statistical Science*, 5:463–472, 1923.
- [45] CA Onyemekwu, X Tan, J Gamboa, K Moghbeli, K Potter, N Prendergast, P Tiberio, F Billings, EW Ely, T Hulgán, et al. Associations between plasma cell-free mitochondrial dna, delirium and coma during sepsis. In *C15.” SURVIVING SEPSIS” IN THE POST PANDEMIC ERA*, pages A3700–A3700. American Thoracic Society, 2022.
- [46] Vishakha Patil, Ganesh Ghalme, Vineet Nair, and Yadati Narahari. Achieving fairness in the stochastic multi-armed bandit problem. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 5379–5386, 2020.
- [47] Dana Pessach and Erez Shmueli. A review on fairness in machine learning. *ACM Computing Surveys (CSUR)*, 55(3):1–44, 2022.
- [48] Krishnapriya Marangattu Prathapan, Claudia Ramos Rivers, Alyce Anderson, Filippos Koutroumpakis, Ioannis E Koutroubakis, Dmitriy Babichenko, Xiaoqing Tan, Gong Tang, Marc Schwartz, Siobhan Proksell, et al. Peripheral blood eosinophilia and long-term severity in pediatric-onset inflammatory bowel disease. *Inflammatory Bowel Diseases*, 26(12):1890–1900, 2020.
- [49] Zhengling Qi, Ying Cui, Yufeng Liu, and Jong-Shi Pang. Estimation of individualized decision rules based on an optimized covariate-dependent equivalent of random outcomes. *SIAM Journal on Optimization*, 29(3):2337–2362, 2019.
- [50] Zhengling Qi, Rui Miao, and Xiaoke Zhang. Proximal learning for individualized treatment regimes under unmeasured confounding. *arXiv preprint arXiv:2105.01187*, 2021.

- [51] Zhengling Qi, Jong-Shi Pang, and Yufeng Liu. On robustness of individualized decision rules. *Journal of the American Statistical Association*, pages 1–15, 2022.
- [52] Min Qian and Susan A Murphy. Performance guarantees for individualized treatment rules. *Annals of Statistics*, 39(2):1180, 2011.
- [53] James M Robins, Miguel Angel Hernan, and Babette Brumback. Marginal structural models and causal inference in epidemiology. *Epidemiology*, 11(5):550–560, 2000.
- [54] Donald B Rubin. Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of Educational Psychology*, 66(5):688, 1974.
- [55] Donald B Rubin. Bayesian inference for causal effects: The role of randomization. *The Annals of Statistics*, pages 34–58, 1978.
- [56] Donald B Rubin. Causal inference using potential outcomes: Design, modeling, decisions. *Journal of the American Statistical Association*, 100(469):322–331, 2005.
- [57] Yasser Sakr, Ulrich Jaschinski, Xavier Wittebole, Tamas Szakmany, Jeffrey Lipman, Silvio A Namendys-Silva, Ignacio Martin-Loeches, Marc Leone, Mary-Nicoleta Lupu, Jean-Louis Vincent, and ICON Investigators. Sepsis in intensive care unit patients: worldwide data from the intensive care over nations audit. *Open Forum Infectious Diseases*, 5(12):ofy313, 2018.
- [58] Prasanna Sattigeri, Samuel C Hoffman, Vijil Chenthamarakshan, and Kush R Varshney. Fairness gan: Generating datasets with fairness properties using a generative adversarial network. *IBM Journal of Research and Development*, 63(4/5):3–1, 2019.
- [59] Christopher W Seymour, Vincent X Liu, Theodore J Iwashyna, Frank M Brunkhorst, Thomas D Rea, André Scherag, Gordon Rubenfeld, Jeremy M Kahn, Manu Shankar-Hari, Mervyn Singer, et al. Assessment of clinical criteria for sepsis: for the third international consensus definitions for sepsis and septic shock (sepsis-3). *The Journal of the American Medical Association*, 315(8):762–774, 2016.
- [60] Manu Shankar-Hari, Gary S Phillips, Mitchell L Levy, Christopher W Seymour, Vincent X Liu, Clifford S Deutschman, Derek C Angus, Gordon D Rubenfeld, Mervyn Singer, et al. Developing a new definition and assessing new clinical criteria for septic shock: for the third international consensus definitions for sepsis and septic shock (sepsis-3). *The Journal of the American Medical Association*, 315(8):775–787, 2016.
- [61] Chengchun Shi, Alin Fan, Rui Song, and Wenbin Lu. High-dimensional a-learning for optimal dynamic treatment regimes. *Annals of statistics*, 46(3):925, 2018.
- [62] Vikesh K Singh, David C Whitcomb, Peter A Banks, Samer AlKaade, Michelle A Anderson, Stephen T Amann, Randall E Brand, Darwin L Conwell, Gregory A Cote, Timothy B Gardner, et al. Acute pancreatitis precedes chronic pancreatitis in the majority of patients: Results from the naps2 consortium. *Pancreatology*, 22(8):1091–1098, 2022.
- [63] Jerzy Splawa-Neyman, Dorota M Dabrowska, and TP Speed. On the application of probability theory to agricultural experiments. essay on principles. section 9. *Statistical Science*, pages 465–472, 1990.
- [64] Erik Sverdrup, Ayush Kanodia, Zhengyuan Zhou, Susan Athey, and Stefan Wager. policytree: Policy learning via doubly robust empirical welfare maximization over trees. *Journal of Open Source Software*, 5(50):2232, 2020.
- [65] Xiaoqing Tan, Judah Abberbock, Priya Rastogi, and Gong Tang. Identifying principal stratum causal effects conditional on a post-treatment intermediate response. In *Conference on Causal Learning and Reasoning*, pages 734–753. PMLR, 2022.
- [66] Xiaoqing Tan, Chung-Chou H. Chang, Ling Zhou, and Lu Tang. A tree-based model averaging approach for personalized treatment effect estimation from heterogeneous data sources. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvari, Gang Niu, and Sivan Sabato, editors, *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 21013–21036. PMLR, 17–23 Jul 2022.

- [67] Xiaoqing Tan, Christopher A Ross, Michael Miller, and Xiaoying Tang. Change-point analysis of putamen and thalamus subregions in premanifest huntington’s disease. In *2018 IEEE 15th International Symposium on Biomedical Imaging (ISBI 2018)*, pages 531–535. IEEE, 2018.
- [68] Xiaoqing Tan, Shu Yang, Wenyu Ye, Douglas E Faries, Ilya Lipkovich, and Zbigniew Kadziola. When doubly robust methods meet machine learning for estimating treatment effects from real-world data: A comparative study. *arXiv preprint arXiv:2204.10969*, 2022.
- [69] Kush R Varshney. Engineering safety in machine learning. In *2016 Information Theory and Applications Workshop (ITA)*, pages 1–5. IEEE, 2016.
- [70] J-L Vincent, Rui Moreno, Jukka Takala, Sheila Willatts, Arnaldo De Mendonça, Hajo Bruining, CK Reinhart, Peter M Suter, and Lambertus G Thijs. The SOFA (sepsis-related organ failure assessment) score to describe organ dysfunction/failure. On behalf of the working group on sepsis-related problems of the european society of intensive care medicine. *Intensive Care Medicine*, 22(7):707–710, 1996.
- [71] Stefan Wager and Susan Athey. Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association*, 113(523):1228–1242, 2018.
- [72] Lan Wang, Yu Zhou, Rui Song, and Ben Sherwood. Quantile-optimal treatment regimes. *Journal of the American Statistical Association*, 113(523):1243–1254, 2018.
- [73] Linbo Wang and Eric Tchetgen Tchetgen. Bounded, efficient and multiply robust estimation of average treatment effects using instrumental variables. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 80(3):531–550, 2018.
- [74] Yuanjia Wang, Haoda Fu, and Donglin Zeng. Learning optimal personalized treatment rules in consideration of benefit and risk: with an application to treating type 2 diabetes patients with insulin therapies. *Journal of the American Statistical Association*, 113(521):1–13, 2018.
- [75] Christopher JCH Watkins and Peter Dayan. Q-learning. *Machine Learning*, 8(3):279–292, 1992.
- [76] Muhammad Bilal Zafar, Isabel Valera, Manuel Gomez Rogriguez, and Krishna P Gummadi. Fairness constraints: Mechanisms for fair classification. In *Artificial Intelligence and Statistics*, pages 962–970. PMLR, 2017.
- [77] Baqun Zhang, Anastasios A Tsiatis, Eric B Laber, and Marie Davidian. A robust method for estimating optimal treatment regimes. *Biometrics*, 68(4):1010–1018, 2012.
- [78] Brian Hu Zhang, Blake Lemoine, and Margaret Mitchell. Mitigating unwanted biases with adversarial learning. In *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society*, pages 335–340, 2018.
- [79] Junzhe Zhang and Elias Bareinboim. Fairness in decision-making—the causal explanation formula. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018.
- [80] Ying-Qi Zhao, Donglin Zeng, Eric B Laber, Rui Song, Ming Yuan, and Michael Rene Kosorok. Doubly robust learning for estimating individualized treatment with censored data. *Biometrika*, 102(1):151–168, 2015.
- [81] Yingqi Zhao, Donglin Zeng, A John Rush, and Michael R Kosorok. Estimating individualized treatment rules using outcome weighted learning. *Journal of the American Statistical Association*, 107(499):1106–1118, 2012.

A Additional Literature Review

Typical model-based methods for deriving IDRs include Q-learning such as [75, 39, 38] and A-learning such as [53, 40] where a model of responses is imposed and the optimal decision rule is obtained by optimizing value function derived from the model. Model-based methods posit a model of responses given observed covariates and treatment assignments, and obtain the optimal IDR by optimizing the corresponding value function derived from the model. Q-learning optimizes the corresponding value function derived from a parametric model of responses given observed covariates and treatment assignments, and it results in an optimal decision rule. A-learning is a semiparametric method, which derives from a model that directly describes the difference between treatments, with the baseline remaining unspecified. On the other hand, model-free methods such as [77, 81, 80] assign values to actions simply through trial and error without pre-specifying a model. Besides, contextual bandit methods (see [6] and references therein) test out different actions and automatically learn which one has the most rewarding outcome for a given situation. See [9, 8, 28, 24] and references therein for a comprehensive review on general IDRs under causal settings.

B Proof of Propositions

Proof of Proposition 1. We observe that to maximize the objective function in (2) is equivalent to maximizing

$$\begin{aligned}
& E_X [G_{S|X} \{E(Y|X, S, A = d(X))\} | X] \\
&= E_X [G_{S|X} \{E(Y|X, S, A = 1)\} \mathbb{1}(d(X) = 1) \\
&\quad + G_{S|X} \{E(Y|X, S, A = -1)\} \mathbb{1}(d(X) = -1)] \\
&= E_X \{ \mathbb{1}(d(X) = 1) [G_{S|X} \{E(Y|X, S, A = 1)\} - G_{S|X} \{E(Y|X, S, A = -1)\}] \\
&\quad + G_{S|X} \{E(Y|X, S, A = -1)\} \} \\
&\propto E_X \{ \mathbb{1}(d(X) = 1) [G_{S|X} \{E(Y|X, S, A = 1)\} - G_{S|X} \{E(Y|X, S, A = -1)\}] \}.
\end{aligned}$$

□

Proof of Proposition 2. Let $d(x) = \text{sgn}\{f(x)\}$, by this transformation, we consider the following objective on a smooth function $f(x)$,

$$\begin{aligned}
& \arg \max_{d \in \mathbb{D}} \frac{1}{n} \sum_{i=1}^n \{ \mathbb{1}(d(x_i) = 1) [g_1(x_i) - g_2(x_i)] \} \\
&= \arg \max_f \frac{1}{n} \sum_{i=1}^n \mathbb{1}[\text{sgn}\{f(x_i)\} = 1] \cdot [g_1(x_i) - g_2(x_i)] \\
&= \arg \min_f \frac{1}{n} \sum_{i=1}^n \mathbb{1}\{1 \cdot f(x_i) < 0\} \cdot [g_1(x_i) - g_2(x_i)] \\
&= \arg \min_f \frac{1}{n} \sum_{i=1}^n \mathbb{1}[\text{sgn}\{g_1(x_i) - g_2(x_i)\} \cdot f(x_i) < 0] \cdot |g_1(x_i) - g_2(x_i)|.
\end{aligned}$$

The sign of the estimated f above is a solution of d to (3).

Hence, the proposed classification-based objective is to minimize

$$\frac{1}{n} \sum_{i=1}^n \mathbb{1}[\text{sgn}\{g_1(x_i) - g_2(x_i)\} \cdot f(x_i) < 0] \cdot |g_1(x_i) - g_2(x_i)|.$$

To this point, we have transformed the optimization problem (2) into a weighted classification problem where for subject i with features x_i , the true label is $\text{sgn}\{g_1(x_i) - g_2(x_i)\}$ and the sample weight is $|g_1(x_i) - g_2(x_i)|$. □

C Details on Modeling and Hyperparameter Tuning

In our implementation, neural networks with mean or quantile losses are used to fit the models with hyperparameters tuned via a 5-fold cross validation in the training data sets. Specifically, implemented in TensorFlow [1], neural networks with mean squared loss is used to model $E(Y|X, S, A)$ separated by the control arm and the treatment arm, respectively. For continuous S , to model $Q_{S|X,A}\{E(Y|X, S, A)\}$, neural networks with quantile loss is used with a prespecified τ , for the control arm and the treatment arm, respectively. In the final weighted classification model, neural

networks with cross-entropy loss is used. Note that the model choices here are flexible. One can perform model selection if they would like to.

Hyperparameter tuning helps prevent overfitting and is essential in machine learning methods or other black-box algorithms such as neural networks. In our implementation, the optimal hyperparameters are obtained via a 5-fold cross validation in the training data sets. Specifically, we consider the number of hidden layers (1, 2, and 3 layers), the number of hidden units in each layer (256, 512, and 1024 nodes), activation function (RELU, Sigmoid, and Tanh), optimizer (Adam, Nadam, and Adadelata), dropout rate (0.1, 0.2, and 0.3), number of epochs (50, 100, and 200), and batch size (32, 64, and 128).

D Additional Simulations

Different quantile criteria. For the quantile criteria, we also consider $\tau = 0.1$ and 0.5 , respectively. Table 5 presents the simulation results for Example 2 with continuous S using 0.1 quantile criterion and 0.5 quantile criterion, respectively. Results show that when τ is small, there is more strength in the proposed method, as the algorithm aims to improve the worst-outcome scenarios. The proposed RISE has the largest gain in objective and value among vulnerable subjects when τ is 0.1 , and has similar performance as the compared approaches when τ is 0.5 .

S as a noise variable. We generate the outcome Y using the following model where S is not involved: $Y = \mathbb{1}(X_1 \leq 0.5)\{8 + 12\mathbb{1}(A = 1) + 16 \exp(X_2) - 26\mathbb{1}(A = 1)X_2\} + \mathbb{1}(X_1 > 0.5)\{13 + 3\mathbb{1}(A_i = 1) + 2 \exp(X_2) - 8\mathbb{1}(A = 1)X_2\} + \epsilon$, where $X_j \sim Unif(0, 1)$, $j = 1, 2$, $A \sim Bernoulli(0.5)$, and $\epsilon \sim N(0, 1)$. For continuous S , $S = \text{expit}\{-2.5(1 - X_1 - X_2)\}$; for discrete S , we consider a binary S that satisfies $\log\{P(S = 1|X)/P(S = 0|X)\} = -2.5(1 - X_1 - X_2)$. Table 6 summarizes the performance of the proposed IDRs compared to the mean criterion for Example 2. The estimated objective and value function are similar for the compared IDRs, which indicates the robustness of the proposed RISE.

Violations of causal assumptions. To further test the robustness of the proposed RISE, we consider scenarios where Assumption 1 may not hold. To test the violation of positivity assumption, using the same setting as in Example 2, we consider an extreme propensity score, or the probability of being treated given X and S . Specifically, we let A satisfy $\log\{P(A_i = 1|X_i)/P(A = 0|X_i)\} = -1.2(-S_i + X_{i1} - X_{i2} + X_{i3} - X_{i4} + X_{i5} - X_{i6})$. To test the unconfoundedness assumption, a random normal noise, $e \sim N(0, 1)$ is added to X_1 in the setting of Example 2. The simulation results are presented in Table 7 and Table 8 respectively.

Table 5: Simulation results for Example 2 with continuous S using 0.1 quantile criterion and 0.5 quantile criterion, respectively. Standard error in parenthesis. The proposed RISE has more strengths when τ is small, as the algorithm aims to improve the worst-outcome scenarios.

Type of S	τ	IDR	Obj. (all)	Obj. (vulnerable)	Value (all)	Value (vulnerable)
Cont.	0.1	Base	7.93 (0.03)	7.92 (0.03)	17.7 (0.02)	8.64 (0.07)
		Exp	8.88 (0.05)	8.85 (0.05)	17.8 (0.02)	10.6 (0.12)
		PT-Base	6.97 (0.02)	6.95 (0.02)	17.9 (0.03)	6.65 (0.04)
		PT-Exp	7.11 (0.02)	7.08 (0.03)	18.0 (0.03)	6.96 (0.05)
		RISE	13.8 (0.01)	13.7 (0.02)	16.9 (0.01)	20.9 (0.03)
Cont.	0.5	Base	17.3 (0.04)	17.2 (0.04)	17.7 (0.02)	23.8 (0.19)
		Exp	17.2 (0.03)	17.4 (0.03)	17.8 (0.02)	22.1 (0.17)
		PT-Base	17.3 (0.05)	17.3 (0.05)	18.0 (0.03)	23.9 (0.26)
		PT-Exp	17.4 (0.05)	17.4 (0.05)	18.1 (0.03)	24.0 (0.25)
		RISE	17.4 (0.04)	17.4 (0.04)	17.8 (0.02)	24.0 (0.22)

Table 6: Simulation results for scenario when S is a noise variable. Vulnerable subjects cannot be defined as S is not important in the example. The estimated objective and value function are similar for the compared IDRs, which indicates the robustness of the proposed RISE.

Type of S	IDR	Obj. (all)	Obj. (vulnerable)	Value (all)	Value (vulnerable)
Disc.	Base	27.5 (0.03)	-	27.5 (0.06)	-
	Exp	27.5 (0.03)	-	27.5 (0.06)	-
	PT-Base	27.5 (0.02)	-	27.5 (0.03)	-
	PT-Exp	27.5 (0.02)	-	27.5 (0.03)	-
	RISE	27.5 (0.03)	-	27.5 (0.06)	-
Cont.	Base	27.2 (0.04)	-	27.3 (0.07)	-
	Exp	27.2 (0.04)	-	27.3 (0.07)	-
	PT-Base	27.2 (0.04)	-	27.3 (0.07)	-
	PT-Exp	27.2 (0.04)	-	27.3 (0.07)	-
	RISE	27.2 (0.04)	-	27.3 (0.07)	-

Table 7: Simulation results for Example 2 where the positivity assumption in Assumption 1b is nearly violated. Standard error in parenthesis.

Type of S	IDR	Obj. (all)	Obj. (vulnerable)	Value (all)	Value (vulnerable)
Disc.	Base	10.0 (0.03)	11.1 (0.03)	19.3 (0.02)	16.1 (0.04)
	Exp	8.80 (0.03)	9.77 (0.04)	19.5 (0.02)	13.6 (0.04)
	PT-Base	9.88 (0.03)	10.7 (0.04)	18.9 (0.02)	15.4 (0.05)
	PT-Exp	8.42 (0.03)	9.14 (0.04)	19.1 (0.02)	12.4 (0.05)
	RISE	13.5 (0.01)	14.0 (0.01)	17.3 (0.01)	22.0 (0.02)
Cont.	Base	11.5 (0.03)	11.5 (0.04)	17.5 (0.03)	13.1 (0.04)
	Exp	10.4 (0.04)	10.4 (0.05)	17.8 (0.04)	10.3 (0.05)
	PT-Base	11.0 (0.04)	10.9 (0.04)	17.7 (0.02)	11.8 (0.04)
	PT-Exp	9.63 (0.03)	9.61 (0.03)	18.0 (0.02)	8.38 (0.03)
	RISE	14.3 (0.01)	14.3 (0.02)	16.9 (0.01)	20.4 (0.02)

Table 8: Simulation results for Example 2 where the unconfoundedness assumption in Assumption 1c is violated. Standard error in parenthesis.

Type of S	IDR	Obj. (all)	Obj. (vulnerable)	Value (all)	Value (vulnerable)
Disc.	Base	7.65 (0.04)	8.44 (0.05)	19.3 (0.03)	11.1 (0.06)
	Exp	8.94 (0.05)	9.91 (0.06)	19.4 (0.02)	13.9 (0.06)
	PT-Base	6.84 (0.03)	7.35 (0.04)	18.9 (0.03)	8.91 (0.05)
	PT-Exp	7.95 (0.05)	8.62 (0.06)	19.1 (0.03)	11.4 (0.06)
	RISE	13.5 (0.01)	14.0 (0.01)	17.4 (0.01)	22.1 (0.02)
Cont.	Base	9.58 (0.03)	9.58 (0.03)	17.9 (0.02)	8.33 (0.05)
	Exp	10.2 (0.04)	10.2 (0.04)	17.8 (0.02)	9.83 (0.06)
	PT-Base	9.27 (0.02)	9.26 (0.03)	17.9 (0.02)	7.51 (0.03)
	PT-Exp	9.34 (0.02)	9.34 (0.03)	18.0 (0.02)	7.72 (0.03)
	RISE	14.2 (0.01)	14.1 (0.02)	16.9 (0.01)	20.1 (0.03)

E Additional Information and Results for Real-data Applications

E.1 Data Availability

The job training dataset [30] is available at <https://users.nber.org/~rdehejia/data/.nswdata2.html>. The ACTG175 dataset [19] is available in the R package `speff2trial`. The sepsis dataset [59] is proprietary and not publicly available. All data used in this work are deidentified.

E.2 Additional Background on the Sepsis Application

Sepsis is leading cause of acute hospital mortality and commonly results in multi-organ dysfunction among ICU patients [57, 45]. Clinically, treatment decisions for sepsis patients are needed to be made within a short period of time due to the rapid deterioration of patient conditions. Lactate and the Sequential Organ Failure Assessment (SOFA) score have been two important indicators of sepsis severity and has been found to be more useful for predicting the outcome of sepsis than other clinical vitals and comorbidity scores [21, 26, 60, 34]. Typically, information of baseline patient characteristics such as age, gender, race, and weight, and common vital signs such as usage of mechanical ventilation, respiratory rate, temperature, intravenous fluids, Glasgow Coma Scale score, platelets, blood urea nitrogen, white blood cell counts, glucose, and creatinine are obtained at the admission of patients. On the other hand, SOFA score combines performance of several organ systems in the body such as neurologic, blood, liver, and kidney [59, 32, 67, 25] and cannot be obtained directly. Lactate labs measures the level of lactic acid in the blood [2, 48, 14] and are less common in routine examination, which could be delayed in ordering. Hence, their information may not be available by the time of treatment decision due to multiple reasons including doctors' delayed ordering, long laboratory processing time, or the rapid deterioration of development of sepsis, which poses tremendous difficulties for early diagnosis and treatment decisions within a short time. According to the new definition of Sepsis-3 [60], a serum lactate level greater than 2 mmol/L is considered to be in critical conditions and is highly likely to indicate a septic shock. Also, a SOFA score greater than 6 has been associated with a higher mortality [70, 18].

E.3 Visualizations

We provide visualizations of features that are important in the estimated decision rules for the three real-data applications in Section 4.2. The Shapely additive explanations (SHAP) [33] is considered to be a united approach to explaining the predictions of any machine learning or black-box models. Figure 2, Figure 3, and Figure 4 presents the SHAP variable importance plots in the final weighted classification model by RISE and Exp, respectively, for the three real-data applications. Correlations between the feature and their SHAP value are highlighted in color. The red color means a feature is positively correlated with assigning treatment $A = 1$ and the blue indicates a negative correlation. Overall, the direction of correlation is similar for RISE and Exp, but their ranking of feature importance may be different.

Fairness in a job training program. Figure 2 presents the SHAP variable importance plots in the final weighted classification model by RISE and Exp, respectively. We observe that whether having a high school diploma and income in 1974 are two important features in the variable important plot by RISE, while incomes in 1974 and in 1975 are important by Exp. It seems that being no degree and low income in 1974 has a higher chance of assigning $A = 1$ (to receive the job training program) by RISE, while low income in 1974 and but a higher income in 1975 may be associated with assigning $A = 1$ by Exp.

Improvement of HIV treatment. Figure 3 presents the SHAP variable importance plots in the final weighted classification model by RISE and Exp, respectively. We observe that age and CD4 T-cell counts are two important features in the variable important plot by RISE, while weight and number of days of previously received antiretroviral therapy are important by Exp. It seems that being of a younger age and high CD4 T-cell count has a higher chance of assigning $A = 1$ (zidovudine combined with didanosine) by RISE, while being of a larger weight and few days of previously received antiretroviral therapy may be associated with assigning the treatment by Exp.

Safe resuscitation for patients with sepsis. Figure 4 presents the SHAP variable importance plots in the final weighted classification model by RISE and Exp, respectively. We observe that Glasgow Coma Scale score, age, and platelets appears to be important features in both the plot by RISE and that by Exp. Other important features in the plot by RISE include temperature and blood urea nitrogen, where in the plot by Exp, respiratory rate and white blood cell counts are of top importance. Being in a low temperature with a high blood urea nitrogen tends to be predicted as $A = 1$ (to assign vasopressors) by RISE while being of higher respiratory rate with high white blood cell counts tends to be predicted as $A = 1$ by Exp.




Figure 2: Visualization for the job training program: SHAP variable importance plots for decision rules RISE (a) and Exp (b), respectively. Covariates (X) are ranked by variable importance in descending order. Correlations between the feature and their SHAP value are highlighted in color. The red color means a feature is positively correlated with assigning treatment $A = 1$ and the blue indicates a negative correlation.




Figure 3: Visualization for the ACTG175 dataset: SHAP variable importance plots for decision rules RISE (a) and Exp (b), respectively. Covariates (X) are ranked by variable importance in descending order. Correlations between the feature and their SHAP value are highlighted in color. The red color means a feature is positively correlated with assigning treatment $A = 1$ and the blue indicates a negative correlation.




Figure 4: Visualization for the sepsis data: SHAP variable importance plots for decision rules RISE (a) and Exp (b), respectively. Covariates (X) are ranked by variable importance in descending order. Correlations between the feature and their SHAP value are highlighted in color. The red color means a feature is positively correlated with assigning treatment $A = 1$ and the blue indicates a negative correlation.