

---

# Learning Pareto-Efficient Decisions with Confidence

---

**Sofia Ek**

Uppsala University, Sweden

**Dave Zachariah**

Uppsala University, Sweden

**Petre Stoica**

Uppsala University, Sweden

## Abstract

The paper considers the problem of multi-objective decision support when outcomes are uncertain. We extend the concept of Pareto-efficient decisions to take into account the uncertainty of decision outcomes across varying contexts. This enables quantifying trade-offs between decisions in terms of tail outcomes that are relevant in safety-critical applications. We propose a method for learning efficient decisions with statistical confidence, building on results from the conformal prediction literature. The method adapts to weak or nonexistent context covariate overlap and its statistical guarantees are evaluated using both synthetic and real data.

## 1 INTRODUCTION

In many decision-making tasks, it is important to balance *trade-offs* between several outcome objectives. For instance, a chosen drug can have intended effects but also side effects; a portfolio can be associated with various risks; and a cost reduction in production may have environmental impacts. These trade-offs can be formally studied by considering an optimization problem with  $m$  different objectives:

$$\max_{x \in \mathcal{X}} (f_1(x), f_2(x), \dots, f_m(x)), \quad (1)$$

where  $x$  is a decision variable and  $f_k(x)$  is the resulting objective for outcome  $k$  that we wish to maximize. A feasible decision  $x^*$  strictly dominates an alternative decision  $x$  if it is superior in at least one objective and not inferior in any other objective. It is said to be *Pareto-efficient* if it is not strictly dominated by any other decision (Miettinen, 2012; Jin and Sendhoff,

2008; Debreu, 1954). By evaluating the objectives in (1) for each Pareto-efficient decision  $x^*$  we can analyze the trade-offs between different outcomes in a specific decision problem.

We are here interested in problems where the objective functions in (1) are *unknown* but we have access to data from past outcomes of decisions. This data is often subject to uncertainty. We therefore consider outcome  $k$  from decision  $x$  to have a *random reward*  $y_k$ . In this paper we address the problem of learning Pareto-efficient decisions with random rewards from past training data.

One approach is to consider the mean reward (conditioned on the decision) as an objective function:  $f_k(x) \equiv \mathbb{E}[y_k|x]$ , which can be approximated by a model learned from data. Rolf et al. (2020) balance two mean objectives that are averaged over randomized decisions in online recommendation systems and data collection in sensitive ecosystems.

However, mean objective functions will not characterize important tail events nor do they represent typical outcomes when the conditional distributions  $p(y_k|x)$  are skewed. A substantial portion of outcome rewards may indeed fall below the mean. In safety-critical scenarios, such as clinical decision support with significant negative rewards, mean objective functions are thus inadequate. An additional complication is that an objective function  $f_k(x)$  learned from data will differ from  $f_k(x)$  and can thus lead to erroneous decisions. This is aggravated by the fact that past training data differs systematically from future test data obtained after a given intervention. (See below for further details.) The problem is also related to that of off-policy learning from logged data for contextual bandit with single objectives (Langford and Zhang, 2008; Swaminathan and Joachims, 2015; Joachims et al., 2018; Su et al., 2020), but here the past policy is unknown.

In order to tackle these challenges, we seek to endow the concept of Pareto-efficiency with a notion of statistical uncertainty (Section 2). Then we develop a method that learns efficient decisions with confidence from training data (Section 3). The method is eval-

---

Proceedings of the 25<sup>th</sup> International Conference on Artificial Intelligence and Statistics (AISTATS) 2022, Valencia, Spain. PMLR: Volume 151. Copyright 2022 by the author(s).

uated using both synthetic and real data (Section 4). The proposed method has the following main features:

- it considers the lower tail rewards at any specified level,
- it provides finite-sample coverage guarantees for the rewards even when the data-generating process is unknown,
- and it enables the study of decision trade-offs using a Pareto-frontier that takes reward uncertainties into account.

The proposed method learns models of conditional reward quantiles and leverages results from the conformal prediction literature to identify Pareto-efficient decisions that have statistical validity (Koenker and Hallock, 2001; Meinshausen and Ridgeway, 2006; Vovk et al., 2005; Shafer and Vovk, 2008). The problem setting lies in the intersection of several fields; decision theory (Chernoff and Moses, 2012), multi-objective optimization (Miettinen, 2012), causal inference (Pearl, 2009; Peters et al., 2017), and conformal prediction (Balasubramanian et al., 2014). The main difference between our proposed method and the method in Rolf et al. (2020) is that we evaluate multiple decisions and random tail rewards instead of randomized binary decisions and expected rewards. Importantly, our method can also provide finite sample guarantees.

## 2 PROBLEM FORMULATION

We consider discrete decisions  $x$  in a space  $\mathcal{X}$ . Each decision has  $m$  outcome rewards, represented by a vector  $\mathbf{y}$  in a reward space  $\mathcal{Y} \subset \mathbb{R}^m$ . The number  $m$  is typically small, the examples below illustrate cases with  $m = 2$  rewards. For generality, we consider decisions taken in different contexts described by a  $d$ -dimensional covariate vector  $\mathbf{z}$  that may affect the outcomes.

Suppose that for any decision  $x$ , we have a vector  $\mathbf{y}^\alpha(x, \mathbf{z})$  in reward space that *lower bounds* each reward in  $\mathbf{y}$  across contexts  $\mathbf{z}$  with a probability of at least  $1 - \alpha$ . Formally, we express this using element-wise inequality as:

$$\mathbb{P}(\mathbf{y}^\alpha(x, \mathbf{z}) \preceq \mathbf{y} \mid x) \geq 1 - \alpha, \quad \forall x \in \mathcal{X}. \quad (2)$$

If a nontrivial bound can be found, we use it to define efficient decisions that take into account the uncertainty of outcomes, including tail events with very low or negative rewards.

We say that decision  $x^*$  *strictly dominates another decision  $x$  with confidence  $1 - \alpha$*  if its reward boundary is not inferior,

$$\mathbf{y}^\alpha(x, \mathbf{z}) \preceq \mathbf{y}^\alpha(x^*, \mathbf{z}),$$




Figure 1: Example of decisions  $x \in \{0, 1, 2, 3, 4\}$  with random outcome rewards  $\mathbf{y} = [y_1, y_2]$ . The small dots show past data of rewards observed from decisions taken across different contexts. From this data we learn Pareto-efficient decisions with confidence 80% across contexts. For a given context specified by  $\mathbf{z}$ , this set was  $\mathcal{X}^\alpha(\mathbf{z}) = \{0, 1, 4\}$  with reward bounds  $\mathbf{y}^\alpha(x, \mathbf{z})$  shown as solid green dots. Decisions 2 and 3 (red crosses) are thus strictly dominated. The set of efficient decisions and its corresponding reward bounds form a frontier  $\mathcal{Y}^\alpha(\mathbf{z})$  (dashed) along which a user can study reward trade-offs.

and that it is superior in at least one outcome  $k$  for which  $y_k^\alpha(x, \mathbf{z}) < y_k^\alpha(x^*, \mathbf{z})$ . The definition above endows the standard notion of Pareto-efficiency with statistical properties. We can now define the object of interest in this paper,  $\mathcal{X}^\alpha(\mathbf{z})$ , the set of all *efficient decisions with confidence  $1 - \alpha$* . Each decision  $x$  in  $\mathcal{X}^\alpha(\mathbf{z})$  is not dominated by any other decision and the uncertain outcome rewards  $\mathbf{y}$  fall below the boundary  $\mathbf{y}^\alpha(x, \mathbf{z})$  with a probability no greater than  $\alpha$  across contexts.

Using  $\mathcal{X}^\alpha(\mathbf{z})$  we can construct a corresponding *frontier*  $\mathcal{Y}^\alpha(\mathbf{z})$  in reward space, which enables a user to quantify trade-offs between different efficient decisions. Figure 1 illustrates such a frontier in an example with a decision space  $\mathcal{X} = \{0, 1, 2, 3, 4\}$  and outcomes rewards  $\mathbf{y} = [y_1, y_2]$ . We can see that two of the decisions are strictly dominated with a confidence of 80%. A user can now analyze the trade-offs between rewards when choosing among the three efficient decisions. (The full setup is described in the numerical experiment section.)

The problem we consider in this paper is to *learn* confidence-based efficient decisions  $\mathcal{X}^\alpha(\mathbf{z})$  and their corresponding frontier  $\mathcal{Y}^\alpha(\mathbf{z})$  from finite training data. We use past data on decisions, contexts and their rewards  $(x_i, \mathbf{y}_i, \mathbf{z}_i)$  drawn independently and identically from an unknown distribution that admits a causal

factorization:

$$\tilde{p}(x, \mathbf{y}, \mathbf{z}) = p(\mathbf{z})p(x|\mathbf{z})p(\mathbf{y}|x, \mathbf{z}). \quad (3)$$

In the case of past experimental data, the decision policy described by  $p(x|\mathbf{z})$  is known. For instance, in medical trials, treatment  $x$  could be assigned at random and independent of context  $\mathbf{z}$ . We also consider using past observational data, where  $p(x|\mathbf{z})$  is unknown and has to be learned from data. In this case, one must assume that  $\mathbf{z}$  contains all relevant data and there are no unblocked confounding variables (Peters et al., 2017, Chapter 9). We do *not*, however, need to assume covariate overlap for all decisions. This is an important difference from both causal inference and contextual bandit literatures (Imbens and Rubin, 2015; Oberst et al., 2020; D’Amour et al., 2021).

We shall now proceed to develop a method that learns  $\alpha$ -efficient decisions  $\mathcal{X}^\alpha(\mathbf{z})$  from finite training data by constructing nontrivial and context-adaptive bounds  $\mathbf{y}^\alpha(x, \mathbf{z})$  that satisfy the statistical guarantee (2).

### 3 METHOD

To construct  $\mathcal{X}^\alpha(\mathbf{z})$  we need boundary vectors  $\mathbf{y}^\alpha(x, \mathbf{z})$  that satisfy (2). We here use the union bound to avoid imposing restrictive modelling assumptions on the joint distribution of rewards and to avoid complex modelling of multivariate quantile functions. Using this bound for the complement of the event in (2) we have:

$$\mathbb{P}\left(\bigcup_{k=1}^m y_k < y_k^\alpha(x, \mathbf{z}) \mid x\right) \leq \sum_{k=1}^m \mathbb{P}(y_k < y_k^\alpha(x, \mathbf{z}) \mid x). \quad (4)$$

Thus, by constructing bounds for each reward that satisfy

$$\mathbb{P}(y_k < y_k^\alpha(x, \mathbf{z}) \mid x) \leq \frac{\alpha}{m}, \quad \forall k,$$

constraint (2) is also satisfied.

We consider now the conditional *quantile function* of an individual reward  $y$ . For continuous rewards, we can express it as

$$q^\alpha(x, \mathbf{z}) = \inf \left\{ \tilde{y} : \frac{\alpha}{m} = \mathbb{P}(y < \tilde{y} \mid x, \mathbf{z}) \right\}.$$

Then  $y^\alpha(x, \mathbf{z}) := q^\alpha(x, \mathbf{z})$  would provide the tightest lower bound for the reward. But the quantile function is unknown and must be estimated from data

$$\mathcal{D} = \left( (x_i, y_i, \mathbf{z}_i) \right)_{i=1}^n,$$

drawn from the training distribution (3), see, e.g., Koenker and Hallock (2001); Meinshausen and Ridgeway (2006). However, using an estimate  $y^\alpha(x, \mathbf{z}) :=$

$\hat{q}^\alpha(x, \mathbf{z})$  in lieu of the unknown function faces two fundamental challenges. First, this choice would not ensure that (2) holds for finite  $n$ . When  $\mathbf{z}$  contains continuous covariates, or is of a moderately large dimension, (2) is at best approximately satisfied in the large-sample case provided that we use a well-specified model of the quantile function. Second, the training data is itself a problem because (3) differs from the distribution that follows when intervening on a *specified* decision variable  $x := x^*$ :

$$p(x, \mathbf{y}, \mathbf{z}) = p(\mathbf{z})\mathbb{1}(x = x^*)p(\mathbf{y}|x, \mathbf{z}), \quad (5)$$

where  $\mathbb{1}(\cdot)$  is the indicator function.

We will here take a different approach by allowing for *any* flexible machine learning method to fit a quantile function and then form a boundary:

$$\boxed{y^\alpha(x, \mathbf{z}) := \hat{q}^\alpha(x, \mathbf{z}) - \kappa^\alpha(x)}, \quad (6)$$

using an *adjustment* variable  $\kappa^\alpha(x)$ . By leveraging recent results in the conformal prediction literature (Vovk et al., 2005; Lei et al., 2015; Romano et al., 2019; Tibshirani et al., 2019), we will show that it is possible to learn tight reward boundaries (6) using flexible methods that ensure the rewards are bounded at the specified probability  $1 - \alpha$  in (2). The methodology is based on randomly splitting the data  $\mathcal{D}$  into two parts, denoted  $\mathcal{D}'$  and  $\mathcal{D}''$ . The splitting allows for a computationally efficient algorithm that at the same time fulfils (2).

Let  $\hat{q}^\alpha(x, \mathbf{z}; \mathcal{D}')$  be any fitted quantile function using  $\mathcal{D}'$ . Next, define the residuals using the remaining data:

$$r_i = \hat{q}^\alpha(x_i, \mathbf{z}_i; \mathcal{D}') - y_i, \quad i \in \mathcal{D}''. \quad (7)$$

Construct a discrete distribution of the residual  $r$ :

$$p(r) = \sum_{i \in \mathcal{D}''} p_i \mathbb{1}(r = r_i) + p_\infty \mathbb{1}(r = \infty), \quad (8)$$

where  $\{p_i\}$  are probability weights.

**Theorem 3.1.** *Consider any  $(x, \mathbf{z})$  drawn from the interventional distribution (5) and let the adjustment variable  $\kappa^\alpha(x)$  be the  $(1 - \frac{\alpha}{m})$ -th quantile of the distribution (8) with the following probability weights*

$$p_i = \frac{w(x_i, \mathbf{z}_i)}{\sum_{j \in \mathcal{D}''} w(x_j, \mathbf{z}_j) + w(x, \mathbf{z})}, \quad (9)$$

$$p_\infty = \frac{w(x, \mathbf{z})}{\sum_{j \in \mathcal{D}''} w(x_j, \mathbf{z}_j) + w(x, \mathbf{z})},$$

using the function

$$w(x, \mathbf{z}) = \frac{\mathbb{1}(x = x^*)}{p(x|\mathbf{z})}. \quad (10)$$

Then the boundary (6) satisfies:

$$\mathbb{P}(y < y^\alpha(x, \mathbf{z}) | x) \leq \frac{\alpha}{m}. \quad (11)$$

Constructing these boundaries for all  $m$  rewards therefore satisfies (2).

*Proof.* The weights are used to adjust densities from the training distribution to the intervened distribution. The first part of the proof builds on the techniques used in Lemma 3 in Tibshirani et al. (2019). We modified it here for the case when  $r$  is only a function of the subset  $\mathcal{D}'$ . The second part of the proof builds the techniques used in Theorem 1 and Corollary 1 in Romano et al. (2019) and Tibshirani et al. (2019), respectively. Our main contribution is the case of one-sided conformal boundaries using split quantile conformal regression under distribution shifts. This differs from the focus of the two cited papers on which the proof is build.

Let  $\mathcal{D}_+$  denote the sequence of training data  $\mathcal{D}$  and a new sample  $(x, y, \mathbf{z})$  drawn from (5), i.e.,

$$\mathcal{D}_+ = (\mathcal{D}', \mathcal{D}'', (x, y, \mathbf{z})).$$

This sequence has a joint distribution with a density function that can be expressed as

$$\begin{aligned} \prod_{i=1}^n \tilde{p}(x_i, y_i, \mathbf{z}_i) \cdot \tilde{p}(x, y, \mathbf{z}) w(x, \mathbf{z}) &= \tilde{p}(\mathcal{D}_+) w(x, \mathbf{z}) \\ &= \tilde{p}(\mathcal{S}_+) w(x, \mathbf{z}), \end{aligned}$$

where  $\mathcal{S}_+$  is an unordered set of elements from  $\mathcal{D}_+$ . Note that there are multiple sequences  $\mathcal{D}_+$  that can be obtained by permuting the elements in  $\mathcal{S}_+$ .

Let  $e_i$  denote the event that the new sample  $(x, y, \mathbf{z})$  equals the  $i$ th sample  $(x_i, y_i, \mathbf{z}_i)$  in the *subsequence*

$$\mathcal{D}_+'' = (\mathcal{D}'', (x, y, \mathbf{z})),$$

given the unordered set  $\mathcal{S}_+$ . The conditional probability of this event can be expressed as

$$\begin{aligned} \mathbb{P}(e_i | \mathcal{S}_+) &= \frac{\mathbb{P}(e_i, \mathcal{S}_+)}{\sum_{j \in \mathcal{D}''} \mathbb{P}(e_j, \mathcal{S}_+) + \mathbb{P}(e, \mathcal{S}_+)} \\ &= \frac{w(x_i, \mathbf{z}_i)}{\sum_{j \in \mathcal{D}''} w(x_j, \mathbf{z}_j) + w(x, \mathbf{z})} \\ &= p_i, \end{aligned}$$

where the first equality follows from the law of total probability and the second from considering all permutations  $\sigma$  over the set  $\mathcal{S}_+$ , such that the final element (new sample) in every resulting sequence is equal to  $(x_i, y_i, \mathbf{z}_i)$ . That is,

$$\begin{aligned} \mathbb{P}(e_i, \mathcal{S}_+) &= \sum_{\sigma: \sigma(n+1) = |\mathcal{D}''| + i} w(x_i, \mathbf{z}_i) \tilde{p}(\mathcal{S}_+) \\ &= w(x_i, \mathbf{z}_i) \tilde{p}(\mathcal{S}_+) n!. \end{aligned}$$

Each sample  $(x_i, y_i, \mathbf{z}_i)$  has a residual  $r_i$  of the form (7). The conditional probability that the residual of the new sample  $r$  equals  $r_i$  is equal to

$$\mathbb{P}(r = r_i | \mathcal{S}_+) = \mathbb{P}(e_i | \mathcal{S}_+) = p_i,$$

which holds since we condition on the full set  $\mathcal{S}_+$  and so  $\mathcal{D}'$  in (7) is fixed. Thus the residual  $r$  given  $\mathcal{S}_+$  follows the discrete distribution:

$$\sum_{i \in \mathcal{D}''} p_i \mathbb{1}(r = r_i) + p_{n+1} \mathbb{1}(r = r_{n+1}). \quad (12)$$

Let  $\tilde{\kappa}^\alpha$  denote its  $(1 - \frac{\alpha}{m})$ -th quantile, so that

$$\mathbb{P}\{r > \tilde{\kappa}^\alpha | \mathcal{S}_+\} \leq \frac{\alpha}{m}.$$

Now in lieu of (12), consider (8) where the only point mass for the final residual is moved to the limit of its range. Thus the quantile of (8) cannot decrease,  $\kappa^\alpha \geq \tilde{\kappa}^\alpha$ , so that

$$\mathbb{P}\{r > \kappa^\alpha | \mathcal{S}_+\} \leq \mathbb{P}\{r > \tilde{\kappa}^\alpha | \mathcal{S}_+\} \leq \frac{\alpha}{m}.$$

Using the fact that  $\mathbb{P}\{r > \kappa^\alpha | \mathcal{S}_+\} = \mathbb{P}\{y < \hat{q} - \kappa^\alpha | \mathcal{S}_+\}$ , we obtain

$$\mathbb{P}\{y < \hat{q} - \kappa^\alpha\} \leq \frac{\alpha}{m},$$

after marginalization.

Since the interventional distribution (5) (and  $w(x, \mathbf{z})$ ) contain  $\mathbb{1}(x = x^*)$ , the resulting probability is conditional on any  $x^* \in \mathcal{X}$  in (11) since all other decision variables have zero probability. Using (4) it follows that (2) is satisfied.  $\square$

There is a wide range of possible models and learning methods to fit  $\hat{q}^\alpha(x, \mathbf{z})$  in (6); from simple linear models to highly flexible neural network or random forest-based models, cf. Koenker and Hallock (2001); Meinshausen and Ridgeway (2006). Whichever method one chooses, the methodology described above ensures that all rewards for each decision will be at least as large as their bounds at a certifiable probability of  $1 - \alpha$ . The choice of fitted  $\hat{q}^\alpha(x, \mathbf{z})$  only affects the *tightness* of the bounds. The efficient decisions  $\mathcal{X}^\alpha(\mathbf{z})$  are thus learned with a model-agnostic confidence of  $1 - \alpha$ .

*Remark 1:* With covariate vectors  $\mathbf{z}$  of a given dimension, the computational cost comprise of fitting  $m \cdot |\mathcal{X}|$  quantile functions and sorting weighted residuals. Each sorting has a run time  $\mathcal{O}(n_x \log(n_x + 1))$ , where  $n_x$  is the number of samples in the set  $\mathcal{D}''$  that have a given decision  $x$ , see Algorithm 1 in the supplementary material. Note that these  $m \cdot |\mathcal{X}|$  tasks can be performed in parallel.

*Remark 2:* In the formulation above, we have set an equal probability level of  $\alpha/m$  to each reward. But

it is entirely possible to specify different levels across rewards that still satisfy (2). This would be relevant in scenarios where certain rewards are more critical to certify than others.

*Remark 3:* The conformal prediction literature is focused on the construction of outcome prediction intervals, but we are here only interested in a lower bound on rewards  $\mathbf{y}$ . See also Linusson et al. (2014).

*Remark 4:* When the training data is only observational and the decisions are taken with unknown policies the weight function (10) must be estimated accurately (Lei and Candès, 2021). This can be done by supervised learning of a model of  $p(\mathbf{z}|x)$ . Alternatively, the weights can be expressed as:

$$w(x, \mathbf{z}) \equiv \frac{\mathbf{1}(x = x^*)}{p(\mathbf{z}|x)p(x)} \sum_{x' \in \mathcal{X}} p(\mathbf{z}|x')p(x'), \quad (13)$$

where a model of  $p(\mathbf{z}|x)$  can be learned in an unsupervised manner.

## 4 NUMERICAL EXPERIMENTS

To illustrate the key concepts of our method for learning  $\mathcal{X}^\alpha(\mathbf{z})$  and verify the statistical properties of the frontier  $\mathcal{Y}^\alpha(\mathbf{z})$ , we conduct numerical experiments with both synthetic and real-world data. For convenience of illustration, we will only consider  $m = 2$  rewards. We split the training data  $\mathcal{D}$  into two parts  $\mathcal{D}'$  and  $\mathcal{D}''$  of equal size.

The conformal inference part of our implementation is based on the code from Romano et al. (2019) which we extended to account for the distributional shift that arises between training and interventional distributions (3) and (5).<sup>1</sup>

### 4.1 Synthetic data

We consider a problem with five decision alternatives  $\mathcal{X} = \{0, 1, 2, 3, 4\}$  and a one-dimensional covariate  $z$ . We begin by specifying the unknown training distribution (3).

The rewards are drawn from  $p(\mathbf{y}|x, z)$  as:

$$\begin{cases} y_1 &= a_x + b_x f\left(\frac{z-55}{9}\right) + u_0, \\ y_2 &= c_x + d_x f\left(\frac{z-50}{8}\right) + u_1, \end{cases} \quad (14)$$

where  $f(\cdot)$  is the sigmoid function and the noise terms  $u_0$  and  $u_1$  are drawn jointly from  $\mathcal{N}(0, 0.04 \cdot \begin{bmatrix} 1 & \rho \\ \rho & 1 \end{bmatrix})$  with a correlation coefficient  $\rho = -0.2$  unless another value is explicitly mentioned. (The rewards are ensured to

be nonnegative by truncation in the rare cases.) The coefficients in (14) are given in Table 1. The covariate is drawn from  $p(z)$  as  $z \sim \mathcal{N}(60, 100)$ .

Table 1: The constants in (14).

	$x = 0$	$x = 1$	$x = 2$	$x = 3$	$x = 4$
$a_x$	2.4	0.7	0.8	2.0	1.2
$b_x$	-1.4	1.5	1.0	-1.2	1.0
$c_x$	0.0	2.2	0.6	0.0	2.2
$d_x$	2.4	-1.5	1.0	2.0	-1.0

In the training data, decisions have been made in different contexts described by  $z$ . We let the decision assignment be as follows

$$x = \begin{cases} 0 & \text{if } 0 \leq s < 0.2, \\ 1 & \text{if } 0.2 \leq s < 0.4, \\ 2 & \text{if } 0.4 \leq s < 0.6, \\ 3 & \text{if } 0.6 \leq s < 0.8, \\ 4 & \text{if } 0.8 \leq s < 1.0, \end{cases}$$

where  $s$  is a random variable. In the case of experimental data, we consider uniform random assignments:  $s \sim \mathcal{U}(0, 1)$  and take  $p(x|z)$  to be known. In the case of observational data, we test the limits of our method by considering an unbalanced scenario where some of the decisions are rare so that the unknown  $p(x|z)$  is near zero (thus yielding weak covariate overlap). This is done by setting

$$s = u \times f\left(\frac{70 - z}{5}\right), \text{ where } u \sim \mathcal{U}(0, 1).$$

We draw  $n = 1000$  training data points to learn  $\mathcal{X}^\alpha(z)$  with confidence level  $1 - \alpha$  using a quantile random forest method (Meinshausen and Ridgeway, 2006), more details are provided in the supplementary material. We verify that the method satisfies (2) at the specified confidence level in the experimental data case in Fig. 2.

In the observational data case, we learn a (two-component) Gaussian mixture model of  $p(z|x)$  in (13). Then (2) will only hold approximately depending on the accuracy of (10), see Lei and Candès (2021). The accuracy of probability weights can be assessed using model validation methods. In Fig. 3 we evaluate bounds obtained with approximate weights and find (2) to be satisfied in this challenging setting with poor data overlap. In Fig. 3a we also include the case where the outcomes are even more correlated by setting  $\rho = -1.0$ , and observe only minor differences. In Fig. 3b, we see that due to the uneven past decision policy, data on decision  $x = 4$  is sparse and the resulting bound becomes notably more conservative than for the other decisions.

<sup>1</sup>The code is available here: <https://github.com/sofiaek/learning-pareto>.

The bounds for  $x = 0$  and  $1$  are the least conservative; the probabilities of rewards exceeding them are very close to  $1 - \alpha$ .

We illustrate the resulting learned frontier  $\mathcal{Y}^\alpha(z)$  for  $1 - \alpha = 80\%$  and two different contextual covariates  $z = 56$  and  $z = 68$  in Figs. 4a and 4b, respectively. The learned efficient decisions for the two different contexts are:  $\mathcal{X}^\alpha(56) = \{0, 4\}$  and  $\mathcal{X}^\alpha(68) = \{0, 1, 2\}$  and thus are context-dependent. (Figure 1 shows  $z = 46$ .) The method adapts to the lack of data for decision  $x = 4$  by setting the lower bound to the minimum possible rewards, represented by a black star in Fig. 4b.

## 4.2 Semi-real data

Next, we consider real-world data from the Tennessee Student/Teacher Achievement Ratio (STAR) study (Achilles et al., 2008; Krueger, 1999). In a four-year class size study teachers and students (in kindergarten up until third grade) were randomly assigned into one of three interventions: small, regular and regular-with-aide class size. Small class size corresponds to 13 to 17 students per teacher, regular class size to 22 to 25 students per teacher and regular-with-aide class size to 22 to 25 students per teacher plus a full-time teacher assistant. The study was thus a randomized controlled trial, with interventions assigned by a given random policy  $p(x)$ .

We follow the procedure in Kallus et al. (2018) and take the class-type at first grade as the treatment, since many students were not part of the study in kindergarten. We use 11 pre-intervention covariates in  $\mathbf{z}$  for each student including: gender, birth month, teacher career, teacher experience and if free lunch was given or not. A more detailed description of the covariates is provided in the supplementary material.

Outcome  $y_1$  is an achievement test score at the end of first grade (the sum of standardized math, reading and listening test scores). Students that were not part of the STAR study in first grade or had missing outcome  $y_1$  are excluded, and the remaining data set includes 6322 students from 75 different schools. To study a second outcome  $y_2$ , the (negative) cost per student, we generate a synthetic value following Hill (2011)

$$\begin{aligned} y_2|x = 0 &\sim \mathcal{N}(-(\mathbf{z}^\top \beta - \omega_0 + \mu), \sigma^2), \\ y_2|x = 1 &\sim \mathcal{N}(-(\exp[(\mathbf{z} + 0.5)^\top \beta] + \mu), \sigma^2), \\ y_2|x = 2 &\sim \mathcal{N}(-(\mathbf{z}^\top \beta - \omega_2 + \mu), \sigma^2), \end{aligned}$$

where  $\mu = 10$  and  $\sigma = 1$ . The parameter  $\beta$  is drawn randomly from  $(0.0, 0.1, 0.2, 0.3, 0.4)$  with probabilities  $(0.4, 0.15, 0.15, 0.15, 0.15)$ . These probabilities are changed compared to Hill (2011), since our model has fewer covariates. We select  $\omega_0$  and  $\omega_2$  so that the effect

of the ‘treatment on the treated’ is 4 and 2, respectively.

The conditional quantiles are estimated with a quantile neural net (Taylor, 2000; Steinwart and Christmann, 2011).

The learned Pareto-frontier for a group of students with a given set of covariates is shown in Fig. 5a. Examples of two such contextual covariates are school in rural areas and teachers with one year teaching experience. Here no decision is dominated at the confidence level 80%, but the trade-offs between test scores and costs become transparent: small class sizes yield the best test scores but are the most costly interventions, while regular class sizes with aide increase test scores by a lesser amount but at a substantially lower cost. A learned Pareto-frontier for a different group of students is shown in Fig. 5b. Their covariates differ for instance by school located in suburban areas and teachers with 10 years teaching experience. Here the decision ‘regular-with-aide’ class size is dominated by the other two decisions.

In the supplementary material, the statistical guarantee for the boundary of the synthetic reward  $y_2$  is validated.

## 5 DISCUSSION

Machine learning methods are typically designed to maximize a single objective. When used for decision making, they can therefore result in unbalanced outcomes that do not align with the objectives of a decision-maker (Christian, 2020). In binary classification tasks, for example, a user may need to balance false-positive and false-negative rates (Tong et al., 2018) and in recommendation systems the balance between welfare and profit is important (Rolf et al., 2020).

In safety-critical applications, such as clinical decision support, it is also important to account for low (even negative) rewards and provide statistical guarantees. In this paper we extended the concept of Pareto-efficiency to include statistical confidence. We then proposed a method that learns efficient decisions and a frontier that takes into account the lower tail rewards at any specified confidence level. Such a frontier can quantify trade-offs and provide valuable insight to decision makers due to its finite-sample statistical guarantees.

In the case when training data is observational rather than experimental, the (unknown) past decision policy must also be learned. Here care must be taken to assess the accuracy of any policy model using validation methods; inaccurate models invalidate the statistical properties of the learned frontier and may lead to erroneous decisions. It is also important to consider systematic differences in the distribution of contextual covariates during training and intervention. While

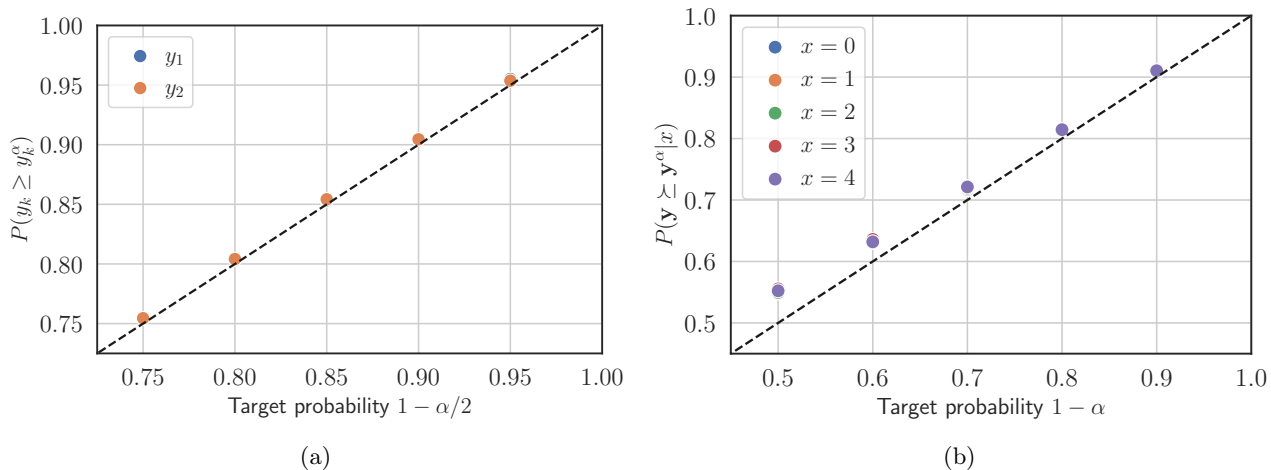


Figure 2: Evaluation of reward boundary  $\mathbf{y}^\alpha(x, z)$  of random rewards  $\mathbf{y}$  using 500 Monte Carlo simulations. Boundaries are designed to satisfy (2). Experimental data scenario with known random assignment probabilities  $p(x)$ . (a) Probability of  $y_k$  exceeding the boundary  $y_k^\alpha(x, z)$  marginalized across decisions. (b) Probability of  $\mathbf{y}$  exceeding the reward boundary  $\mathbf{y}^\alpha(x, z)$  under decision  $x$ . Note that the dots are overlapping in both figures.




Figure 3: Evaluation of reward boundary  $\mathbf{y}^\alpha(x, z)$  of random rewards  $\mathbf{y}$  using 500 Monte Carlo simulations. Observational data scenario where  $p(z|x)$  in (13) is estimated. (a) Probability of rewards exceeding the boundaries  $\mathbf{y}^\alpha(x, z)$  marginalized across decisions for two values of  $\rho$  in (14), marginalized across decisions. (b) Probability of  $\mathbf{y}$  exceeding the reward boundary  $\mathbf{y}^\alpha(x, z)$  under decision  $x$ . Note that several dots overlap in both figures.




Figure 4: Efficient decisions with confidence 80% in the observational data scenario for two different contextual covariates (a) and (b). In (a) the efficient set is  $\mathcal{X}^\alpha(z) = \{0, 4\}$  and decisions 1, 2 and 3 are thus strictly dominated, while in (b) efficient set is  $\mathcal{X}^\alpha(z) = \{0, 1, 2\}$  and decisions 3 and 4 are strictly dominated. In (b) there was no data for decision  $x = 4$  and the method adapts by setting the minimum possible reward as a boundary (highlighted as a black star). The dashed lines illustrate the frontiers  $\mathcal{Y}^\alpha(z)$  for each context.

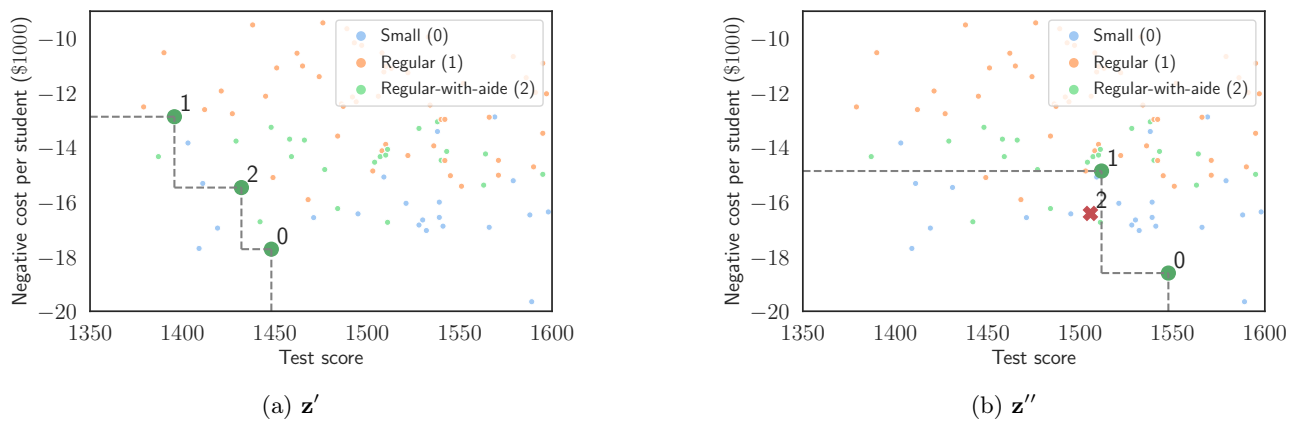


Figure 5: Frontiers  $\mathcal{Y}^\alpha(\mathbf{z})$  for student groups with two different sets of covariates learned from semi-real data with three different class size interventions. (a) All decisions are efficient so that  $\mathcal{X}^\alpha(\mathbf{z}') = \{0, 1, 2\}$ . (b) Efficient decisions are  $\mathcal{X}^\alpha(\mathbf{z}'') = \{0, 1\}$  since intervention 2 is dominated by 1. Confidence level 80%.



any known covariate shift can readily be incorporated by adjusting the weights, unknown shifts can produce misleading results.

### Acknowledgments

This research was partially supported by the Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by Knut and Alice Wallenberg Foundation, and the Swedish Research Council under contracts 2018-05040 and 2021-05022.

### References

- C.M. Achilles, Helen Pate Bain, Fred Bellott, Jayne Boyd-Zaharias, Jeremy Finn, John Folger, John Johnston, and Elizabeth Word. Tennessee’s Student Teacher Achievement Ratio (STAR) project, 2008. URL <https://doi.org/10.7910/DVN/SIWH9F>.
- Vineeth Balasubramanian, Shen-Shyang Ho, and Vladimir Vovk. *Conformal prediction for reliable machine learning: theory, adaptations and applications*. Newnes, 2014.
- Herman Chernoff and Lincoln E Moses. *Elementary decision theory*. Courier Corporation, 2012.
- B. Christian. *The Alignment Problem: Machine Learning and Human Values*. W. W. Norton, 2020.
- Gerard Debreu. Valuation equilibrium and Pareto optimum. *Proceedings of the National Academy of Sciences of the United States of America*, 40(7):588, 1954.
- Alexander D’Amour, Peng Ding, Avi Feller, Lihua Lei, and Jasjeet Sekhon. Overlap in observational studies with high-dimensional covariates. *Journal of Econometrics*, 221(2):644–654, 2021.
- Jennifer L Hill. Bayesian nonparametric modeling for causal inference. *Journal of Computational and Graphical Statistics*, 20(1):217–240, 2011.
- Guido W Imbens and Donald B Rubin. *Causal inference in statistics, social, and biomedical sciences*. Cambridge University Press, 2015.
- Yaochu Jin and Bernhard Sendhoff. Pareto-based multiobjective machine learning: An overview and case studies. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, 38(3):397–415, 2008.
- Thorsten Joachims, Adith Swaminathan, and Maarten de Rijke. Deep learning with logged bandit feedback. In *International Conference on Learning Representations*, 2018.
- Nathan Kallus, Aahlad Manas Puli, and Uri Shalit. Removing hidden confounding by experimental grounding. *Advances in neural information processing systems*, 31, 2018.
- Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Roger Koenker and Kevin F Hallock. Quantile regression. *Journal of economic perspectives*, 15(4): 143–156, 2001.
- Alan B Krueger. Experimental estimates of education production functions. *The quarterly journal of economics*, 114(2):497–532, 1999.
- John Langford and Tong Zhang. The epoch-greedy algorithm for multi-armed bandits with side information. In *Advances in neural information processing systems*, pages 817–824, 2008.
- Jing Lei, Alessandro Rinaldo, and Larry Wasserman. A conformal prediction approach to explore functional data. *Annals of Mathematics and Artificial Intelligence*, 74(1):29–43, 2015.
- Lihua Lei and Emmanuel J Candès. Conformal inference of counterfactuals and individual treatment effects. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 2021.
- Henrik Linusson, Ulf Johansson, and Tuve Löfström. Signed-error conformal regression. In *Pacific-Asia Conference on Knowledge Discovery and Data Mining*, pages 224–236. Springer, 2014.
- Nicolai Meinshausen and Greg Ridgeway. Quantile regression forests. *Journal of Machine Learning Research*, 7(6), 2006.
- Kaisa Miettinen. *Nonlinear multiobjective optimization*, volume 12. Springer Science & Business Media, 2012.
- Michael Oberst, Fredrik Johansson, Dennis Wei, Tian Gao, Gabriel Brat, David Sontag, and Kush Varshney. Characterization of overlap in observational studies. In *International Conference on Artificial Intelligence and Statistics*, pages 788–798. PMLR, 2020.
- Judea Pearl. *Causality*. Cambridge university press, 2009.
- Jonas Peters, Dominik Janzing, and Bernhard Schölkopf. *Elements of causal inference: foundations and learning algorithms*. The MIT Press, 2017.
- Esther Rolf, Max Simchowitz, Sarah Dean, Lydia T Liu, Daniel Bjorkegren, Moritz Hardt, and Joshua Blumenstock. Balancing competing objectives with noisy data: Score-based classifiers for welfare-aware machine learning. In *International Conference on Machine Learning*, pages 8158–8168. PMLR, 2020.
- Yaniv Romano, Evan Patterson, and Emmanuel Candès. Conformalized quantile regression. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.

- Andrew M Saxe, James L McClelland, and Surya Ganguli. Exact solutions to the nonlinear dynamics of learning in deep linear neural networks. In *International Conference on Learning Representations*, 2014.
- Glenn Shafer and Vladimir Vovk. A tutorial on conformal prediction. *Journal of Machine Learning Research*, 9(3), 2008.
- Ingo Steinwart and Andreas Christmann. Estimating conditional quantiles with the help of the pinball loss. *Bernoulli*, 17(1):211–225, 2011.
- Yi Su, Maria Dimakopoulou, Akshay Krishnamurthy, and Miroslav Dudík. Doubly robust off-policy evaluation with shrinkage. In *International Conference on Machine Learning*, pages 9167–9176. PMLR, 2020.
- Adith Swaminathan and Thorsten Joachims. Counterfactual risk minimization: Learning from logged bandit feedback. In *International Conference on Machine Learning*, pages 814–823, 2015.
- James W Taylor. A quantile regression neural network approach to estimating the conditional density of multiperiod returns. *Journal of Forecasting*, 19(4): 299–311, 2000.
- Ryan J Tibshirani, Rina Foygel Barber, Emmanuel Candes, and Aaditya Ramdas. Conformal prediction under covariate shift. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- Xin Tong, Yang Feng, and Jingyi Jessica Li. Neyman-Pearson classification algorithms and NP receiver operating characteristics. *Science advances*, 4(2): eao1659, 2018.
- Vladimir Vovk, Alex Gammerman, and Glenn Shafer. *Algorithmic learning in a random world*. Springer Science & Business Media, 2005.

---

## Supplementary Material: Learning Pareto-Efficient Decisions with Confidence

---

### A NUMERICAL EXPERIMENTS

More details of the numerical experiments described in Section 4 is provided in this material. In Algorithm 1 the computation of the Pareto-efficient decisions  $\mathcal{X}^\alpha(\mathbf{z})$  and the corresponding frontier  $\mathcal{Y}^\alpha(\mathbf{z})$  are summarized. A quantile function algorithm  $\hat{q}$  and a weight function type  $\hat{w}$  to be learned from data should be provided before running the algorithm. For the weights, one could for example use (13).

---

#### Algorithm 1 Pareto-efficient decisions

---

**Input:** Test sample  $\mathbf{z}$ , confidence level  $1 - \alpha$ , training data  $\mathcal{D}$

```

1: Randomly split  $\mathcal{D}$  into  $\mathcal{D}'$  and  $\mathcal{D}''$ 
2: for every  $x \in \mathcal{X}$  do ▷ Calculate lower bound of each reward
3:   Set  $x^* := x$  and  $\mathbf{y}^\alpha(x, \mathbf{z}) := \mathbf{0}$ 
4:   Learn weights  $\hat{w}(x, \mathbf{z}; \mathcal{D}')$ 
5:   Compute  $\hat{w}(x_i, \mathbf{z}_i; \mathcal{D}')$  for all  $i \in \mathcal{D}''$ 
6:   Compute  $\hat{w}(x, \mathbf{z}; \mathcal{D}')$  for test sample  $\mathbf{z}$ 
7:   for each reward  $k = 1, \dots, m$  do
8:     Learn quantile function  $\hat{q}^\alpha(x, \mathbf{z}; \mathcal{D}')$ 
9:     Compute  $r_i = \hat{q}^\alpha(x_i, \mathbf{z}_i; \mathcal{D}') - y_{k,i}$  for all  $i \in \mathcal{D}''$ 
10:    Construct  $p(r) = \sum_{i \in \mathcal{D}''} p_i \mathbb{1}(r = r_i) + p_\infty \mathbb{1}(r = \infty)$ , where  $\{p_i\}$  are given in (9)
11:    Set  $\kappa^\alpha(x)$  to the  $(1 - \frac{\alpha}{m})$ -th quantile of  $p(r)$ 
12:    Set  $y_k^\alpha(x, \mathbf{z}) := \hat{q}^\alpha(x, \mathbf{z}; \mathcal{D}') - \kappa^\alpha(x)$ 
13:  end for
14: end for
15: for every  $x \in \mathcal{X}$  do ▷ Find Pareto-efficient decisions
16:   if  $\mathbf{y}^\alpha(x, \mathbf{z})$  is Pareto-efficient then
17:     Add  $x$  to the set of all Pareto-efficient decisions  $\mathcal{X}^\alpha(\mathbf{z})$ 
18:     Add  $\mathbf{y}^\alpha(x, \mathbf{z})$  to the frontier  $\mathcal{Y}^\alpha(\mathbf{z})$ 
19:   end if
20: end for
Output:  $\mathcal{X}^\alpha(\mathbf{z})$  and  $\mathcal{Y}^\alpha(\mathbf{z})$ 

```

---

#### A.1 Synthetic data

In Subsection 4.1, we draw  $n = 1000$  training data points to learn  $\mathcal{X}^\alpha(z)$  with confidence level  $1 - \alpha$  using a quantile random forest method (Meinshausen and Ridgeway, 2006). The random forest has 1000 trees and the minimum number of samples in each leaf node is set to 1. In the main paper, we show results to verify that the method satisfies (2). Each individual  $y_k$  should also satisfy (11), with confidence level  $1 - \alpha/m$ , here  $m = 2$ . In Fig. 6 we verify that each reward  $y_k$  in the reward vector  $\mathbf{y}$  satisfies this in the experimental data case.

In the observational data case, we learn a Gaussian mixture model of  $p(z|x)$  and (11) will only hold approximately. In Fig. 7 we evaluate each reward boundary obtained with these approximate weights and find (11) to be satisfied for both. The bounds for  $x = 0$  and 1 are closest to the requested confidence  $1 - \alpha/2$ . Note that data on decision  $x = 4$  is sparse due to the uneven past decision policy. The resulting bound for this decision still satisfies (11), but is typically more conservative than for the other decisions.




Figure 6: Evaluation of reward boundary  $y_k^\alpha(x, z)$  of random rewards  $y_k$  using 500 Monte Carlo simulations. Experimental data scenario with known random assignment probabilities  $p(x)$ . Boundaries are designed to satisfy (11). (a) Probability of  $y_1$  exceeding the boundary  $y_1^\alpha(x, z)$  under decision  $x$ . (b) Probability of  $y_2$  exceeding the reward boundary  $y_2^\alpha(x, z)$  under decision  $x$ . Note that the dots are overlapping in both figures.

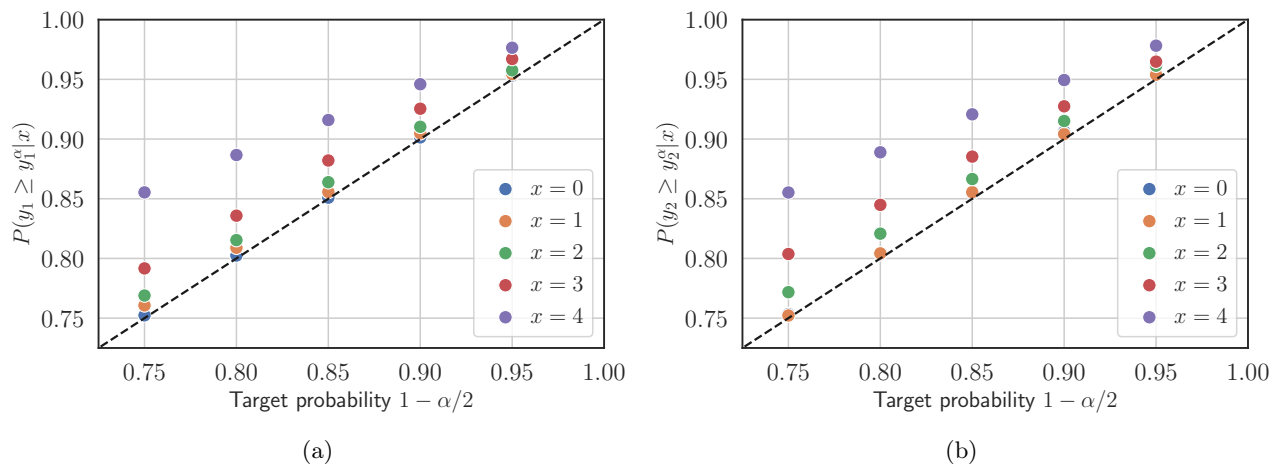


Figure 7: Evaluation of reward boundary  $y_k^\alpha(x, z)$  of random rewards  $y_k$  using 500 Monte Carlo simulations. Observational data scenario where  $p(z|x)$  in (13) is estimated. (a) Probability of  $y_1$  exceeding the boundary  $y_1^\alpha(x, z)$  under decision  $x$ . (b) Probability of  $y_2$  exceeding the reward boundary  $y_2^\alpha(x, z)$  under decision  $x$ . Note that several dots for  $x = 0$  and  $x = 1$  overlap in the figures.

## A.2 Semi-real data

For the experiments with real data in Subsection 4.2 we use 11 pre-intervention covariates, for more details see Table 2. The categorical covariates use one-hot encoding and the other covariates are scaled to zero mean and unit variance in the quantile neural net.

The neural net consists of three fully connected layers. The first layer takes the covariate vector as input. We then use 16 nodes in each of the hidden layers, and two nodes in the last layer. We are here only interested in a lower bound on the rewards and only one of the outputs are actually used in the calculation of  $y_k^\alpha(x, \mathbf{z})$ . Rectified linear unit (ReLU) is used as the activation function for the hidden layers and no activation function is used for the last layer. The weights are initialized in each layer with a (semi) orthogonal matrix (Saxe et al., 2014) and all biases are initialized to 0.

To fit the parameters of the network the pinball loss is minimized (Steinwart and Christmann, 2011). The Adam optimizer (Kingma and Ba, 2014) is used as the stochastic optimization algorithm. For the optimizer a batch size of 64, a fixed learning rate of 0.05 and weight decay  $10^{-3}$  is used. The network is trained for maximum 100 epochs, the actual number is tuned by cross validation (20 percent of the data is held out).

Table 2: The covariates used in the experiments with real data.

Variable name	Description	Categorical
gender	Student gender	yes
race	Student race/ethnicity	yes
g1surban	School urbanicity	yes
g1promote	Recommended for promotion from grade 1 to grade 2	yes
g1specin	Received special instruction	yes
g1freelunch	Free or reduced lunch given	yes
birthmonth	Student birth month	no
g1present	Days present at school	no
g1absent	Days absent from school	no
g1tcareer	Teachers career level	no
g1tyears	Teachers total teaching experience in years	no

The statistical guarantee in (11) for the boundary of the synthetic reward  $y_2$  is validated in Fig. 8. The method is tested in 500 experiments, where the test and training set come from 5-fold cross-validation repeated 100 times. We find that the reward  $y_2$  satisfies (11).

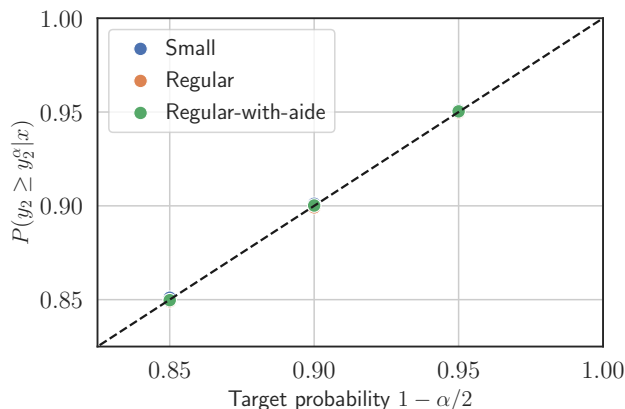


Figure 8: Evaluation of reward boundary  $y_2^\alpha(x, \mathbf{z})$  of random reward  $y_2$  under decision  $x$  using 500 Monte Carlo simulations. Experimental data scenario with known random policy  $p(x)$ . Note that several dots overlap in the figure.