
Robust Data-driven Prescriptiveness Optimization

Mehran Poursoltani¹ Erick Delage² Angelos Georghiou³

Abstract

The abundance of data has led to the emergence of a variety of optimization techniques that attempt to leverage available side information to provide more anticipative decisions. The wide range of methods and contexts of application have motivated the design of a universal unitless measure of performance known as the coefficient of prescriptiveness. This coefficient was designed to quantify both the quality of contextual decisions compared to a reference one and the prescriptive power of side information. To identify policies that maximize the former in a data-driven context, this paper introduces a distributionally robust contextual optimization model where the coefficient of prescriptiveness substitutes for the classical empirical risk minimization objective. We present a bisection algorithm to solve this model, which relies on solving a series of linear programs when the distributional ambiguity set has an appropriate nested form and polyhedral structure. Studying a contextual shortest path problem, we evaluate the robustness of the resulting policies against alternative methods when the out-of-sample dataset is subject to varying amounts of distribution shift.

1. Introduction

Stochastic programming is perceived as one of the fundamental methods devised for decision-making under uncertainty (see Shapiro et al. 2021 and Birge & Louveaux 2011). Given a cost function $h(\mathbf{x}, \boldsymbol{\xi})$ that depends on a decision $\mathbf{x} \in \mathbb{R}^{n_x}$ and a random vector $\boldsymbol{\xi} \in \mathbb{R}^{n_\xi}$, the stochastic

¹Desautels Faculty of Management, McGill University, Montréal, Canada ²GERAD & Department of Decision Sciences, HEC Montréal, Montréal, Canada ³Department of Business and Public Administration, University of Cyprus, Nicosia, Cyprus. Correspondence to: Mehran Poursoltani <mehran.poursoltani@mcgill.ca>.

Proceedings of the 41st International Conference on Machine Learning, Vienna, Austria. PMLR 235, 2024. Copyright 2024 by the author(s).

programming (SP) problem is defined as

$$(SP) \quad \mathbf{x}^* \in \arg \min_{\mathbf{x} \in \mathcal{X}} \mathbb{E}_F[h(\mathbf{x}, \boldsymbol{\xi})], \quad (1)$$

where \mathcal{X} is a convex feasible set, $h(\mathbf{x}, \boldsymbol{\xi})$ is a cost function that is assumed convex in \mathbf{x} for all $\boldsymbol{\xi}$, and $\boldsymbol{\xi}$ is assumed to be drawn from the distribution F . The solution methods for this problem mainly rely on either assuming a priori distribution for F or exploiting a set of independent and identically distributed observations. In the latter case, a set of i.i.d observations of the random vector $\boldsymbol{\xi}$ denoted by $\mathcal{S} := \{\boldsymbol{\xi}_i\}_{i=1}^N$ can be used to formulate the following sample average approximation problem:

$$(SAA) \quad \mathbf{x}^* \in \arg \min_{\mathbf{x} \in \mathcal{X}} \frac{1}{N} \sum_{i=1}^N h(\mathbf{x}, \boldsymbol{\xi}_i), \quad (2)$$

where we assume a uniform distribution over the observed data. Recently, the availability of large datasets has played a critical role in redirecting the optimization methods devised for decision-making under uncertainty towards taking advantage of so-called “side information” or “covariates”. This paradigm encourages decision-makers to benefit from the available data beyond the desired random variables to make more anticipative decisions. For instance, a portfolio manager who optimizes her investments in the stock market may consider a variety of available micro and macroeconomic indicators as side information to make more anticipative decisions (see Brandt et al. 2009 and Bazier-Matte & Delage 2020), while a traffic path planner can utilize side information like time of day, weather status and holiday/work day to find the best route through the city (see Bertsimas & Kallus 2020). This gives rise to the following contextual stochastic optimization (CSO) problem:

$$(CSO) \quad \mathbf{x}^*(\boldsymbol{\zeta}) \in \arg \min_{\mathbf{x} \in \mathcal{X}} \mathbb{E}_F[h(\mathbf{x}, \boldsymbol{\xi})|\boldsymbol{\zeta}], \quad (3)$$

where $\boldsymbol{\zeta} \in \mathbb{R}^{n_\zeta}$ denotes the given vector of “covariates”, or so-called “features”. In this case, any observed random vector $\boldsymbol{\xi}_i$ is accompanied by a vector of covariates $\boldsymbol{\zeta}_i \in \mathbb{R}^{n_\zeta}$. The difficulty of this problem shows up when the conditional probability distribution function $F_{\boldsymbol{\xi}|\boldsymbol{\zeta}}$ is unknown, and only a set of i.i.d observations $\mathcal{T} := \{(\boldsymbol{\zeta}_i, \boldsymbol{\xi}_i)\}_{i=1}^N$ is available. In this case, a data-driven variant of the CSO problem can be written as

$$\mathbf{x}^*(\boldsymbol{\zeta}) \in \arg \min_{\mathbf{x} \in \mathcal{X}} \mathbb{E}_{\hat{F}_{\boldsymbol{\xi}|\boldsymbol{\zeta}}} [h(\mathbf{x}, \boldsymbol{\xi})], \quad (4)$$

where $\hat{F}_{\xi|\zeta}$ is a conditional probability model for ξ given ζ inferred from the available data, e.g. by training a random forest (Breiman, 2001), or estimated via kernel density estimation (Ban & Rudin, 2019). To deal with possible overfitting in the presence of limited data or possible distribution shifts due to unexpected events, one can formulate a distributionally robust contextual stochastic optimization (DRCSO) model, which in general, takes the following form

$$(DRCSO) \quad \mathbf{x}^*(\zeta) \in \arg \min_{\mathbf{x} \in \mathcal{X}} \sup_{F \in \mathcal{D}} \mathbb{E}_F[h(\mathbf{x}, \xi)|\zeta] \quad (5)$$

where \mathcal{D} is the ambiguity set containing the set of admissible distributions (see Duchi et al. 2020; Bertsimas & Van Parys 2022; Kannan et al. 2020; Nguyen et al. 2022; Esteban-Pérez & Morales 2022; Srivastava et al. 2021, and literature within).

Recently, Bertsimas & Kallus 2020 proposed to compare the performance of different CSO (or DRCSO) approaches, by measuring the ‘‘coefficient of prescriptiveness’’, defined as:

$$\mathcal{P}_F(\mathbf{x}(\cdot)) := 1 - \frac{\mathbb{E}_F[h(\mathbf{x}(\zeta), \xi)] - \mathbb{E}_F[\min_{\mathbf{x}' \in \mathcal{X}} h(\mathbf{x}', \xi)]}{\mathbb{E}_F[h(\hat{\mathbf{x}}, \xi)] - \mathbb{E}_F[\min_{\mathbf{x}' \in \mathcal{X}} h(\mathbf{x}', \xi)]}, \quad (6)$$

where $\hat{\mathbf{x}} := \operatorname{argmin}_{\mathbf{x}} \mathbb{E}_{\hat{F}}[h(\mathbf{x}, \xi)]$ with \hat{F} as the empirical distribution that puts equal weights on each observed data point $\{\xi_i\}_{i=1}^N$ (i.e. the solution of SAA). The idea behind the coefficient of prescriptiveness is that it measures the performance of a given policy $\mathbf{x}(\zeta)$ relative to the constant decision $\hat{\mathbf{x}}$ which is agnostic to the side information ζ , and to the fully anticipative policy which achieves the progressive optimal value of $\mathbb{E}_F[\min_{\mathbf{x}' \in \mathcal{X}} h(\mathbf{x}', \xi)]$. It is easy to see that a high value of \mathcal{P}_F indicates that the policy can leverage the contextual information of ζ with $\mathcal{P}_F = 1$ indicating that the policy is achieving the fully anticipative performance in terms of ξ . In contrast, a low value of \mathcal{P}_F indicates that the policy is not able to exploit (or even is misled by) the available information. This behavior is reminiscent of R^2 , the ‘‘coefficient of determination’’, typically used in the context of predictive models, a connection which we discuss in the next section.

Following the introduction of the coefficient of prescriptiveness, this metric has been employed in several pieces of research to demonstrate the potential of proposed data-driven policies for leveraging the available side information. One can refer to Bertsimas et al. 2016 for such a comparison in the context of inventory management, Stratigakos et al. 2022 for energy trading, Notz & Pibernik 2022 for flexible capacity planning, and Kallus & Mao 2023 for shortest path and portfolio optimization problems. We note that, in the current literature, \mathcal{P}_F is only used as a benchmark metric for assessing the performance of policies computed using different approaches, e.g., in Bertsimas & Kallus 2020, the metric compares policies computed (amongst others) using CSO

where the conditional probability is estimated by random forests and kernel density estimation. Given its prevalence as a performance measure, it is natural to question whether it is possible and useful to directly optimize the coefficient of prescriptiveness.

While one can show that maximizing \mathcal{P}_F reduces to solving the CSO problem, one may wonder how the \mathcal{P}_F measure should be robustified in order to improve out-of-sample performance. In this work, we introduce for the first time a distributionally robust version of \mathcal{P}_F . We establish connections to other models in the literature and present an efficient algorithm to maximize it when the conditional probability model is discrete (such as with a random forest or with a Kernel density estimator). The rest of the paper is organized as follows. Section 2 motivates the optimization of the coefficient of prescriptiveness by explicating its relationship to the coefficient of determination in the field of statistics. Section 3 introduces a robust data-driven prescriptiveness optimization model that can be used to maximize a distributionally robust version of the coefficient of prescriptiveness. We reformulate this problem as a convex optimization problem that can reduce to a linear program when the ambiguity set takes the form of a so-called ‘‘nested Conditional Value-at-Risk (CVaR) set’’. A bisection method is proposed to solve the latter, as well as an acceleration scheme; finally, Section 4 presents the numerical experiments, where we evaluate the robustness of the resulting policies against benchmark ones in a shortest path problem when the out-of-sample dataset confronts a distribution shift. All proofs are relegated in Appendix A.

2. Motivation for optimizing \mathcal{P} and its robustification

As argued in Bertsimas & Kallus 2020, in the context of predictive models, where one wishes to predict the value of $\xi \in \mathbb{R}$ based on a list of covariates ζ using a statistical model $f : \mathbb{R}^{n_\zeta} \rightarrow \mathbb{R}$, one popular metric that is employed takes the form of the so-called ‘‘coefficient of determination’’:

$$R^2(f(\cdot)) := 1 - \frac{\mathbb{E}_{\hat{F}}[(f(\zeta) - \xi)^2]}{\mathbb{E}_{\hat{F}}[(\hat{\xi} - \xi)^2]},$$

where $\hat{\xi} := \mathbb{E}_{\hat{F}}[\xi]$ is the empirical mean of ξ in the data set and \hat{F} is the empirical joint distribution of (ζ, ξ) . The popularity of R^2 compared to mean squared error as a measure of performance can be partially attributed to being unitless. It is upper bounded by 1, with a value closer to 1, indicating that most of the variation of ξ can be modeled using $f(\cdot)$. On the flip side, when strictly smaller than 0, its absolute value measures the percentage of additional variations that are introduced by the predictive model, thus indicating a degradation of predictive power when compared to the simple sample average $\hat{\xi}$.

The coefficient of prescriptiveness can be viewed as an attempt to introduce an analogous measure in the contextual optimization setting. More specifically, it reduces to R^2 when $n_x = 1$ and $h(x, \xi) := (x - \xi)^2$, namely:

$$\begin{aligned} \mathcal{P}_{\hat{F}}(x(\cdot)) &= 1 - \frac{\mathbb{E}_{\hat{F}}[(x(\zeta) - \xi)^2] - \mathbb{E}_{\hat{F}}[\min_{x'}(x' - \xi)^2]}{\mathbb{E}_{\hat{F}}[(\hat{x} - \xi)^2] - \mathbb{E}_{\hat{F}}[\min_{x'}(x' - \xi)^2]} \\ &= R^2(x(\cdot)), \end{aligned}$$

since $\mathbb{E}_{\hat{F}}[\min_{x'}(x' - \xi)^2] = 0$ and $\hat{x} := \operatorname{argmin}_{\mathbf{x}} \mathbb{E}_{\hat{F}}[(x - \xi)^2] = \hat{\xi}$. Hence, the coefficient of prescriptiveness has a similar interpretation as R^2 . Namely, \mathcal{P}_F is upper bounded by 1, and as it gets closer to 1, it indicates how successful the data-driven policy has been in closing the gap between the SAA solution that makes no use of covariate information and a hypothetical policy that would have access to full information about ξ .

One can also find traces in the literature of attempts to measure $R^2(x(\cdot))$ out-of-sample. Namely, [Campbell & Thompson 2008](#) studies whether excess stock return predictors can outperform historical averages in terms of out-of-sample explanatory power of such predictors. This measure can be captured using

$$R_F^2(f(\cdot), \hat{F}) := 1 - \frac{\mathbb{E}_F[(f(\zeta) - \xi)^2]}{\mathbb{E}_F[(\hat{\xi} - \xi)^2]} = \mathcal{P}_F(f(\cdot)),$$

which naturally leads to the question of whether $R^2(f(\cdot))$ is a good approximation for $R_F^2(f(\cdot), \hat{F})$ in a data-driven environment (potentially susceptible to distribution shifts). If not, then one must turn to employing more robust estimation methods.

3. Robust Data-driven Prescriptiveness Optimization

In order to tackle the robustification and optimization of \mathcal{P} , we consider a more general version of this measure, which relaxes the assumption that the benchmark is the solution to (2) and widens the scope of our analysis. To this end, we define the prescriptiveness competitive ratio (PCR) of a policy $\mathbf{x}(\cdot)$ with respect to a reference policy $\bar{\mathbf{x}}$ as:

$$\mathcal{V}_F(\mathbf{x}(\cdot), \bar{\mathbf{x}}) := \begin{cases} 1 - \frac{\mathbb{E}_F[h(\mathbf{x}(\zeta), \xi)] - \mathbb{E}_F[\min_{\mathbf{x}' \in \mathcal{X}} h(\mathbf{x}', \xi)]}{\mathbb{E}_F[h(\bar{\mathbf{x}}, \xi)] - \mathbb{E}_F[\min_{\mathbf{x}' \in \mathcal{X}} h(\mathbf{x}', \xi)]} & \text{if } \mathbb{E}_F[h(\bar{\mathbf{x}}, \xi)] - \mathbb{E}_F[\min_{\mathbf{x}' \in \mathcal{X}} h(\mathbf{x}', \xi)] > 0 \\ 1 & \text{if } \mathbb{E}_F[h(\bar{\mathbf{x}}, \xi)] = \mathbb{E}_F[\min_{\mathbf{x}' \in \mathcal{X}} h(\mathbf{x}', \xi)] \\ & \text{and } \mathbb{E}_F[h(\mathbf{x}(\zeta), \xi)] = \mathbb{E}_F[\min_{\mathbf{x}' \in \mathcal{X}} h(\mathbf{x}', \xi)] \\ -\infty & \text{otherwise} \end{cases} . \quad (7)$$

Indeed, the coefficient of prescriptiveness can be considered a special case when $\bar{\mathbf{x}} := \hat{\mathbf{x}}$:

$$\begin{aligned} \mathcal{V}_F(\mathbf{x}(\cdot), \hat{\mathbf{x}}) &= 1 - \frac{\mathbb{E}_F[h(\mathbf{x}(\zeta), \xi)] - \mathbb{E}_F[\min_{\mathbf{x}' \in \mathcal{X}} h(\mathbf{x}', \xi)]}{\mathbb{E}_F[h(\hat{\mathbf{x}}, \xi)] - \mathbb{E}_F[\min_{\mathbf{x}' \in \mathcal{X}} h(\mathbf{x}', \xi)]} \\ &= \mathcal{P}_F(\mathbf{x}(\cdot)). \end{aligned}$$

when $\mathbb{E}_F[h(\hat{\mathbf{x}}, \xi)] - \mathbb{E}_F[\min_{\mathbf{x}' \in \mathcal{X}} h(\mathbf{x}', \xi)] > 0$, while the two other cases follow from the natural extension of the definition of $\mathcal{P}_F(\mathbf{x}(\cdot))$. In contrast to $\mathcal{P}_F(\mathbf{x}(\cdot))$ which benchmarks policy $\mathbf{x}(\cdot)$ only to the SAA solution, the definition of \mathcal{V}_F allows to benchmark against any other static policy. This allows our model to accommodate situations where more sophisticated statistical tools might be used to obtain the reference decision (e.g. regularized or distributionally robust SAA approaches ([Lam, 2019](#); [Mohajerin Esfahani & Kuhn, 2018](#); [Van Parys et al., 2021](#)), variance-based regularized solution schemes ([Duchi et al., 2020](#)), or data-pooled solutions schemes ([Gupta & Kallus, 2022](#))).¹

In a finite sample regime, where \hat{F} might fail to capture the true underlying distribution, or in a situation where we expect distribution shifts, one should be interested in a distributionally robust estimation of the PCR (or equivalently of the coefficient of prescriptiveness), which takes the form of:

$$\begin{aligned} \mathcal{V}_{\mathcal{D}}(\mathbf{x}(\cdot), \bar{\mathbf{x}}) &:= \inf_{F \in \mathcal{D}} \mathcal{V}_F(\mathbf{x}(\cdot), \bar{\mathbf{x}}) = \inf_{F \in \mathcal{D}} \mathcal{P}_F(\mathbf{x}(\cdot)) \\ &\text{when } \bar{\mathbf{x}} := \hat{\mathbf{x}}. \end{aligned}$$

PCR where \mathcal{D} is a set of distribution over the joint space (ζ, ξ) , and the notation $\mathcal{V}_{\mathcal{D}}$ is overloaded to denote the distributional robust PCR measure. Furthermore, one might be interested in identifying the policy that maximizes the PCR in the form of the following distributionally robust optimization problem:

$$(DRPCR) \quad \max_{\mathbf{x}(\cdot) \in \mathcal{H}} \mathcal{V}_{\mathcal{D}}(\mathbf{x}(\cdot), \bar{\mathbf{x}})$$

where $\mathcal{H} \subseteq \{\mathbf{x} : \mathbb{R}^{n_\zeta} \rightarrow \mathcal{X}\}$. The following lemma provides interpretable bounds for the value of $\mathcal{V}_{\mathcal{D}}$.

Lemma 3.1. *If $\bar{\mathbf{x}} \in \mathcal{H}$, then the optimal value of DRPCR is necessarily in the interval $[0, 1]$.*

Lemma 3.1 can be interpreted as follows. First, if $\mathbf{x}(\zeta)$ achieves a $\mathcal{V}_{\mathcal{D}}(\mathbf{x}(\cdot), \bar{\mathbf{x}}) = 1$ then the policy is guaranteed to exploit ζ just as efficiently as if it had full information about ξ (namely achieves the fully anticipative performance). On the other end of the spectrum, $\mathcal{V}_{\mathcal{D}}(\mathbf{x}(\cdot), \bar{\mathbf{x}}) = 0$ indicates that the policy can potentially fail to exploit any of the

¹In fact, one can go a step further and define $\mathcal{V}_F(\mathbf{x}(\cdot), \bar{\mathbf{x}}(\cdot))$ where $\bar{\mathbf{x}}(\cdot)$ is not a static policy. For example, $\bar{\mathbf{x}}(\cdot)$ could be a simple rule-based policy such as the order-up-to policy in inventory control. For ease of exposition, we treat the benchmark policy $\bar{\mathbf{x}}$ as a static policy for the remainder of the paper.

information present in ζ . When $\bar{x} \in \mathcal{H}$, one can always prevent negative PCR by falling back to the benchmark policy \bar{x} .

Next, we show that in an environment where the distribution is known, the optimal policy obtained from CSO is an optimal solution to DRPCR. Before proceeding, we first make the following assumption.

Assumption 3.2. The policy set \mathcal{H} contains all possible mappings, i.e. $\mathcal{H} := \{x : \mathbb{R}^{n_\zeta} \rightarrow \mathcal{X}\}$.

Lemma 3.3. *Given that Assumption 3.2 is satisfied, if the distribution set is a singleton, i.e. $\mathcal{D} = \{\bar{F}\}$, then the optimal policy obtained from the CSO problem that employs \bar{F} maximizes DRPCR.*

While Lemma 3.3 implies that DRPCR reduces to CSO when the distribution is known thus making the question of PCR optimization and performance irrelevant, this is not the case anymore for larger ambiguity sets \mathcal{D} .

In this section, we first present a convex reformulation of DRPCR and then provide a reformulation of the problem for the nested CVaR ambiguity set. Finally, we propose a decomposition algorithm for solving the problem based on a bisection algorithm.

3.1. Convex formulation for DRPCR

The following proposition provides a convex reformulation of DRPCR.

Proposition 3.4. *Given that $\bar{x} \in \mathcal{H}$, DRPCR is equivalent to*

$$\max_{x(\cdot) \in \mathcal{H}, \gamma} \quad \gamma \quad (8a)$$

$$\text{s.t.} \quad Q(x(\cdot), \gamma) \leq 0 \quad (8b)$$

$$0 \leq \gamma \leq 1. \quad (8c)$$

where

$$Q(x(\cdot), \gamma) := \sup_{F \in \mathcal{D}} \mathbb{E}_F \left[h(x(\zeta), \xi) - \left((1 - \gamma)h(\bar{x}, \xi) + \gamma \min_{x' \in \mathcal{X}} h(x', \xi) \right) \right]$$

is a convex non-decreasing function of γ . Moreover, problem (8) is a convex optimization problem when \mathcal{H} is convex.

From the reformulation (8) one can draw interesting insights regarding the connection of DRPCR and risk-averse regret minimization, see Poursoltani et al. 2023. For $\gamma = 1$, the problem reduces to the ex-post risk-averse regret minimization problem. In contrast, for $\gamma = 0$, one can interpret the problem as regretting the performance of the policy compared to a policy with less information. In the notation of Poursoltani et al. 2023, this will lead to a risk-averse regret problem with $\Delta = -1$.

3.2. The nested CVaR ambiguity set \mathcal{D}

In the following, we consider a discrete empirical distribution \bar{F} and restrict \mathcal{D} to be a nested CVaR ambiguity set. This ambiguity set is motivated by the works on nested dynamic risk measures (see Riedel 2004, Detlefsen & Scandolo 2005 and Ruszczyński & Shapiro 2006) as will be explained shortly. We formalize our approach through the following assumption.

Assumption 3.5. There is a discrete distribution \bar{F} , with $\{\zeta_\omega\}_{\omega \in \Omega_\zeta}$ and $\{\xi_\omega\}_{\omega \in \Omega_\xi}$ as the set of distinct scenarios for ζ and ξ respectively, such that the distribution set \mathcal{D} takes the form of the ‘‘nested CVaR ambiguity set’’ with respect to $\mathbb{P}_{\bar{F}}$ and defined as

$$\bar{\mathcal{D}}(\bar{F}, \alpha) := \left\{ F \in \mathcal{M}(\Omega_\zeta \times \Omega_\xi) \left| \begin{array}{l} \mathbb{P}_F(\zeta = \zeta_\omega) = \\ \mathbb{P}_{\bar{F}}(\zeta = \zeta_\omega) \quad \forall \omega \in \Omega_\zeta, \\ \mathbb{P}_F(\xi = \xi_{\omega'} | \zeta_\omega) \leq \\ (1/(1 - \alpha)) \mathbb{P}_{\bar{F}}(\xi = \xi_{\omega'} | \zeta_\omega) \\ \forall \omega \in \Omega_\zeta, \omega' \in \Omega_\xi \end{array} \right. \right\}. \quad (9)$$

where $\mathcal{M}(\Omega_\zeta \times \Omega_\xi)$ is the set of all distributions supported on over the joint space $\{\zeta_\omega\}_{\omega \in \Omega_\zeta} \times \{\xi_\omega\}_{\omega \in \Omega_\xi}$.

The structure of $\bar{\mathcal{D}}(\bar{F}, \alpha)$ implies that there is no ambiguity in the marginal distribution of the observed random variable ζ . Rather, the ambiguity is solely on the unobserved random variable ξ and is sized using the parameter α . The nested CVaR ambiguity set owes its name from Ruszczyński & Shapiro 2006 and the fact that for any function $g(x, \xi)$:

$$\begin{aligned} & \sup_{F \in \bar{\mathcal{D}}(\bar{F}, \alpha)} \mathbb{E}_F [g(x(\zeta), \xi)] \\ &= \sup_{F \in \bar{\mathcal{D}}(\bar{F}, \alpha)} \sum_{\omega \in \Omega_\zeta} \sum_{\omega' \in \Omega_\xi} \mathbb{P}_F(\zeta = \zeta_\omega) \cdot \\ & \quad \mathbb{P}_F(\xi = \xi_{\omega'} | \zeta = \zeta_\omega) g(x(\zeta_\omega), \xi_{\omega'}) \\ &= \mathbb{E}_{\bar{F}} [\text{CVaR}_{\bar{F}}^\alpha (g(x(\zeta), \xi) | \zeta)]. \end{aligned}$$

For $\alpha = 0$, the problem reduces to $\min_{x(\cdot) \in \mathcal{H}} \mathbb{E}_{\bar{F}} [h(x(\zeta), \xi)]$, effectively recovering the CSO policy. On the other spectrum, for $\alpha = 1$ the problem reduces to $\min_{x(\cdot) \in \mathcal{H}} \mathbb{E}_{\bar{F}} [\max_{\omega: \bar{F}(\xi = \xi_\omega | \zeta) > 0} h(x(\zeta), \xi_\omega)]$, which implies that for each realization of ζ_ω the decision $x(\zeta_\omega)$ is robust against all admissible realizations of ξ given ζ_ω .

The nested CVaR representation and full policy space Assumption 3.2 can be exploited to optimize $Q(x(\cdot), \gamma)$.

Proposition 3.6. *Under Assumption 3.5, problem (8) can*

thus be reformulated as

$$\max_{\gamma} \quad \gamma \quad (10a)$$

$$\text{s.t.} \quad \sum_{\omega \in \Omega_{\zeta}} \mathbb{P}_{\bar{F}}(\zeta = \zeta_{\omega}) \phi_{\omega}(\gamma) \leq 0 \quad (10b)$$

$$0 \leq \gamma \leq 1, \quad (10c)$$

where $\phi_{\omega}(\gamma)$ is a non-decreasing function (when $\bar{x} \in \mathcal{X}$) capturing the optimal value of:

$$\min_{\mathbf{x} \in \mathcal{X}, t, s \geq 0} \quad t + \frac{1}{1-\alpha} \sum_{\omega' \in \Omega_{\xi}} \mathbb{P}_{\bar{F}}(\xi = \xi_{\omega'} | \zeta = \zeta_{\omega}) s_{\omega'} \quad (11a)$$

$$\text{s.t.} \quad s_{\omega'} \geq h(\mathbf{x}, \xi_{\omega'}) - \left((1-\gamma)h(\bar{\mathbf{x}}, \xi_{\omega'}) + \gamma \min_{\mathbf{x}' \in \mathcal{X}} h(\mathbf{x}', \xi_{\omega'}) \right) - t, \quad \forall \omega' \in \Omega_{\xi} \quad (11b)$$

and can be reduced to a linear program when \mathcal{X} is polyhedral and $h(\mathbf{x}, \xi_{\omega'})$ is linear programming representable.

In practice, \bar{F} is often composed of an empirical distribution \hat{F}_{ζ} and a trained conditional distribution $\hat{F}_{\xi|\zeta}$. Given an optimal solution γ^* to problem (13), one should then define the extended optimal policy $\mathbf{x}(\zeta)$ beyond $\{\zeta_{\omega}\}_{\omega \in \Omega_{\zeta}}$ using the optimal solution of problem (11) with $\hat{F}_{\xi|\zeta}$.

This being said, whether problem (11) is reduceable to a linear program or, more generally, a convex optimization model, its size scales with $|\Omega_{\zeta}| \cdot |\Omega_{\xi}|$, which can be computationally challenging. We therefore propose a decomposition algorithm to efficiently solve the problem. Let $\psi(\gamma) := \sum_{\omega \in \Omega_{\zeta}} \mathbb{P}_{\bar{F}}(\zeta = \zeta_{\omega}) \phi_{\omega}(\gamma)$. Using the definition of $\phi_{\omega}(\gamma)$, we observe that for fixed γ one can evaluate $\psi(\gamma)$ by solving $|\Omega_{\zeta}|$ distinct problem (11) for each $\omega \in \Omega_{\zeta}$. Moreover, given that each $\phi_{\omega}(\gamma)$ is non-decreasing (see Proposition 3.6), one concludes that $\psi(\gamma)$ is non-decreasing. Hence, one can design a bisection algorithm on γ to solve the DRPCR problem (8). Namely, each step consists in identifying the mid-point $\tilde{\gamma}$ of an interval known to contain the optimal value of γ , and verifying whether $\tilde{\gamma}$ is feasible by evaluating $\psi(\tilde{\gamma})$ to decide which of the two sub-interval below or above $\tilde{\gamma}$ contains γ^* , see Figure 2 (left) in Appendix B. The details of this algorithm are presented in Algorithm 1. It's efficiency relies on the difficulty of executing step 7, i.e. evaluation $\phi_{\omega}(\gamma)$ for each ω . The following lemma provides formal guaranties regarding the convergence rate of Algorithm 1.

Lemma 3.7. *Algorithm 1 terminates in $\lceil \log_2(1/\epsilon) \rceil$ iterations. Moreover, if \mathcal{X} is polyhedral and $h(\mathbf{x}, \xi)$ linear programming representable, the algorithm terminates in polynomial time with respect to $\log(1/\epsilon)$, $|\Omega_{\zeta}|$, $|\Omega_{\xi}|$, n_x , n_{ξ} , the size of the LP representation of \mathcal{X} and of $h(\mathbf{x}, \xi)$.*

Appendix B further proposes an accelerated bisection algorithm for the case when \mathcal{X} is convex. Namely, it derives the sub-gradient of $\psi(\gamma)$ and exploits its convexity to tighten the interval for γ^* at each iteration.

Algorithm 1 Bisection algorithm for DRPCR

```

1: Input: Tolerance  $\epsilon > 0$ 
2: Set  $\gamma^- := 0, \gamma^+ := 1$ 
3: while  $\gamma^+ - \gamma^- > \epsilon$  do
4:   Set  $\tilde{\gamma} := (\gamma^+ + \gamma^-)/2$ 
5:   //Solve  $\min_{\mathbf{x}(\cdot) \in \mathcal{H}} Q(\mathbf{x}(\cdot), \tilde{\gamma})$  to get optimal value  $\tilde{\psi}(\tilde{\gamma})$ 
6:   for  $\omega \in \Omega_{\zeta}$  do
7:     Solve problem (11) with  $\omega$  and  $\tilde{\gamma}$  to get optimal value  $\phi_{\omega}(\tilde{\gamma})$ 
8:   end for
9:   if  $\tilde{\psi} := \sum_{\omega \in \Omega_{\zeta}} \mathbb{P}_{\bar{F}}(\zeta = \zeta_{\omega}) \phi_{\omega}(\tilde{\gamma}) \leq 0$  then
10:    Set  $\gamma^- := \tilde{\gamma}$ 
11:  else
12:    Set  $\gamma^+ := \tilde{\gamma}$ 
13:  end if
14: end while
15: Return  $\gamma^* := \gamma^-$ 
    
```

3.3. Generalized nested ambiguity set \mathcal{D}

One can generalize the results of the previous section by considering a generalized version of the ambiguity set formalized in the following assumption.

Assumption 3.8. For a discrete distribution \bar{F} , the distribution set \mathcal{D} takes the form of the ‘‘generalized nested ambiguity set’’ with respect to $\mathbb{P}_{\bar{F}}$ and defined as

$$\bar{\mathcal{D}}(\bar{F}, r_{\zeta}, r_{\xi}) := \left\{ F \in \mathcal{M}(\Omega_{\zeta} \times \Omega_{\xi}) \mid \begin{array}{l} d_{\zeta}(F_{\zeta}, \bar{F}_{\zeta}) \leq r_{\zeta} \\ d_{\xi}(F_{\xi|\zeta_{\omega}}, \bar{F}_{\xi|\zeta_{\omega}}) \leq r_{\xi} \quad \forall \omega \in \Omega_{\zeta} \end{array} \right\}. \quad (12)$$

where $d_{\zeta}(F_{\zeta}, \bar{F}_{\zeta})$ and $d_{\xi}(F_{\xi|\zeta_{\omega}}, \bar{F}_{\xi|\zeta_{\omega}})$ are two convex divergence measures, i.e. non-negative, convex in their first argument and minimized when the two probability measures are equal, applied on marginal distribution of F and the conditional distribution of ξ given ζ , respectively.

The structure of (12) allows one to control the ambiguity about both the marginal distribution of ζ and the conditional distributions of ξ using the parameters r_{ζ} and r_{ξ} to bound the maximum divergence respectively. In particular, it reduces to the nested CVaR ambiguity set when using

$$d_{\zeta}(F_{\zeta}, \bar{F}_{\zeta}) := \inf \left\{ s \mid \mathbb{P}_{F_{\zeta}}(\zeta = \zeta_{\omega}) \leq \frac{1}{1-s} \mathbb{P}_{\bar{F}_{\zeta}}(\zeta = \zeta_{\omega}), \forall \omega \in \Omega \right\},$$

$$d_\xi(F_{\xi|\zeta}, \bar{F}_{\xi|\zeta}) := \inf \left\{ s \mid \mathbb{P}_{F_{\xi|\zeta}}(\boldsymbol{\xi} = \boldsymbol{\xi}_{\omega'}) \leq \frac{1}{1-s} \mathbb{P}_{\bar{F}_{\xi|\zeta}}(\boldsymbol{\xi} = \boldsymbol{\xi}_{\omega'}), \forall \omega' \in \Omega' \right\},$$

$$r_\zeta := 0, \text{ and } r_\xi := \alpha.$$

In the following, to simplify presentation, given that Ω_ζ and Ω_ξ are finite, we let $\mathbf{p} \in \mathbb{R}^{|\Omega_\zeta|}$ denote the vector of probabilities $p_\omega := \mathbb{P}_F(\zeta = \zeta_\omega)$ and $\mathbf{q}^\omega \in \mathbb{R}^{|\Omega_\xi|}$ denote the probabilities $q_{\omega'}^\omega := \mathbb{P}_F(\xi = \xi_{\omega'} | \zeta = \zeta_\omega)$, and similarly for $\bar{\mathbf{p}}$ and $\bar{\mathbf{q}}^\omega$ to capture the same probabilities under \bar{F} . We will further abuse notation and denote $d_\zeta(\mathbf{p}, \bar{\mathbf{p}}) := d_\zeta(F_\zeta, \bar{F}_\zeta)$ and $d_\xi(\mathbf{q}^\omega, \bar{\mathbf{q}}^\omega) := d_\xi(F_{\xi|\zeta_\omega}, \bar{F}_{\xi|\zeta_\omega})$. The following proposition generalizes Proposition 3.6.

Proposition 3.9. *Under Assumption 3.8, problem (8) can thus be reformulated as*

$$\max_{\gamma} \quad \gamma \quad (13a)$$

$$\text{s.t.} \quad \sup_{\mathbf{p} \in \mathcal{Z}} \sum_{\omega \in \Omega_\zeta} p_\omega \bar{\phi}_\omega(\gamma) \leq 0 \quad (13b)$$

$$0 \leq \gamma \leq 1, \quad (13c)$$

where $\mathcal{Z} := \{\mathbf{p} : \mathbf{p} \geq 0, \mathbf{e}^\top \mathbf{p} = 1, d_\zeta(\mathbf{p}, \bar{\mathbf{p}}) \leq r_\zeta\}$ and $\bar{\phi}_\omega(\gamma)$ is a non-decreasing function (when $\bar{\mathbf{x}} \in \mathcal{X}$) capturing the optimal value of:

$$\min_{\mathbf{x}, t, \alpha, \mathbf{s}} \quad t + r_\xi \alpha + d_*(\mathbf{s}, \alpha, \bar{\mathbf{q}}^\omega) \quad (14a)$$

$$\text{s.t.} \quad s_{\omega'} \geq h(\mathbf{x}, \boldsymbol{\xi}_{\omega'}) - \left((1-\gamma)h(\bar{\mathbf{x}}, \boldsymbol{\xi}_{\omega'}) + \gamma \min_{\mathbf{x}' \in \mathcal{X}} h(\mathbf{x}', \boldsymbol{\xi}_{\omega'}) \right) - t, \forall \omega' \in \Omega_\xi \quad (14b)$$

$$\mathbf{x} \in \mathcal{X}, \alpha \geq 0, \quad (14c)$$

where $d_*(\mathbf{s}, \alpha, \bar{\mathbf{q}}^\omega) := \sup_{\mathbf{q}} \mathbf{s}^\top \mathbf{q} - \alpha d_\xi(\mathbf{q}, \bar{\mathbf{q}}^\omega)$ is the perspective of the convex conjugate of $d_\xi(\mathbf{q}, \bar{\mathbf{q}}^\omega)$.

Algorithm 1 can be applied in the generalized setting with the simple modification that problem (11) in step 7 is replaced with the convex problem (14), and step 9 must compute $\sup_{\mathbf{p} \in \mathcal{Z}} \sum_{\omega \in \Omega_\zeta} p_\omega \bar{\phi}_\omega(\gamma)$, which now requires solving a convex optimization problem.

4. Experiments

In this section, we present a numerical study that compares the performance of DRPCR against three other data-driven benchmark methods to evaluate its robustness to perturbations of the data generating process. Specifically, we will observe how these models react to the situation where one faces a distribution shift for $\boldsymbol{\xi}$. In a vehicle routing problem with travel time uncertainties, this can be interpreted as a shift in the distribution of the travel times, for instance, when a special event is happening in the town. Alternatively, one can think of an inventory management problem

where the manager faces a shift in the demand distribution, e.g., an unforeseen increase in demand for sanitizer during the first days of an epidemic. In general, there are numerous reasons why distribution shifts considerations might be needed depending on the context. In this regard, we refer the reader to Schrouff et al. 2022 and Filos et al. 2020 for such considerations in healthcare and autonomous driving applications.

The application that we consider for our numerical experiments is a shortest path problem described in Kallus & Mao 2023. A directed graph is defined as $\mathcal{G} = (\mathcal{V}, \mathcal{A})$, where \mathcal{V} denotes the set of nodes and $\mathcal{A} \in \mathcal{V} \times \mathcal{V}$ is the set of arcs, i.e., ordered pairs (i, j) of nodes describing the existence of a directed path from node i to node j . The corresponding travel time of such an arc is assumed to be $\xi_{(i,j)}$. The objective of this problem is to identify the shortest path from an origin (node o) to a destination (node d). Moving away from an ideal world of known parameters gives rise to a stochastic version of this problem. In this setting, the traveling times along the arcs $\boldsymbol{\xi} \in \mathbb{R}^{|\mathcal{A}|}$ are uncertain; however, one might still have access to side information or observed covariates. In this case, aiming at minimizing the expected travel time leads to the following CSO problem:

$$\mathbf{x}^*(\boldsymbol{\zeta}) \in \arg \min_{\mathbf{x} \in \mathcal{X}} \mathbb{E}_{\hat{F}_{\xi|\zeta}}[\mathbf{x}^\top \boldsymbol{\xi}], \quad (15)$$

where

$$\mathcal{X} = \left\{ \mathbf{x} \in \mathbb{R}^{|\mathcal{A}|} \left| \begin{array}{l} x_{(i,j)} \in \{0, 1\} \quad \forall (i, j) \in \mathcal{A} \\ \sum_{j:(o,j) \in \mathcal{A}} x_{(o,j)} - \sum_{j:(j,o) \in \mathcal{A}} x_{(j,o)} = 1 \\ \sum_{j:(d,j) \in \mathcal{A}} x_{(d,j)} - \sum_{j:(j,d) \in \mathcal{A}} x_{(j,d)} = -1 \\ \sum_{j:(i,j) \in \mathcal{A}} x_{(i,j)} - \sum_{j:(j,i) \in \mathcal{A}} x_{(j,i)} = 0 \quad \forall i \in \mathcal{V} \setminus \{o, d\} \end{array} \right. \right\},$$

and $x_{(i,j)} = 1$ if we decide to travel from node i to node j and $x_{(i,j)} = 0$ otherwise. Unlike Kallus & Mao 2023, we enforce the integrality constraints. Furthermore, $\hat{F}_{\xi|\zeta}$ denotes the conditional distribution inferred from the training dataset.

As discussed in Section 1, DRCSO is a method proposed for robustifying the policies against distributional uncertainties in the data-driven context. Consequently, one can consider DRCSO, as an alternative to CSO, for solving this shortest-path problem. Using the nested CVaR ambiguity set introduced in Assumption 3.5 as the ambiguity set of DRCSO, one gets the model below:

$$(DRCSO) \quad \mathbf{x}^*(\boldsymbol{\zeta}) \in \arg \min_{\mathbf{x} \in \mathcal{X}} \sup_{F_{\xi|\zeta} \in \bar{\mathcal{D}}(\hat{F}_{\xi|\zeta}, \alpha)} \mathbb{E}_{F_{\xi|\zeta}}[\mathbf{x}^\top \boldsymbol{\xi}], \quad (16)$$

where

$$\bar{\mathcal{D}}(\hat{F}_{\xi|\zeta}, \alpha) := \{F_{\xi|\zeta} \in \mathcal{M}(\Omega_{\xi}) : \mathbb{P}_{F_{\xi|\zeta}}(\xi = \xi_{\omega'}) \leq (1/(1-\alpha))\mathbb{P}_{\hat{F}_{\xi|\zeta}}(\xi = \xi_{\omega'}) \forall \omega' \in \Omega_{\xi}\},$$

and α is the control parameter for the size of the ambiguity set. Staying in the DRCSO context, one can exploit a worst-case regret minimization approach instead of worst-case expected travel time. In our experiments, we look into the optimal solutions arising from an ex-post regret minimization setting, introduced as a $\Delta = 1$ regret minimization model in Poursoltani et al. 2023. This leads to the following distributionally robust contextual regret optimization (DRCRO) problem:

$$(DRCRO) \quad \mathbf{x}^*(\zeta) \in \arg \min_{\mathbf{x} \in \mathcal{X}} \sup_{F_{\xi|\zeta} \in \bar{\mathcal{D}}(\hat{F}_{\xi|\zeta}, \alpha)} \mathbb{E}_{F_{\xi|\zeta}}[\mathbf{x}^{\top} \xi - \min_{\mathbf{x}' \in \mathcal{X}} \mathbf{x}'^{\top} \xi]. \quad (17)$$

In this case, the decision maker compares her travel time to the one resulting from a benchmark decision that knows the future realization of ξ . The ultimate goal is to minimize the worst-case expectation of this gap, so-called ‘‘worst-case expected regret’’, where the ambiguity set is nested CVaR. Finally, we solve our introduced DRPCR problem under nested CVaR ambiguity set, where the $Q(\mathbf{x}(\cdot), \gamma)$ function takes the form of:

$$Q(\mathbf{x}(\cdot), \gamma) := \sup_{F \in \bar{\mathcal{D}}(\tilde{F}, \alpha)} \mathbb{E}_F \left[\mathbf{x}(\zeta)^{\top} \xi - \left((1-\gamma)\hat{\mathbf{x}}^{\top} \xi + \gamma \min_{\mathbf{x}' \in \mathcal{X}} \mathbf{x}'^{\top} \xi \right) \right], \quad (18)$$

where \tilde{F} denotes the distribution derived from the training dataset, composed of the empirical distribution \hat{F}_{ζ} of ζ and the inferred conditional distribution $\hat{F}_{\xi|\zeta}$, while $\hat{\mathbf{x}} := \arg \min_{\mathbf{x}} \mathbb{E}_{\tilde{F}}[h(\mathbf{x}, \xi)]$ with \tilde{F} that puts equal weights on each observed data point $\{\xi_i\}_{i=1}^N$ (i.e. the SAA solution). Based on an optimal solution γ^* for the DRPCR problem, one can retrieve an optimal policy using:

$$\mathbf{x}^*(\zeta) \in \arg \min_{\mathbf{x} \in \mathcal{X}} \sup_{F_{\xi|\zeta} \in \bar{\mathcal{D}}(\hat{F}_{\xi|\zeta}, \alpha)} \mathbb{E}_{F_{\xi|\zeta}} \left[\mathbf{x}^{\top} \xi - \left((1-\gamma^*)\hat{\mathbf{x}}^{\top} \xi + \gamma^* \min_{\mathbf{x}' \in \mathcal{X}} \mathbf{x}'^{\top} \xi \right) \right], \quad (19)$$

which can be obtained by solving (11) with γ^* and replacing $\mathbb{P}_{\tilde{F}}(\xi = \xi_{\omega'} | \zeta = \zeta_{\omega})$ with $\mathbb{P}_{\hat{F}_{\xi|\zeta}}(\xi_{\omega'})$.

We adapt our numerical experiments to the graph (\mathcal{G}) structure employed in Kallus & Mao 2023 with the same origin (o) and destination (d); therefore, we study a graph with the size of 45 nodes ($|\mathcal{V}| = 45$) and 97 arcs ($|\mathcal{A}| = 97$). We assume there exist 200 covariates ($n_{\zeta} = 200$) and the vector composed of travel times ξ and covariates ζ follow a multivariate normal distribution. Specifically, each covariate

ζ_i follows a normal distribution with a mean of zero and standard deviation of one (i.e. $\zeta_i \sim \mathcal{N}(0, 1)$). Similarly, each travel time $\xi_{(i,j)}$ is normal with a standard deviation that matches the deviation present in Kallus & Mao 2023’s dataset yet both the correlation and mean vector are treated differently. Starting with correlation, we introduce a new correlation structure for $(\zeta, \xi)^2$ by instantiating a random correlation matrix (see Appendix D for details).

Our treatment of the mean of ξ embodies our objective to study robustness to distribution shifts. Namely, while the data generating process for the training set employs the same mean vector as in Kallus & Mao 2023, our validation data set and out-of-sample test set will measure the performance of proposed policies on generating processes where the mean of ξ as been perturbed, i.e. $\mathbb{E}[\xi_{(i,j)}] := (1 + \delta_{(i,j)})\mu_{(i,j)}$. Six tests were conducted for different levels of mean perturbations: no distribution shift $\delta_{(i,j)} = 0$, which does not allow for any perturbation, along with tests that take into account shifts with $\delta_{(i,j)}$ generated i.i.d. according to a uniform distribution on $[0\%, m]$, where $m \in \mathcal{M} := \{20\%, 30\%, 40\%, 50\%, 60\%$ represents the maximum possible perturbation. Furthermore, the perturbation experienced in the validation set is independent of the test set. This is to simulate situations where the level of robustness would be calibrated on a data set where a distribution shift of similar size is observed as the shift experienced out-of-sample.

Experiments for each perturbation range contain 50 instances generated by resampling the training, validation, and test data sets. Both the training and validation datasets consist of 400 data points, while the test set contains 1000 data points and is used to measure the ‘‘out-of-sample’’ performance. The training dataset is used for learning purposes, which allows us to infer the conditional probabilities of $\hat{F}_{\xi|\zeta}$ once a new covariate vector ζ is observed. From a wide range of existing predictive tools for inference of $\hat{F}_{\xi|\zeta}$, Bertsimas & Kallus 2020 compare methods such as k-nearest-neighbors regression (Hastie et al. 2001), local linear regressions (Cleveland & Devlin 1988), classification and regression trees (CART; Breiman et al. 1984), and random forests (RF; Breiman 2001). In their experiments, the best coefficient of prescriptiveness belongs to random forests. We exploit the code provided in Kallus & Mao 2023 to train random forests over our training datasets and then use it as the conditional distribution estimator $\hat{F}_{\xi|\zeta}$ for our validation and out-of-sample data points. The validation dataset is used to calibrate the size of the ambiguity set (α) for the DRCSO, DRCRO, and DRPCR models. The procedure for calibrating α and the associated optimal γ for the

²This was done after observing that with Kallus & Mao 2023’s dataset the optimal uninformed decisions produced nearly the same performance as the optimal hindsight decisions that exploited full information about realized travel times.




Figure 1. Shortest path problem: (a) statistics of the out-of-sample coefficient of prescriptiveness (lower values indicate worse performance). (b) statistics of $\mathbb{E}_{\tilde{F}}[\|\mathbf{x}^*(\zeta) - \hat{\mathbf{x}}\|_1]$ where \tilde{F} is the out-of-sample distribution (lower values reflect a closer proximity to the SAA solution).

DRPCR model and to calibrate α for the DRCSO and DRCRO models are described in Appendix C (see algorithms 2 and 3 respectively). We define the set of discretized α values as $\mathcal{A} := \mathcal{A}_1 \cap \mathcal{A}_2$, where \mathcal{A}_1 includes 20 logarithmically spaced values in $[0.01, 0.99]$ and \mathcal{A}_2 includes 20 evenly spaced values in $[0, 1)$. For CSO, Algorithm 3 can also be used with $\mathcal{A} = \{0\}$. From a computational point of view, the training of the DRPCR algorithm took on average less than 36 minutes per instance, compared to closer to 3 minutes for DRCSO and DRCRO. The main difference comes from the extra Step 6 of Algorithm 2, which requires solving the DRPCR problem for each candidate for α , and took on average 50 seconds to solve, and needs to be repeated for all $\alpha \in \mathcal{A}$. Once the optimal α^* and γ^* are determined by Algorithm 2 for a given training/validation dataset, equation (19) provides the optimal policy $x^*(\cdot)$ for any covariate ζ received in real-time. Similarly, equations (16) and (17) can be employed to derive the optimal DRCSO and DRCRO policies associated with the real-time input of ζ , relying on the calibrated values of α^* obtained from Algorithm 3. All optimization problems are implemented in Python and solved using Gurobi 8.1.1 on a machine featuring an Intel processor Xeon(R) CPU E5-2687W v3 @ 3.10GHz 3.10 GHz (2 processors) and 128 GB RAM. The code used for the numerical experiments is available at https://github.com/erickdelage/robust_prescriptive_opt.

Figure 1(a) reports the coefficients of prescriptiveness $\mathcal{V}_F(x^*(\cdot), \hat{\mathbf{x}})$, where F is the test dataset, for the four policies and perturbation levels. More details on the average out-of-sample performance are also presented in Table 2 in Appendix E. We observe the following: (i) When considering a particular optimization model, the coefficient of

prescriptiveness decreases as the magnitude of the distribution shift increases. Indeed, these policies face a more serious robustness challenge as they approach more extreme scenarios beyond what was seen in the train dataset. (ii) When the test set follows the same distribution as the train set, all four policies roughly demonstrate similar performance; however, when this set experiences a distribution shift, DRPCR policies differentiate their performance compared to the alternative ones. (iii) Imposing a more severe distribution shift accentuates this differentiation. For instance, when the mean travel times across the edges are perturbed up to 50% in the test set, DRPCR policies provide a positive coefficient of prescriptiveness, at least over 75% of instances. On the contrary, the alternative policies fail to reach a positive ratio over almost a similar number of instances. This observation is further amplified in the case of 60% perturbation. In this scenario, while CSO, DRCSO, and DRCRO policies fail to return a positive out-of-sample coefficient of prescriptiveness, DRPCR still can reach a positive median of 4% which can go up to 18% at its best.

Figure 1(b) depicts the statistics of the 1-norm distance metric of the adaptable policies from the optimal SAA solution. This illustration elucidates the reason behind the superior performance of the DRPCR method compared to others. Indeed, it is notable that DRPCR implicitly utilizes \bar{x} as an anchor for the adaptable policy. In other words, under large distribution shifts, it is able to learn under what context it is worth staying closer to \bar{x} , where the relative regret is zero, as explained below Lemma 3.1. This phenomenon highlights how the DRPCR applies a completely different form of regularization, compared to prior DRO models.

For further insights, readers are directed to Appendix F,

Table 1. Average runtime (minutes) per instance

Method	Type of Problem	Algorithm	Levels of Mean Perturbations (m)					
			0%	20%	30%	40%	50%	60%
DRC SO	relaxed $x(\cdot)$	3	3.08	2.38	3.06	3.14	2.38	2.38
DRC RO	relaxed $x(\cdot)$	3	2.96	2.28	2.96	3.02	2.28	2.28
DRPCR	relaxed $x(\cdot)$	2	32.76	25.22	32.86	33.58	25.28	25.22
DRC SO	binary $x(\cdot)$	3	3.08	2.64	3.04	3.42	3.02	3.40
DRC RO	binary $x(\cdot)$	3	3.32	3.42	3.20	3.44	3.20	3.44
DRPCR	binary $x(\cdot)$	2	34.28	38.08	33.32	37.92	33.14	37.92

where an additional set of experiments is presented. In line with the approach in Kallus & Mao 2023, the integrality constraint of $x(\cdot)$ is relaxed in this supplementary investigation. One can refer to Table 1 for a comparison of runtime of Algorithms 2 and 3 under both relaxed and binary policies.

5. Conclusion

The proposed DRPCR model offers an innovative method for calibrating contextual optimization problems and introduces a unique form of regularization, differing from previous DRO models and achieving significantly improved out-of-sample performance. Unfortunately, in its current form, the approach requires considerable training time, i.e. approximately tenfold that of the DRC SO and DRC RO models. This is due to the bisection algorithm needing $\lceil \log_2(1/\epsilon) \rceil$ steps to converge. In order to improve tractability (at the expense of optimality), one might consider limiting the admissible policy $x(\cdot)$ to those with affine dependence on the side information ζ . The resulting DRPCR takes the form of a smaller optimization problem with size proportional to the number of dimensions of ζ rather than $|\Omega_\zeta|$. Affine policies might also facilitate the use of more general non-nested ambiguity sets, which constitute a current limitation of the proposed DRPCR model. Finally, we expect that additional empirical evaluation with other data generating and application environments would certainly benefit our understanding of the value of the presented DRPCR approach.

Impact Statement

While we expect this work to contribute in the long term to many potential societal consequences, none are direct consequences of the results presented in this paper.

Acknowledgments

The authors gratefully acknowledge support from the Institut de Valorisation des Données (IVADO), from the Cana-

dian Natural Sciences and Engineering Research Council [RGPIN-2022-05261], and the Canada Research Chair program [CRC-2018-00105].

References

Ban, G.-Y. and Rudin, C. The big data newsvendor: Practical insights from machine learning. *Operations Research*, 67(1):90–108, 2019.

Bazier-Matte, T. and Delage, E. Generalization bounds for regularized portfolio selection with market side information. *INFOR: Information Systems and Operational Research*, 58(2):374–401, 2020.

Bertsimas, D. and Kallus, N. From predictive to prescriptive analytics. *Management Science*, 66(3):1025–1044, 2020.

Bertsimas, D. and Van Parys, B. Bootstrap robust prescriptive analytics. *Mathematical Programming*, 195(1-2):39–78, 2022.

Bertsimas, D., Kallus, N., and Hussain, A. Inventory management in the era of big data. *Production and Operations Management*, 25(12):2006–2009, 2016.

Birge, J. R. and Louveaux, F. *Introduction to stochastic programming*. Springer Science & Business Media, 2011.

Brandt, M. W., Santa-Clara, P., and Valkanov, R. Parametric Portfolio Policies: Exploiting Characteristics in the Cross-Section of Equity Returns. *The Review of Financial Studies*, 22(9):3411–3447, 02 2009.

Breiman, L. Random forests. *Machine learning*, 45:5–32, 2001.

Breiman, L., Friedman, J., Stone, C., and Olshen, R. Classification algorithms and regression trees. *Classification and regression trees*, 15(2):246, 1984.

Campbell, J. Y. and Thompson, S. B. Predicting excess stock returns out of sample: Can anything beat the historical average? *The Review of Financial Studies*, 21(4):1509–1531, 2008.

Cleveland, W. S. and Devlin, S. J. Locally weighted regression: an approach to regression analysis by local fitting. *Journal of the American statistical association*, 83(403):596–610, 1988.

Detlefsen, K. and Scandolo, G. Conditional and dynamic convex risk measures. *Finance and Stochastics*, 9(4):539–561, 2005.

- Duchi, J. C., Hashimoto, T., and Namkoong, H. Distributionally robust losses for latent covariate mixtures. *CoRR*, abs/2007.13982, 2020.
- Esteban-Pérez, A. and Morales, J. M. Distributionally robust stochastic programs with side information based on trimmings. *Mathematical Programming*, 195(1):1069–1105, 2022.
- Filos, A., Tigkas, P., McAllister, R., Rhinehart, N., Levine, S., and Gal, Y. Can autonomous vehicles identify, recover from, and adapt to distribution shifts? In *International Conference on Machine Learning*, pp. 3145–3153. PMLR, 2020.
- Grötschel, M., Lovász, L., and Schrijver, A. The ellipsoid method and its consequences in combinatorial optimization. *Combinatorica*, 1(2):169–197, 1981.
- Gupta, V. and Kallus, N. Data pooling in stochastic optimization. *Management Science*, 68(3):1595–1615, 2022.
- Hastie, T., Tibshirani, R., and Friedman, J. *The Elements of Statistical Learning*. Springer Series in Statistics. Springer New York Inc., New York, NY, USA, 2001.
- Kallus, N. and Mao, X. Stochastic optimization forests. *Management Science*, 69(4):1975–1994, 2023.
- Kannan, R., Bayraksan, G., and Luedtke, J. R. Residuals-based distributionally robust optimization with covariate information. *arXiv preprint arXiv:2012.01088*, 2020.
- Karmarkar, N. A new polynomial-time algorithm for linear programming. *Combinatorica*, 4(4):373–395, 1984.
- Lam, H. Recovering best statistical guarantees via the empirical divergence-based distributionally robust optimization. *Operations Research*, 67(4):1090–1105, 2019.
- Mohajerin Esfahani, P. and Kuhn, D. Data-driven distributionally robust optimization using the Wasserstein metric: Performance guarantees and tractable reformulations. *Mathematical Programming*, 171(1):115–166, 2018.
- Nguyen, V. A., Zhang, F., Blanchet, J., Delage, E., and Ye, Y. Robustifying conditional portfolio decisions via optimal transport. *arXiv preprint arXiv:2103.16451*, 2022.
- Notz, P. M. and Pibernik, R. Prescriptive analytics for flexible capacity management. *Management Science*, 68(3):1756–1775, 2022. doi: 10.1287/mnsc.2020.3867.
- Poursoltani, M., Delage, E., and Georghiou, A. Risk-averse regret minimization in multistage stochastic programs. *Operations Research*, <https://doi.org/10.1287/opre.2022.2429>, 2023.
- Riedel, F. Dynamic coherent risk measures. *Stochastic Processes and their Applications*, 112(2):185–200, 2004.
- Ruszczynski, A. and Shapiro, A. Conditional risk mappings. *Mathematics of Operations Research*, 31(3):544–561, 2006.
- Schrouff, J., Harris, N., Koyejo, O., Alabdulmohsin, I., Schneider, E., Opsahl-Ong, K., Brown, A., Roy, S., Mincu, D., Chen, C., et al. Maintaining fairness across distribution shift: do we have viable solutions for real-world applications? *arXiv preprint arXiv:2202.01034*, 2022.
- Shapiro, A. Interchangeability principle and dynamic equations in risk averse stochastic programming. *Operations Research Letters*, 45(4):377–381, 2017.
- Shapiro, A., Dentcheva, D., and Ruszczyński, A. *Lectures on stochastic programming: modeling and theory*. SIAM, 2021.
- Sion, M. On general minimax theorems. *Pacific Journal of Mathematics*, 8(1):171–176, 1958.
- Srivastava, P. R., Wang, Y., Hanasusanto, G. A., and Ho, C. P. On data-driven prescriptive analytics with side information: A regularized nadaraya-watson approach. *arXiv preprint arXiv:2110.04855*, 2021.
- Stratigakos, A., Camal, S., Michiorri, A., and Kariniotakis, G. Prescriptive trees for integrated forecasting and optimization applied in trading of renewable energy. *IEEE Transactions on Power Systems*, 37(6):4696–4708, 2022.
- Van Parys, B. P. G., Esfahani, P. M., and Kuhn, D. From data to decisions: Distributionally robust optimization is optimal. *Management Science*, 67(6):3387–3402, 2021.

A. Proofs

A.1. Proof of Lemma 3.1

This follows simply from $\mathcal{V}_F(\mathbf{x}(\cdot), \bar{\mathbf{x}})$ being bound above by 1 for all policy $\mathbf{x}(\cdot)$ and all distribution F due to:

$$\begin{aligned}\mathcal{V}_F(\mathbf{x}(\cdot), \bar{\mathbf{x}}) &= 1 - \frac{\mathbb{E}_F[h(\mathbf{x}(\zeta), \boldsymbol{\xi})] - \mathbb{E}_F[\min_{\mathbf{x}' \in \mathcal{X}} h(\mathbf{x}', \boldsymbol{\xi})]}{\mathbb{E}_F[h(\bar{\mathbf{x}}, \boldsymbol{\xi})] - \mathbb{E}_F[\min_{\mathbf{x}' \in \mathcal{X}} h(\mathbf{x}', \boldsymbol{\xi})]} \\ &\leq 1 - \frac{\mathbb{E}_F[\min_{\mathbf{x}' \in \mathcal{X}} h(\mathbf{x}', \boldsymbol{\xi})] - \mathbb{E}_F[\min_{\mathbf{x}' \in \mathcal{X}} h(\mathbf{x}', \boldsymbol{\xi})]}{\mathbb{E}_F[h(\bar{\mathbf{x}}, \boldsymbol{\xi})] - \mathbb{E}_F[\min_{\mathbf{x}' \in \mathcal{X}} h(\mathbf{x}', \boldsymbol{\xi})]} = 1\end{aligned}$$

when $\mathbb{E}_F[h(\hat{\mathbf{x}}, \boldsymbol{\xi})] - \mathbb{E}_F[\min_{\mathbf{x}' \in \mathcal{X}} h(\mathbf{x}', \boldsymbol{\xi})] > 0$, and otherwise equal to 1 or $-\infty$ both bounded above by 1. Hence,

$$\max_{\mathbf{x}(\cdot) \in \mathcal{H}} \inf_{F \in \mathcal{D}} \mathcal{V}_F(\mathbf{x}(\cdot), \bar{\mathbf{x}}) \leq 1.$$

Moreover, if $\bar{\mathbf{x}} \in \mathcal{H}$, then we have that

$$\max_{\mathbf{x}(\cdot) \in \mathcal{H}} \mathcal{V}_D(\mathbf{x}(\cdot), \bar{\mathbf{x}}) \geq \mathcal{V}_D(\bar{\mathbf{x}}, \bar{\mathbf{x}}) = \begin{cases} 0 & \text{if } \mathbb{E}_F[h(\bar{\mathbf{x}}, \boldsymbol{\xi})] - \mathbb{E}_F[\min_{\mathbf{x}' \in \mathcal{X}} h(\mathbf{x}', \boldsymbol{\xi})] > 0 \\ 1 & \text{otherwise.} \end{cases} \quad \square$$

A.2. Proof of Lemma 3.3

Let $\tilde{\mathbf{x}}(\cdot)$ be a CSO optimal policy, then necessarily $\tilde{\mathbf{x}}(\cdot) \in \mathcal{H}$ since $\tilde{\mathbf{x}}(\zeta) \in \mathcal{X}$ for all ζ . This confirms that $\tilde{\mathbf{x}}(\cdot)$ is feasible in DRPCR. Next, we can demonstrate optimality through:

$$\mathcal{V}_D(\tilde{\mathbf{x}}(\cdot), \bar{\mathbf{x}}) = \mathcal{V}_{\bar{F}}(\tilde{\mathbf{x}}(\cdot), \bar{\mathbf{x}}) \geq \max_{\mathbf{x}(\cdot) \in \mathcal{H}} \mathcal{V}_{\bar{F}}(\mathbf{x}(\cdot), \bar{\mathbf{x}}) = \max_{\mathbf{x}(\cdot) \in \mathcal{H}} \mathcal{V}_D(\mathbf{x}(\cdot), \bar{\mathbf{x}}),$$

since for all $\mathbf{x}(\cdot) \in \mathcal{H}$, we have that $\mathbb{E}_F[h(\mathbf{x}(\zeta), \boldsymbol{\xi})|\zeta] \geq \min_{\mathbf{x}(\cdot) \in \mathcal{H}} \mathbb{E}_F[h(\mathbf{x}(\zeta), \boldsymbol{\xi})|\zeta] = \mathbb{E}_F[h(\tilde{\mathbf{x}}(\zeta), \boldsymbol{\xi})|\zeta]$ for all ζ , which we can show implies that $\mathcal{V}_F(\tilde{\mathbf{x}}(\cdot), \bar{\mathbf{x}}) \geq \mathcal{V}_F(\mathbf{x}(\cdot), \bar{\mathbf{x}})$. More specifically, if $\mathbb{E}_F[h(\hat{\mathbf{x}}, \boldsymbol{\xi})] = \mathbb{E}_F[\min_{\mathbf{x}' \in \mathcal{X}} h(\mathbf{x}', \boldsymbol{\xi})]$, then either $\mathbb{E}_F[h(\mathbf{x}(\zeta), \boldsymbol{\xi})] = \mathbb{E}_F[\min_{\mathbf{x}' \in \mathcal{X}} h(\mathbf{x}', \boldsymbol{\xi})]$ thus

$$\mathbb{E}_F[\min_{\mathbf{x}' \in \mathcal{X}} h(\mathbf{x}', \boldsymbol{\xi})] = \mathbb{E}_F[h(\tilde{\mathbf{x}}(\zeta), \boldsymbol{\xi})] \leq \mathbb{E}_F[h(\mathbf{x}(\zeta), \boldsymbol{\xi})] = \mathbb{E}_F[\min_{\mathbf{x}' \in \mathcal{X}} h(\mathbf{x}', \boldsymbol{\xi})]$$

meaning that $\mathcal{V}_F(\tilde{\mathbf{x}}(\cdot), \bar{\mathbf{x}}) = \mathcal{V}_F(\mathbf{x}(\cdot), \bar{\mathbf{x}}) = 1$, or $\mathbb{E}_F[h(\mathbf{x}(\zeta), \boldsymbol{\xi})] > \mathbb{E}_F[\min_{\mathbf{x}' \in \mathcal{X}} h(\mathbf{x}', \boldsymbol{\xi})]$ thus $\mathcal{V}_F(\tilde{\mathbf{x}}(\cdot), \bar{\mathbf{x}}) \geq -\infty = \mathcal{V}_F(\mathbf{x}(\cdot), \bar{\mathbf{x}})$. Alternatively, the case where $\mathbb{E}_F[h(\hat{\mathbf{x}}, \boldsymbol{\xi})] > \mathbb{E}_F[\min_{\mathbf{x}' \in \mathcal{X}} h(\mathbf{x}', \boldsymbol{\xi})]$ is straightforward as the function

$$f(y) := 1 - \frac{y - \mathbb{E}_F[\min_{\mathbf{x}' \in \mathcal{X}} h(\mathbf{x}', \boldsymbol{\xi})]}{\mathbb{E}_F[h(\bar{\mathbf{x}}, \boldsymbol{\xi})] - \mathbb{E}_F[\min_{\mathbf{x}' \in \mathcal{X}} h(\mathbf{x}', \boldsymbol{\xi})]}$$

is strictly decreasing. □

A.3. Proof of Proposition 3.4

We first present the DRPCR in epigraph form:

$$\max_{\gamma, \mathbf{x}(\cdot) \in \mathcal{H}} \gamma \tag{20a}$$

$$\text{s.t. } \mathcal{V}_F(\mathbf{x}(\cdot), \bar{\mathbf{x}}) \geq \gamma, \quad \forall F \in \mathcal{D} \tag{20b}$$

$$0 \leq \gamma \leq 1 \tag{20c}$$

where we added the redundant constraint $\gamma \in [0, 1]$ since Lemma 3.1 ensures that the optimal value of DRPCR is in this interval.

Focusing on constraint (20b), we can then consider two cases for the definition of $\mathcal{V}_F(\mathbf{x}(\cdot), \bar{\mathbf{x}})$. In the case that $\mathbb{E}_F[h(\bar{\mathbf{x}}, \boldsymbol{\xi})] - \mathbb{E}_F[\min_{\mathbf{x}' \in \mathcal{X}} h(\mathbf{x}', \boldsymbol{\xi})] > 0$, one can multiply both sides of the inequality to equivalently obtain:

$$\mathbb{E}_F[h(\mathbf{x}(\zeta), \boldsymbol{\xi})] - \mathbb{E}_F[\min_{\mathbf{x}' \in \mathcal{X}} h(\mathbf{x}', \boldsymbol{\xi})] \leq (1 - \gamma) \left(\mathbb{E}_F[h(\bar{\mathbf{x}}, \boldsymbol{\xi})] - \mathbb{E}_F[\min_{\mathbf{x}' \in \mathcal{X}} h(\mathbf{x}', \boldsymbol{\xi})] \right)$$

which is equivalent, when rearranging the terms, to:

$$\mathbb{E}_F[h(\mathbf{x}(\zeta), \xi) - (1 - \gamma)h(\bar{\mathbf{x}}, \xi) - \gamma \min_{\mathbf{x}' \in \mathcal{X}} h(\mathbf{x}', \xi)] \leq 0. \quad (21)$$

In the second case where $\mathbb{E}_F[h(\bar{\mathbf{x}}, \xi)] = \mathbb{E}_F[\min_{\mathbf{x}' \in \mathcal{X}} h(\mathbf{x}', \xi)]$, then constraint (20b) is equivalent to:

$$\mathbb{E}_F[h(\mathbf{x}(\zeta), \xi)] = \mathbb{E}_F[\min_{\mathbf{x}' \in \mathcal{X}} h(\mathbf{x}', \xi)] \quad \& \quad \gamma \leq 1,$$

yet $\gamma \leq 1$ is redundant while the former condition can equivalently be posed as (21). We are left with

$$\mathbb{E}_F[h(\mathbf{x}(\zeta), \xi) - (1 - \gamma)h(\bar{\mathbf{x}}, \xi) - \gamma \min_{\mathbf{x}' \in \mathcal{X}} h(\mathbf{x}', \xi)] \leq 0, \quad \forall F \in \mathcal{D},$$

which can equivalently be described by $Q(\mathbf{x}(\cdot), \gamma) \leq 0$. One can further conclude that $Q(\mathbf{x}(\cdot), \gamma) \leq 0$ is convex and non-decreasing in γ given that it is the supremum of a set of affine non-decreasing functions:

$$Q(\mathbf{x}(\cdot), \gamma) = \sup_{F \in \mathcal{D}} \mathbb{E}_F[h(\mathbf{x}(\zeta), \xi) - h(\bar{\mathbf{x}}, \xi)] + \gamma(\mathbb{E}_F[h(\bar{\mathbf{x}}, \xi) - \min_{\mathbf{x}' \in \mathcal{X}} h(\mathbf{x}', \xi)]),$$

with $h(\bar{\mathbf{x}}, \xi) \geq \min_{\mathbf{x}' \in \mathcal{X}} h(\mathbf{x}', \xi)$ for all ξ since $\bar{\mathbf{x}} \in \mathcal{X}$. □

A.4. Proof of Proposition 3.6

Letting $g(\mathbf{x}, \xi, \gamma) := h(\mathbf{x}, \xi) - ((1 - \gamma)h(\bar{\mathbf{x}}, \xi) + \gamma \min_{\mathbf{x}' \in \mathcal{X}} h(\mathbf{x}', \xi))$, we have that

$$\begin{aligned} \psi(\gamma) &:= \min_{\mathbf{x}(\cdot) \in \mathcal{H}} Q(\mathbf{x}(\cdot), \gamma) \\ &= \min_{\mathbf{x}(\cdot) \in \mathcal{H}} \sup_{F \in \bar{\mathcal{D}}(\bar{F}, \alpha)} \mathbb{E}_F[g(\mathbf{x}(\zeta), \xi, \gamma)] \\ &= \min_{\mathbf{x}(\cdot) \in \mathcal{H}} \mathbb{E}_{\bar{F}} \left[\text{CVaR}_{\bar{F}}^\alpha \left(g(\mathbf{x}(\zeta), \xi, \gamma) \mid \zeta \right) \right] \\ &= \min_{\mathbf{x}(\cdot) \in \mathcal{H}} \mathbb{E}_{\bar{F}} \left[\inf_t t + \frac{1}{1 - \alpha} \mathbb{E}_{\bar{F}} \left[\max \left(0, g(\mathbf{x}(\zeta), \xi, \gamma) - t \right) \mid \zeta \right] \right] \\ &= \mathbb{E}_{\bar{F}} \left[\inf_{\mathbf{x} \in \mathcal{X}, t} t + \frac{1}{1 - \alpha} \mathbb{E}_{\bar{F}} \left[\max \left(0, g(\mathbf{x}, \xi, \gamma) - t \right) \mid \zeta \right] \right], \end{aligned}$$

where we exploit the infimum representation of CVaR and the interchangeability property of expected value operators (see Shapiro 2017 and reference therein). Given that \bar{F} is a discrete distribution as described in Assumption 3.5, one can compute $\psi(\gamma)$ by solving for each scenario ζ_ω with $\omega \in \Omega_\zeta$ the problem (11b). Based on the solution of problem (11) for each $\omega \in \Omega_\zeta$, one can obtain $\psi(\gamma) := \sum_{\omega \in \Omega_\zeta} \mathbb{P}_{\bar{F}}(\zeta = \zeta_\omega) \phi_\omega(\gamma)$ together with a potentially feasible policy $\mathbf{x}(\zeta) := \mathbf{x}_{\omega(\zeta)}^*$, where $\omega(\zeta) = \operatorname{argmin}_{\omega \in \Omega_\zeta} \|\zeta - \zeta_\omega\|$ and \mathbf{x}_ω refers to the minimizer of problem (11).

The function $\phi_\omega(\gamma)$ is non-decreasing in γ since γ only appears in constraint (11b), which can be rewritten as:

$$s_{\omega'} \geq h(\mathbf{x}, \xi_{\omega'}) - h(\bar{\mathbf{x}}, \xi_{\omega'}) + \left(h(\bar{\mathbf{x}}, \xi_{\omega'}) - \min_{\mathbf{x}' \in \mathcal{X}} h(\mathbf{x}', \xi_{\omega'}) \right) \gamma - t, \quad \forall \omega' \in \Omega_\xi.$$

Since the right-hand side of this constraint is non-decreasing in γ , due to $h(\bar{\mathbf{x}}, \xi_{\omega'}) \geq \min_{\mathbf{x}' \in \mathcal{X}} h(\mathbf{x}', \xi_{\omega'})$, one can conclude that the minimum of (11) cannot decrease when γ is increased, since the feasible set is reduced.

We further note that problem (11) can be reduced to a linear program when \mathcal{X} is polyhedral and $h(\mathbf{x}, \xi_{\omega'})$ is linear programming representable for all $\omega' \in \Omega_\xi$. For example, in the context of a portfolio optimization, where \mathcal{X} is the

probability simplex and $h(\mathbf{x}, \boldsymbol{\xi}) := -\boldsymbol{\xi}^T \mathbf{x}$, we have that problem (8) reduces to:

$$\begin{aligned}
 & \max_{\gamma, \{\mathbf{x}^\omega, t^\omega, \mathbf{s}^\omega\}_{\omega \in \Omega_\zeta}} \quad \gamma \\
 & \text{subject to} \quad \sum_{\omega \in \Omega_\zeta} \mathbb{P}_{\bar{F}}(\zeta = \zeta_\omega) \left(t^\omega + \frac{1}{1-\alpha} \sum_{\omega' \in \Omega_\xi} \mathbb{P}_{\bar{F}}(\boldsymbol{\xi} = \boldsymbol{\xi}_{\omega'} | \zeta = \zeta_\omega) s_{\omega'}^\omega \right) \leq 0 \\
 & \quad s_{\omega'}^\omega \geq \boldsymbol{\xi}_{\omega'}^T \mathbf{x}^\omega - (1-\gamma) \boldsymbol{\xi}_{\omega'}^T \bar{\mathbf{x}} - \gamma \min_{\mathbf{x}' \in \mathcal{X}} \boldsymbol{\xi}_{\omega'}^T \mathbf{x}' - t^\omega, \quad \forall \omega' \in \Omega_\xi, \omega \in \Omega_\zeta \\
 & \quad s^\omega \geq 0, \quad \forall \omega \in \Omega_\zeta \\
 & \quad \sum_{i=1}^{n_x} \mathbf{x}_i^\omega = 1, \quad \forall \omega \in \Omega_\zeta \\
 & \quad \mathbf{x}^\omega \geq 0, \quad \forall \omega \in \Omega_\zeta \\
 & \quad 0 \leq \gamma \leq 1.
 \end{aligned}$$

Alternatively, in the context of a shortest path problem (see Section 4 for details), we have that problem (8) reduces to a mixed integer linear program:

$$\begin{aligned}
 & \max_{\gamma, \{\mathbf{x}^\omega, t^\omega, \mathbf{s}^\omega\}_{\omega \in \Omega_\zeta}} \quad \gamma \\
 & \text{subject to} \quad \sum_{\omega \in \Omega_\zeta} \mathbb{P}_{\bar{F}}(\zeta = \zeta_\omega) \left(t^\omega + \frac{1}{1-\alpha} \sum_{\omega' \in \Omega_\xi} \mathbb{P}_{\bar{F}}(\boldsymbol{\xi} = \boldsymbol{\xi}_{\omega'} | \zeta = \zeta_\omega) s_{\omega'}^\omega \right) \leq 0 \\
 & \quad s_{\omega'}^\omega \geq \boldsymbol{\xi}_{\omega'}^T \mathbf{x}^\omega - (1-\gamma) \boldsymbol{\xi}_{\omega'}^T \bar{\mathbf{x}} - \gamma \min_{\mathbf{x}' \in \mathcal{X}} \boldsymbol{\xi}_{\omega'}^T \mathbf{x}' - t^\omega, \quad \forall \omega' \in \Omega_\xi, \omega \in \Omega_\zeta \\
 & \quad s^\omega \geq 0, \quad \forall \omega \in \Omega_\zeta \\
 & \quad \sum_{j:(o,j) \in \mathcal{A}} x_{(o,j)}^\omega - \sum_{j:(j,o) \in \mathcal{A}} x_{(j,o)}^\omega = 1, \quad \forall \omega \in \Omega_\zeta \\
 & \quad \sum_{j:(d,j) \in \mathcal{A}} x_{(d,j)}^\omega - \sum_{j:(j,d) \in \mathcal{A}} x_{(j,d)}^\omega = -1, \quad \forall \omega \in \Omega_\zeta \\
 & \quad \sum_{j:(i,j) \in \mathcal{A}} x_{(i,j)}^\omega - \sum_{j:(j,i) \in \mathcal{A}} x_{(j,i)}^\omega = 0, \quad \forall i \in \mathcal{V} \setminus \{o, d\}, \omega \in \Omega_\zeta \\
 & \quad \mathbf{x}_{(i,j)}^\omega \in \{0, 1\}, \quad \forall (i, j) \in \mathcal{A}, \omega \in \Omega_\zeta \\
 & \quad 0 \leq \gamma \leq 1. \quad \square
 \end{aligned}$$

A.5. Proof of Lemma 3.7

One can first easily verify that in Algorithm 1, we have that $\Delta := \gamma^+ - \gamma^-$ is initially equal to 1 and reduces by a factor of 2 at every iteration. The algorithm therefore necessarily terminates after $\lceil \log_2(1/\epsilon) \rceil$ iterations. When \mathcal{X} is polyhedral and $h(\mathbf{x}, \boldsymbol{\xi})$ is linear programming representable, problem (13) reduces to a linear program that can be solved in polynomial time with respect to $|\Omega_\xi|$, n_x , n_ξ , the size of the LP representation of \mathcal{X} , and of $h(\mathbf{x}, \boldsymbol{\xi})$ (see Grötschel et al. 1981 and Karmarkar 1984). Given that this problem is solved $|\Omega_\zeta|$ at each iteration of the algorithm. We conclude that the total run time of the algorithm is polynomial with respect to all of these quantities. \square

A.6. Proof of Proposition 3.9

The proof follows similar steps as Proposition 3.6. Let $g(\mathbf{x}, \boldsymbol{\xi}, \gamma) := h(\mathbf{x}, \boldsymbol{\xi}) - ((1 - \gamma)h(\bar{\mathbf{x}}, \boldsymbol{\xi}) + \gamma \min_{\mathbf{x}' \in \mathcal{X}} h(\mathbf{x}', \boldsymbol{\xi}))$, we have that

$$\begin{aligned}
 \psi(\gamma) &= \min_{\mathbf{x}(\cdot) \in \mathcal{H}} \sup_{F \in \bar{\mathcal{D}}(F, \alpha)} \mathbb{E}_F \left[g(\mathbf{x}(\zeta), \boldsymbol{\xi}, \gamma) \right] \\
 &= \min_{\mathbf{x}(\cdot) \in \mathcal{H}} \sup_{\mathbf{p}: \mathbf{p} \geq 0, \sum_{\omega \in \Omega_\zeta} p_\omega = 1, d_\zeta(\mathbf{p}, \bar{\mathbf{p}}) \leq r_\zeta} \sum_{\omega \in \Omega_\zeta} p_\omega \sup_{\mathbf{q}: \mathbf{q} \geq 0, \sum_{\omega' \in \Omega_\xi} q_{\omega'} = 1, d_\xi(\mathbf{q}, \bar{\mathbf{q}}^\omega) \leq r_\xi} \sum_{\omega' \in \Omega_\xi} q_{\omega'} g(\mathbf{x}(\zeta_\omega), \boldsymbol{\xi}_{\omega'}, \gamma) \\
 &= \sup_{\mathbf{p}: \mathbf{p} \geq 0, \sum_{\omega \in \Omega_\zeta} p_\omega = 1, d_\zeta(\mathbf{p}, \bar{\mathbf{p}}) \leq r_\zeta} \sum_{\omega \in \Omega_\zeta} p_\omega \min_{\mathbf{x} \in \mathcal{X}} \sup_{\mathbf{q}: \mathbf{q} \geq 0, \sum_{\omega' \in \Omega_\xi} q_{\omega'} = 1, d_\xi(\mathbf{q}, \bar{\mathbf{q}}^\omega) \leq r_\xi} \sum_{\omega' \in \Omega_\xi} q_{\omega'} g(\mathbf{x}(\zeta_\omega), \boldsymbol{\xi}_{\omega'}, \gamma) \\
 &= \sup_{\mathbf{p}: \mathbf{p} \geq 0, \sum_{\omega \in \Omega_\zeta} p_\omega = 1, d_\zeta(\mathbf{p}, \bar{\mathbf{p}}) \leq r_\zeta} \sum_{\omega \in \Omega_\zeta} p_\omega \bar{\phi}_\omega(\gamma), \tag{22}
 \end{aligned}$$

with

$$\bar{\phi}_\omega(\gamma) := \min_{\mathbf{x} \in \mathcal{X}} \sup_{\mathbf{q}: \mathbf{q} \geq 0, \sum_{\omega' \in \Omega_\xi} q_{\omega'} = 1, d_\xi(\mathbf{q}, \bar{\mathbf{q}}^\omega) \leq r_\xi} \sum_{\omega' \in \Omega_\xi} q_{\omega'} g(\mathbf{x}(\zeta_\omega), \boldsymbol{\xi}_{\omega'}, \gamma). \tag{23}$$

Denoting $\mathbf{g} \in \mathbb{R}^{|\Omega'|}$ with $g_{\omega'} := g(\mathbf{x}(\zeta_\omega), \boldsymbol{\xi}_{\omega'}, \gamma)$, one can derive a reformulation of the inner supremum in (23) as an infimum following:

$$\begin{aligned}
 \sup_{\mathbf{q}: \mathbf{q} \geq 0, \sum_{\omega' \in \Omega_\xi} q_{\omega'} = 1, d_\xi(\mathbf{q}, \bar{\mathbf{q}}^\omega) \leq r_\xi} \mathbf{q}^T \mathbf{g} &= \sup_{\mathbf{q}} \inf_{\lambda \geq 0, t, \alpha \geq 0} \mathbf{q}^T \mathbf{g} + \boldsymbol{\lambda}^T \mathbf{q} + t(1 - \mathbf{1}^T \mathbf{q}) + \alpha(r_\xi - d_\xi(\mathbf{q}, \bar{\mathbf{q}}^\omega)) \\
 &= \inf_{\lambda \geq 0, t, \alpha \geq 0} \sup_{\mathbf{q}} \mathbf{q}^T \mathbf{g} + \boldsymbol{\lambda}^T \mathbf{q} + t(1 - \mathbf{1}^T \mathbf{q}) + \alpha(r_\xi - d_\xi(\mathbf{q}, \bar{\mathbf{q}}^\omega)) \\
 &= \inf_{\lambda \geq 0, t, \alpha \geq 0} t + r_\xi \alpha + d_*(\mathbf{g} + \boldsymbol{\lambda} - t, \alpha, \bar{\mathbf{q}}^\omega),
 \end{aligned}$$

where $d_*(\mathbf{v}, \alpha, \bar{\mathbf{q}}^\omega) := \sup_{\mathbf{q}} \mathbf{v}^T \mathbf{q} - \alpha d_\xi(\mathbf{q}, \bar{\mathbf{q}}^\omega)$ is the perspective of the convex conjugate of $d_\xi(\mathbf{q}, \bar{\mathbf{q}}^\omega)$.

By joining this reformulation with the minimization in \mathbf{x} , we get

$$\min_{\mathbf{x} \in \mathcal{X}, t, \alpha \geq 0, \mathbf{s}} t + r_\xi \alpha + d_*(\mathbf{s}, \alpha, \bar{\mathbf{q}}^\omega) \tag{24a}$$

$$\text{subject to } s_{\omega'} \geq g(\mathbf{x}(\zeta_\omega), \boldsymbol{\xi}_{\omega'}, \gamma) - t, \forall \omega' \in \Omega' \tag{24b}$$

which concludes the proof. \square

B. Acceleration strategy for Algorithm 1

One can possibly accelerate the convergence rate on the bisection Algorithm 1 by exploiting the fact that $\psi(\cdot)$ is a convex function when \mathcal{X} is convex. Indeed, for the current interval $[\gamma^-, \gamma^+]$, $\psi(\gamma)$ can be under- and over-estimated, see Figure 2 (right). The procedure can be described as follows. First, we construct a line that will underestimate ψ by identifying a subgradient of the function at $\tilde{\gamma}$. This can be computed analytically since

$$\begin{aligned}
 \psi(\gamma) &:= \mathbb{E}_{\bar{F}} \left[\min_{\mathbf{x} \in \mathcal{X}} \text{CVaR}_\alpha \left(h(\mathbf{x}, \boldsymbol{\xi}) - \left((1 - \gamma)h(\bar{\mathbf{x}}, \boldsymbol{\xi}) + \gamma \min_{\mathbf{x}' \in \mathcal{X}} h(\mathbf{x}', \boldsymbol{\xi}) \right) \middle| \zeta \right) \right] \\
 &= \mathbb{E}_{\bar{F}} \left[\min_{\mathbf{x} \in \mathcal{X}} \sup_{F \in \bar{\mathcal{D}}(F, \alpha)} \mathbb{E}_F \left[h(\mathbf{x}, \boldsymbol{\xi}) - \left((1 - \gamma)h(\bar{\mathbf{x}}, \boldsymbol{\xi}) + \gamma \min_{\mathbf{x}' \in \mathcal{X}} h(\mathbf{x}', \boldsymbol{\xi}) \right) \middle| \zeta \right] \right] \\
 &\geq \mathbb{E}_{\bar{F}} \left[\sup_{F \in \bar{\mathcal{D}}(F, \alpha)} \min_{\mathbf{x} \in \mathcal{X}} \mathbb{E}_F \left[h(\mathbf{x}, \boldsymbol{\xi}) - \left((1 - \gamma)h(\bar{\mathbf{x}}, \boldsymbol{\xi}) + \gamma \min_{\mathbf{x}' \in \mathcal{X}} h(\mathbf{x}', \boldsymbol{\xi}) \right) \middle| \zeta \right] \right] \\
 &\geq \mathbb{E}_{\bar{F}} \left[\min_{\mathbf{x} \in \mathcal{X}} \mathbb{E}_{F_{\xi|\zeta}^*} \left[h(\mathbf{x}, \boldsymbol{\xi}) - \left((1 - \gamma)h(\bar{\mathbf{x}}, \boldsymbol{\xi}) + \gamma \min_{\mathbf{x}' \in \mathcal{X}} h(\mathbf{x}', \boldsymbol{\xi}) \right) \middle| \zeta \right] \right] \\
 &= \mathbb{E}_{\bar{F}} \left[\underbrace{\min_{\mathbf{x} \in \mathcal{X}} \mathbb{E}_{F_{\xi|\zeta}^*} [h(\mathbf{x}, \boldsymbol{\xi}) - h(\bar{\mathbf{x}}, \boldsymbol{\xi})]}_{\text{offset "a"}} + \gamma \underbrace{\mathbb{E}_{\bar{F}} \left[\mathbb{E}_{F_{\xi|\zeta}^*} \left[h(\bar{\mathbf{x}}, \boldsymbol{\xi}) - \min_{\mathbf{x}' \in \mathcal{X}} h(\mathbf{x}', \boldsymbol{\xi}) \right] \right]}_{\text{slope "b"}} \right],
 \end{aligned}$$

where $F_{\xi|\zeta}^*$ is the conditional probability given ζ of any member (hopefully a maximizer) of $\bar{\mathcal{D}}(\bar{F}, \alpha)$. Note that the first inequality is tight based on Sion's minimax theorem (see [Sion 1958](#)) given that $\bar{\mathcal{D}}(\bar{F}, \alpha)$ is compact, while the second is tight as long as $F_{\xi|\zeta}^*$ achieves the supremum. Such a maximizer can be identified using:

$$F_{\xi|\zeta}^* \in \underset{\substack{F_{\xi|\zeta} \in \mathcal{M}(\Omega_\zeta) : \\ \mathbb{P}_{F_{\xi|\zeta}}(\xi) \leq (1-\alpha)^{-1} \mathbb{P}_{\bar{F}}(\xi|\zeta), \forall \xi}}{\operatorname{argmax}} \mathbb{E}_{F_{\xi|\zeta}} \left[h(\mathbf{x}_\gamma^*(\zeta), \xi) - \left((1-\gamma)h(\bar{\mathbf{x}}, \xi) + \gamma \min_{\mathbf{x}' \in \mathcal{X}} h(\mathbf{x}', \xi) \right) \right]$$

where $\mathbf{x}_\gamma^*(\zeta)$ is the minimizer of (11) with $\zeta_\omega = \zeta$ since $(\mathbf{x}_\gamma^*(\cdot), F^*)$, with F^* as the composition of \bar{F} marginalized on ζ and $F_{\xi|\zeta}^*$,³ is a saddle point of:

$$g(\mathbf{x}(\cdot), F) := \mathbb{E}_F \left[h(\mathbf{x}(\zeta), \xi) - \left((1-\gamma)h(\bar{\mathbf{x}}, \xi) + \gamma \min_{\mathbf{x}' \in \mathcal{X}} h(\mathbf{x}', \xi) \right) \right].$$

Such a $F_{\xi|\zeta}^*$ can be obtained as a side product of solving problem (11) using the optimal dual variables associated with constraint (11b). If we denote by $\gamma_u := a/b$ then the right bound of the interval can be updated to $\gamma^{+'} := \min(\gamma^+, \gamma_u)$.




Figure 2. Visualization of the basic (left) and accelerated (right) bisection algorithm. The blue squared brackets indicate the current estimated interval containing the optimal γ^* and the red squared brackets indicate the interval in the next iterations. The right graph also visualizes the over and under estimators of $\psi(\gamma)$.

The second step is to construct an overestimator. If $\psi(\tilde{\gamma}) > 0$, then we evaluate $\psi(\gamma^-)$ and construct the line that passes through $(\gamma^-, \psi(\gamma^-))$ and $(\tilde{\gamma}, \psi(\tilde{\gamma}))$. If $\psi(\tilde{\gamma}) < 0$ then we evaluate $\psi(\gamma^+)$ and construct the line that passes through $(\gamma^+, \psi(\gamma^+))$ and $(\tilde{\gamma}, \psi(\tilde{\gamma}))$. We denote the point for which the line evaluates to zero as γ_o , and update the left bound of the interval to $\gamma^{-'} := \max(\gamma^-, \gamma_o)$. Hence, the new interval is given by $[\gamma^{-'}, \gamma^+]$, which would potentially significantly reduce the search space.

We conclude this section by commenting that the accelerated bisection algorithm could require up to two evaluations of the ψ function at each iteration instead of a single one as described in the original algorithm.

³Namely, $P_{F^*}(\xi) = P_{\bar{F}}(\xi)$ and $P_{F^*}(\xi|\zeta) = P_{F_{\xi|\zeta}^*}(\xi)$ for all ζ .

C. Algorithms for calibrating the size of the ambiguity sets

Algorithm 2 Algorithm for calibrating the size of the ambiguity set (α) for DRPCR

- 1: Input: Training dataset $\{\zeta_j, \xi_j\}_{j=1}^{N_{train}}$ and validation dataset $\{\zeta_j, \xi_j\}_{j=1}^{N_{validation}}$ and $\mathcal{A} := \{\alpha_i\}_{i=1}^n \subset [0, 1]$
 - 2: Train a random forest model $\hat{F}_{\xi|\zeta}$ on $\{\zeta_j, \xi_j\}_{j=1}^{N_{train}}$
 - 3: Let \tilde{F} be the composition of $\hat{F}_{\xi|\zeta}$ with empirical distribution \hat{F}_ζ of ζ in the training set $\{\zeta_j\}_{j=1}^{N_{train}}$
 - 4: **for** $i = 1, \dots, n$ **do**
 - 5: //Construct $\hat{x}_i^*(\cdot)$ with α_i and \tilde{F}
 - 6: Solve DRPCR with α_i and \tilde{F} to get γ_i^*
 - 7: //Evaluate $\hat{x}_i^*(\cdot)$ on empirical distribution of realizations in $\{\zeta_j, \xi_j\}_{j=1}^{N_{validation}}$
 - 8: **for** $j = 1, \dots, N_{validation}$ **do**
 - 9: Solve (11) with γ_i^* , α_i , and replacing $\mathbb{P}_{\tilde{F}}(\xi = \xi_{\omega'} | \zeta = \zeta_\omega)$ with $\mathbb{P}_{\hat{F}_{\xi|\zeta_j}}(\xi_{\omega'})$ to get optimal \mathbf{x}_j^*
 - 10: Let $\hat{x}_i^*(\zeta_j) := \mathbf{x}_j^*$
 - 11: **end for**
 - 12: Set $s^i := \mathcal{P}_{\hat{F}}^\alpha(\hat{x}_i^*(\cdot))$ for empirical distribution \hat{F} on $\{\zeta_j, \xi_j\}_{j=1}^{N_{validation}}$
 - 13: **end for**
 - 14: Let $i^* := \operatorname{argmax}_i s^i$ and set $\alpha^* := \alpha_{i^*}$, $\gamma^* := \gamma_{i^*}$, and $\mathbf{x}^*(\cdot) := \hat{x}_{i^*}^*(\cdot)$
 - 15: Return $(\alpha^*, \gamma^*, \mathbf{x}^*(\cdot))$
-

Algorithm 3 Algorithm for calibrating the size of the ambiguity set (α) for CVaR-loss/CVaR-regret

- 1: Input: Training dataset $\{\zeta_j, \xi_j\}_{j=1}^{N_{train}}$ and validation dataset $\{\zeta_j, \xi_j\}_{j=1}^{N_{validation}}$ and $\mathcal{A} := \{\alpha_i\}_{i=1}^n \subset [0, 1]$
 - 2: Train a random forest model $\hat{F}_{\xi|\zeta}$ on $\{\zeta_j, \xi_j\}_{j=1}^{N_{train}}$
 - 3: **for** $i = 1, \dots, n$ **do**
 - 4: //Evaluate $\hat{x}_i^*(\cdot)$ on empirical distribution of realizations in $\{\zeta_j, \xi_j\}_{j=1}^{N_{validation}}$
 - 5: **for** $j = 1, \dots, N_{validation}$ **do**
 - 6: Solve (16)/(17) with α_i for $\zeta := \zeta_j$ in validation set to get optimal \mathbf{x}_j^*
 - 7: Let $\hat{x}_i^*(\zeta_j) := \mathbf{x}_j^*$
 - 8: **end for**
 - 9: Set $s^i := \mathcal{P}_{\hat{F}}^\alpha(\hat{x}_i^*(\cdot))$ for empirical distribution \hat{F} on $\{\zeta_j, \xi_j\}_{j=1}^{N_{validation}}$
 - 10: **end for**
 - 11: Let $i^* := \operatorname{argmax}_i s^i$ and set $\alpha^* := \alpha_{i^*}$ and $\mathbf{x}^*(\cdot) := \hat{x}_{i^*}^*(\cdot)$
 - 12: Return $(\alpha^*, \mathbf{x}^*(\cdot))$
-

D. Generation of random covariance matrix with arbitrary variances

A random covariance matrix for the random vector of (ζ, ξ) is generated based on a two-step procedure that follows. The first step consists in generating a random symmetric positive-definite matrix described in Algorithm 4, a method implemented in the `sklearn.datasets.make_spd_matrix` function of scikit-learn machine learning library in Python.

Algorithm 4 Algorithm for generating random symmetric positive-definite matrix

- 1: Input: Dimension of the square matrix $n_\zeta + n_\xi$
 - 2: Generate random square matrix $A_{n_\zeta+n_\xi}$ sampling from the uniform distribution $\mathcal{U}_{[0,1]}$
 - 3: Construct the symmetric matrix $M = A^\top A$
 - 4: Decompose M with Singular Value Decomposition (SVD) method as $M = U\Sigma V^\top$
 - 5: Generate random diagonal matrix S sampling from the uniform distribution $\mathcal{U}_{[0,1]}$
 - 6: Construct $\Sigma' = S + J$ where J is the square matrix of ones with the size of $n_\zeta + n_\xi$
 - 7: Get the symmetric positive-definite matrix as $M' = U\Sigma'V^\top$
 - 8: Return M'
-

Given the vector of standard deviations for (ζ, ξ) denoted by $[\sigma_\zeta^\top \ \sigma_\xi^\top]^\top$ and also a random symmetric positive-definite matrix generated by Algorithm 4, one can implement the second stage described in Algorithm 5 to get a random covariance matrix with arbitrary standard deviations of $[\sigma_\zeta^\top \ \sigma_\xi^\top]^\top$.

Algorithm 5 Algorithm for generating random covariance matrix with arbitrary standard deviations

- 1: Input: Random symmetric positive-definite matrix (M) and vector of standard deviations $[\sigma_\zeta^\top \ \sigma_\xi^\top]^\top$
 - 2: Convert matrix M into its associated correlation matrix $\text{Corr} = (\text{diag}(M))^{-\frac{1}{2}} M (\text{diag}(M))^{-\frac{1}{2}}$
 - 3: Get the arbitrary covariance matrix of $\text{Cov} = \text{diag}\left(\begin{bmatrix} \sigma_\zeta \\ \sigma_\xi \end{bmatrix}\right) (\text{Corr}) \text{diag}\left(\begin{bmatrix} \sigma_\zeta \\ \sigma_\xi \end{bmatrix}\right)$
 - 4: Return Cov
-

E. Average out-of-sample coefficient of prescriptiveness

Table 2. Average out-of-sample coefficient of prescriptiveness

Problem Type	Method	Level of Perturbation					
		0%	20%	30%	40%	50%	60%
Relaxed $x(\cdot)$	CSO	0.45	0.30	0.19	0.04	-0.13	-0.31
	DRCRO	0.45	0.30	0.18	0.04	-0.13	-0.31
	DRCRO	0.45	0.30	0.18	0.04	-0.13	-0.32
	DRPCR	0.45	0.31	0.23	0.13	0.05	0.01
Binary $x(\cdot)$	CSO	0.45	0.30	0.19	0.04	-0.13	-0.31
	DRCRO	0.44	0.30	0.19	0.06	-0.09	-0.25
	DRCRO	0.44	0.30	0.19	0.05	-0.11	-0.28
	DRPCR	0.44	0.32	0.24	0.15	0.07	0.02

F. Additional experiments

While the experiments in Section 4 consider an exact version of the shortest path problem, to be closer to the setting proposed in Kallus & Mao 2023, we also conduct a second set of experiments where $x(\cdot)$ represents relaxed variables. Figure 3 (a) illustrates the coefficients of prescriptiveness obtained from the optimal relaxed policies. These results, in general, are aligned with the ones spotted in Figure 1 (a); however, one remarks the following. Firstly, the results derived from CSO remain exactly the same as the binary case. This stems from the fact that optimal relaxed CSO decisions are known to be integral for the stochastic shortest path problems; conversely, this is not the case for DRCRO, DRCRO, and DRPCR where robustness breaks the linearity of the objective. Secondly, comparing Figures 1 (a) and 3 (a) reveals that forcing DRCRO and DRCRO to propose binary policies enhances their out-of-sample performance, surpassing those of CSO. Indeed, this setting seems to provide these two approaches the chance to better prepare for potential distribution shifts; however, despite their enhanced performance, the highest degree of robustness to distribution shift remains associated with DRPCR policies. Thirdly, this comparative analysis yields counter-intuitive empirical evidence that out-of-sample performance might be slightly improved when imposing integrality constraints on the three robust models. We hypothesize that this might be caused by the additional flexibility of the relaxed models, which makes them more susceptible to overfitting their assumed stochastic models.




Figure 3. Shortest path problem (relaxed version): (a) statistics of the out-of-sample coefficient of prescriptiveness (lower values indicate worse performance). (b) statistics of $\mathbb{E}_{\tilde{F}}[\|\mathbf{x}^*(\zeta) - \hat{\mathbf{x}}\|_1]$ where \tilde{F} is the out-of-sample distribution (lower values reflect a closer proximity to the SAA solution).