

Distributionally Robust Optimization under Decision-Dependent Ambiguity Set with an Application to Machine Scheduling

Nilay Noyan

Industrial Engineering Program, Sabancı University, Istanbul, Turkey, nnoyan@sabanciuniv.edu

Gábor Rudolf

Department of Industrial Engineering, Koc University, Istanbul, Turkey, grudolf@ku.edu.tr

Miguel Lejeune

Department of Decision Sciences, George Washington University, USA, mlejeune@gwu.edu

December 18, 2019

ABSTRACT: We introduce a new class of distributionally robust optimization problems under decision-dependent ambiguity sets. In particular, as our ambiguity sets, we consider balls centered on a decision-dependent probability distribution. The balls are based on a class of *earth mover's distances* that includes both the total variation distance and the Wasserstein metrics. We discuss the main computational challenges in solving the problems of interest, and provide an overview of various settings leading to tractable formulations. Some of the arising side results, such as the mathematical programming expressions for robustified risk measures in a discrete space, are also of independent interest. Finally, we rely on state-of-the-art modeling techniques from humanitarian logistics and machine scheduling to arrive at potentially practical applications, and present a numerical study for a novel risk-averse scheduling problem with controllable processing times.

Keywords: stochastic programming; distributionally robust optimization; decision-dependent ambiguity; earth mover's distances; Wasserstein metric; endogenous uncertainty; decision-dependent probabilities; risk-averse; robustified risk; stochastic scheduling; robust scheduling; robust pre-disaster; random link failures; network interdiction

1. Introduction The classical stochastic programming literature relies on the assumption that the probability distribution of uncertain model parameters is given as a model input, often as set of scenarios along with their probabilities. However, in many decision-making applications the true parameter distribution is unknown. *Distributionally robust optimization* (DRO) is a recent and appreciated approach to hedge against such distributional uncertainty. Instead of assuming that there is a known underlying probability distribution, in DRO one considers an *ambiguity set* that consists of probability distributions, and solves a minimax-type problem to determine decisions that provide hedging against the worst-case parameter distribution in the ambiguity set (see, e.g., [Goh and Sim, 2010](#); [Wiesemann et al., 2014](#)).

Another common fundamental assumption in the stochastic programming literature is that the underlying probability space is independent of the decisions. In other words, it is usually assumed that the probability distributions of random model parameters are exogenously given. In the DRO setting this attitude translates to the assumption that the specified ambiguity set of distributions is decision-independent. However, in certain situations decisions can directly affect the distribution of the parameters, either by changing the parameter realizations or by changing the probabilities of underlying random events that occur after the decisions are taken. This phenomenon is known as *decision-dependent* or *endogenous uncertainty*. For example, in the context of pre-disaster planning, if the links of a transportation network are subject to random failure in case of a disaster, then the investment decisions on strengthening such links (seismic retrofitting of bridges/viaducts on links) can reduce the failure probabilities and improve network survivability ([Peeta et al., 2010](#)).

In our study we aim to simultaneously address distributional and decision-dependent uncertainty. We next provide a brief overview of the relevant literature on these two concepts.

Distributionally robust optimization. The two most widely used types of ambiguity sets in the

DRO literature are moment-based and statistical distance-based ones (for a review, see Postek et al., 2016). Moment-based ambiguity sets contain all probability distributions that satisfy certain general moment conditions (see, e.g., Delage and Ye, 2010; Zymler et al., 2013; Wiesemann et al., 2014). A common example is the ambiguity set consisting of distributions that exactly match the empirical first and second moments; however, such exact moment-based ambiguity sets typically do not contain the true distribution. In addition, as very different distributions can have the same (or similar) lower moments, and the use of higher moments can be impractical, moment-based approaches often lead to overly conservative solutions.

In the present work we therefore limit our attention to statistical distance-based ambiguity sets. These sets consist of probability distributions that are in the vicinity of a nominal distribution—often the empirical one—thought to approximate the true distribution. The vicinity is defined here as a ball centered on the nominal distribution. A wide variety of statistical distances, which provide a measure of dissimilarity between two probability distributions, have been employed to construct such balls. These range from Wasserstein metrics (see, e.g., Pflug and Wozabal, 2007; Esfahani and Kuhn, 2018), the Prohorov metric (see, e.g., Erdoğan and Iyengar, 2006), and ζ -structures (Zhao and Guan, 2018) such as the Kolmogorov–Smirnov statistic and the bounded Lipschitz metric, to the class of ϕ -divergences (see, e.g., Jiang and Guan, 2015). Distances in the latter class are frequently employed in a data-driven context (for an overview see Bayraksan and Love, 2015), and include the Kullback–Leibler divergence (see, e.g., Calafiore, 2007; Hu and Hong, 2012), the Burg entropy (Wang et al., 2016), the total variation distance, the Hellinger distance, the χ^2 distance, and the modified- χ^2 distance. Several studies have highlighted the fact that utilizing a distance-based approach in DRO leads to desirable statistical properties, including consistency and good out-of-sample performance (see, e.g., Lam, 2016; Esfahani and Kuhn, 2018; Van Parys et al., 2017). On a related note, many of the popular regularization methods that are utilized in the machine learning literature to improve out-of-sample performance have recently been shown to be equivalent to statistical distance-based DRO models (see, e.g., Blanchet et al., 2017; Gao et al., 2017). A significant number of DRO studies focus on ϕ -divergences, which are often shown to work well if the uncertain parameters are known to be supported on a discrete set. However, when the possible realizations of parameters form a continuous spectrum, the use of ϕ -divergences can be problematic (as pointed out by, e.g., Blanchet et al., 2017; Gao and Kleywegt, 2016) due to ignoring the metric structure of the realization space, and limiting the support of the measures in the ambiguity set.

An attractive feature of any distance-based approach is that one can control the degree of conservatism simply by adjusting the radius. When using certain distances, such the Wasserstein-1 metric, an appropriate choice of radius can also guarantee that, with a prescribed level of confidence, the true probability distribution belongs to the ambiguity set (see Esfahani and Kuhn, 2018; Zhao and Guan, 2015). This is in contrast to, for example, the Kullback–Leibler divergence, which does not permit the construction of a confidence set that includes the true probability distribution (Esfahani and Kuhn, 2018). We refer to Gao and Kleywegt (2016) for a more elaborate discussion of the pros and cons associated with various ambiguity sets, and in particular the advantages of Wasserstein metrics over ϕ -divergences. Due to these advantages the use of Wasserstein distances in DRO has seen a recent sharp increase, including studies by Pflug and Wozabal (2007), Wozabal (2014), Gao and Kleywegt (2016), Gao and Kleywegt (2017), Ji and Lejeune (2017), Luo and Mehrotra (2017), Esfahani and Kuhn (2018), Ji and Lejeune (2018), and Zhao and Guan (2018). In line with these developments our focus will be on a general class of *earth mover’s distances*, introduced in a discrete context by Rubner et al. (1998). Our chosen class includes both the total variation distance and the Wasserstein-1 metric (also known simply as the Wasserstein metric, or the Kantorovich–Rubinstein metric, Kantorovich and Rubinshtein, 1958), allows the construction of ambiguity sets based on higher-order Wasserstein distances, and also has favorable tractability properties.

Decision-dependent uncertainty. As highlighted in Haus et al. (2017), while decision-dependent uncertainty—endogenous uncertainty—is straightforward to express in the framework of Markov decision processes, its use in stochastic programming remains a tough endeavor, and is far from being a well-resolved issue. Hellemo et al. (2018) discuss the modeling and applications of decision-dependent uncertainty in mathematical programming, and present a taxonomy of stochastic programming approaches with decision-dependent uncertainty. The relevant literature primarily focuses on two types of optimization problems (Goel and Grossmann, 2006): problems with decision-dependent information revelation, and problems with decision-dependent probabilities. In problems of the first type, decisions can partially resolve the uncertainty, affect the timing of uncertainty resolution, and alter the set of possible future random outcomes. In problems of the second-type, decisions alter the probability measures. The first problem type has been addressed more widely (see, e.g., Jonsbråten et al., 1998; Goel and Grossmann, 2004; 2006; Khaligh and MirHassani, 2016) in the literature. Accordingly, in our study, we aim to contribute to the literature by focusing on problems of the second type, where decisions can affect the likelihood of underlying random future events and/or can affect the possible realizations of the random parameters.

Stochastic problems with decision-dependent probability measures are notoriously difficult to model and solve, and, not surprisingly, the relevant literature is quite sparse. Dupacova (2006) briefly discusses optimization under decision-dependent uncertainty, without providing specific formulations or solution methods. Studies that feature algorithmic developments are relatively recent, and typically rely on additional structural properties that are specific to their problems of interest. A significant part of the literature focuses on one particular stochastic pre-disaster investment problem, where the links of a transportation network are subject to probabilistic failures. This problem—originally introduced by Peeta et al. (2010)—aims to use a limited budget to increase the survival probabilities of selected links in such a way that the total expected shortest-path distance between a number of origin-destination pairs is minimized. Modeling the problem in a straightforward fashion involves expressing probabilities as non-linear functions of decision variables, which gives rise to highly non-linear models that are often intractable. Several relevant studies (Flach and Poggi, 2010; Laumanns et al., 2014; Schichl and Sellmann, 2015; Haus et al., 2017) have instead focused on developing efficient alternative solution methods for this particular pre-disaster investment problem. Among these studies, the working papers by Laumanns et al. (2014) and Haus et al. (2017) consider a class of problems where the decisions are binary, and the inherent uncertainty is characterized by a set of binary vectors whose components are independent random variables. They develop effective and exact mixed-integer programming formulations for this class by introducing novel distribution shaping and scenario bundling techniques. These techniques enable an efficient characterization of the decision-dependent scenario probabilities via a set of linear constraints. More recently, Dentcheva and Ruszczyński (2018) have established a new framework, which addresses decision-dependent uncertainty in the context of risk-averse two-stage optimization models with partial information and decision-dependent observation distribution.

DRO with decision-dependent uncertainty. In this study, we incorporate decision-dependent uncertainty into distributionally robust stochastic programming problems via decision-dependent ambiguity sets. Until recently, DRO with decision-dependent ambiguity sets has been an almost untouched research area. Zhang et al. (2016) consider decision-dependent ambiguity sets defined via parametric moment conditions with generic cone constraints. Adopting the total variation metric, the authors establish quantitative stability results for the ambiguity set, the optimal values and solutions. Royset and Wets (2017) utilize recent developments from the variational theory of bivariate functions to establish convergence results for approximations of a class of DRO problems with decision-dependent ambiguity sets. Their discussion covers a variety of ambiguity sets, including moment-based and stochastic dominance-based ones. A major part of their toolset relies on the so-called hypo-distance between CDFs, which

is shown to be a metrization of weak convergence. The recent working paper by Luo and Mehrotra (2018), contemporaneously with our initial report on this subject (Noyan et al., 2018), provides formulations for problems that feature distributional ambiguity sets with decision-dependent radii, as well as for problems with ambiguity sets defined by decision-dependent bounds on moments. Finally, we highlight the recent interest in the tangentially related area of traditional robust optimization models with decision-dependent uncertainty sets (Bertsimas and Vayanos, 2014; Lappas and Gounaris, 2018; Nohadani and Sharma, 2018).

Our contributions. We present a unified modeling framework for a class of DRO problems with decision-dependent earth mover’s distance-based ambiguity sets. Our models typically give rise to non-convex non-linear programs, which are in general very hard to solve. However, we provide an overview of several settings where it is possible to obtain tractable formulations. Some of the side results that make these formulations possible are also of independent interest, including novel mathematical programming expressions for robustified risk measures in a discrete space. We also discuss potential practical applications, utilizing state-of-the-art modeling techniques from the fields of humanitarian logistics and machine scheduling. In particular, we emphasize that we also contribute to the machine scheduling literature by introducing a new risk-averse machine scheduling problem, which simultaneously addresses distributional and decision-dependent uncertainty of controllable stochastic processing times. As a proof of concept for our modeling framework, we also present a computational study, which is one of the very first in the literature for DRO problems with endogenous uncertainty.

Outline. The rest of the paper is organized as follows. In Section 2 we establish necessary notation and recall some basic definitions. Section 3 describes the class of DRO problems of interest. Sections 4–6 are dedicated to developing the corresponding mathematical programming formulations, with Section 5 in particular dedicated to the aforementioned side results about robustified risk measures in a discrete space. Section 7 highlights potential applications: a class of stochastic network problems (with an expanded treatment in the Appendix), and a novel risk-averse scheduling problem with controllable processing times. In Section 8 we present our computational study, and Section 9 contains our concluding remarks regarding future research directions.

2. Preliminaries The set of the first n positive integers is denoted by $[n] = \{1, \dots, n\}$, while the positive part of a number $\eta \in \mathbb{R}$ is denoted by $[\eta]_+ = \max\{\eta, 0\}$. The extended real numbers are denoted by $\bar{\mathbb{R}} = \mathbb{R} \cup \{-\infty, +\infty\}$.

The family of all probability measures on a measurable space (Ω, \mathcal{A}) is denoted by $\mathcal{P}(\Omega, \mathcal{A})$. Let us denote by $\mathcal{L}^m(\Omega, \mathcal{A})$ the family of all measurable mappings from (Ω, \mathcal{A}) to $(\mathbb{R}^m, \mathcal{A}_B^m)$, where \mathcal{A}_B^m is the σ -algebra of m -dimensional Borel sets, and denote the set of m -dimensional random vectors by $\mathcal{V}^m(\Omega, \mathcal{A}) = \mathcal{P}(\Omega, \mathcal{A}) \times \mathcal{L}^m(\Omega, \mathcal{A})$. In a pair $[\mathbb{P}, \boldsymbol{\xi}] \in \mathcal{V}^m(\Omega, \mathcal{A})$ we view the mapping $\boldsymbol{\xi} : \Omega \rightarrow \mathbb{R}^m$ as a random variable on the probability space $(\Omega, \mathcal{A}, \mathbb{P})$, with corresponding CDF $F_{[\mathbb{P}, \boldsymbol{\xi}]}$. If we denote the family of all m -variate CDFs by \mathcal{F}^m , then we trivially have $\{F_{[\mathbb{P}, \boldsymbol{\xi}]} \mid [\mathbb{P}, \boldsymbol{\xi}] \in \mathcal{V}^m(\Omega, \mathcal{A})\} \subset \mathcal{F}^m$, and equality holds if and only if (Ω, \mathcal{A}) is a continuous space, i.e., if it admits a standard continuous uniform random variable. One such continuous space is the standard Borel space $((0, 1), \mathcal{A}_B)$ on the unit interval; we will denote the Borel probability measure on this standard space (i.e., the restriction of the Lebesgue measure to \mathcal{A}_B) by \mathbb{B} . For $1 \leq p \leq \infty$ we introduce the standard L_p -space $L_p^S = \{X \in \mathcal{L}^1((0, 1), \mathcal{A}_B, \mathbb{B}) : \|X\|_{L_p} < \infty\}$, where $\|\cdot\|_{L_p}$ is the L_p -norm for random variables on the standard probability space. In the remainder of the paper we use the common convention whereby p and q refer to a dual pair of values that satisfy $1 \leq p < \infty$, $1 < q \leq \infty$, and $\frac{1}{p} + \frac{1}{q} = 1$.

We define the *law* of a random vector $[\mathbb{P}, \boldsymbol{\xi}] \in \mathcal{V}^m(\Omega, \mathcal{A})$ as the push-forward probability measure $\text{law}[\mathbb{P}, \boldsymbol{\xi}] \in \mathcal{P}(\mathbb{R}^m, \mathcal{A}_B^m)$ given by $[\text{law}[\mathbb{P}, \boldsymbol{\xi}]](A) = \mathbb{P}(\boldsymbol{\xi} \in A)$ for $A \in \mathcal{A}_B^m$. Two random vectors have the

same law if and only if they have the same CDF. By definition, for any measurable space (Ω, \mathcal{A}) we have $\{\text{law}[\mathbb{P}, \boldsymbol{\xi}] : [\mathbb{P}, \boldsymbol{\xi}] \in \mathcal{V}^m(\Omega, \mathcal{A})\} \subset \mathcal{P}(\mathbb{R}^m, \mathcal{A}_B^m)$, and equality again holds for continuous spaces such as $((0, 1), \mathcal{A}_B)$.

Finally, we establish some notational conventions for working with a finite sample space $\Omega = \{\omega^1, \dots, \omega^n\}$. Probability measures on $(\Omega, 2^\Omega)$ are denoted by blackboard bold (double struck) characters, and the probabilities of elementary events by corresponding lowercase letters, e.g., given $\mathbb{P} \in \mathcal{P}(\Omega, 2^\Omega)$ we write $p^i = \mathbb{P}(\{\omega^i\})$. Similarly, we use uppercase letters for scalar-valued random variables, and use the corresponding lowercase letters for their realizations, e.g., given $Z : \Omega \rightarrow \mathbb{R}$ we write $z^i = Z(\omega^i)$. Finally, random vectors are typically denoted by bold Greek letters, and upper indices are used to refer to their realizations, e.g., given $\boldsymbol{\xi} : \Omega \rightarrow \mathbb{R}^m$ we write $\boldsymbol{\xi}^i = \boldsymbol{\xi}(\omega^i)$.

2.1 Earth mover's distances We now introduce a general class of earth mover's distances (EMDs). Consider a function $\delta : \mathbb{R}^m \times \mathbb{R}^m \rightarrow \mathbb{R}_+$, typically chosen to be a symmetric measure of dissimilarity (or distance) between m -dimensional real vectors. We will always assume that δ is *reflexive*, i.e., that $\delta(\mathbf{x}, \mathbf{x}) = 0$ holds for all $\mathbf{x} \in \mathbb{R}^m$. If the stronger condition $\delta(\mathbf{x}_1, \mathbf{x}_2) = 0 \Leftrightarrow \mathbf{x}_1 = \mathbf{x}_2$ holds for all $\mathbf{x}_1, \mathbf{x}_2 \in \mathbb{R}^m$, we say that δ is *definite*. We remark that the choice of \mathbb{R}^m as the native space of realizations is somewhat arbitrary, and for the purposes of the following definitions the space \mathbb{R}^m could be replaced with a general ground set; in the literature the ground set is commonly assumed to be a Polish space with distinguished metric δ . However, we restrict ourselves to working with real vectors, as all of the examples and applications that we discuss later will naturally fit this framework.

The function δ , which measures dissimilarities among vectors in the ground space \mathbb{R}^m , induces an EMD $\Delta : \mathcal{P}(\mathbb{R}^m, \mathcal{A}_B^m) \times \mathcal{P}(\mathbb{R}^m, \mathcal{A}_B^m) \rightarrow \bar{\mathbb{R}}_+$ that measures dissimilarities among probability distributions on \mathbb{R}^m . The EMD between distributions $\mathbb{P}_1, \mathbb{P}_2 \in \mathcal{P}(\mathbb{R}^m, \mathcal{A}_B^m)$ is given by

$$\Delta(\mathbb{P}_1, \mathbb{P}_2) = \inf_{\mathbb{P}^* \in \Pi(\mathbb{P}_1, \mathbb{P}_2)} \int_{\mathbb{R}^m \times \mathbb{R}^m} \delta(\mathbf{x}_1, \mathbf{x}_2) \mathbb{P}^*(d(\mathbf{x}_1, \mathbf{x}_2)), \quad (1)$$

where the infimum is taken over the family of distributions with marginals \mathbb{P}_1 and \mathbb{P}_2 ,

$$\Pi(\mathbb{P}_1, \mathbb{P}_2) = \left\{ \mathbb{P}^* \in \mathcal{P}(\mathbb{R}^m \times \mathbb{R}^m, \mathcal{A}_B^m \times \mathcal{A}_B^m) : \begin{array}{l} \mathbb{P}^*(S \times \mathbb{R}^m) = \mathbb{P}_1(S), \\ \mathbb{P}^*(\mathbb{R}^m \times S) = \mathbb{P}_2(S) \end{array} \text{ for all } S \in \mathcal{A}_B^m \right\}.$$

The above definition can naturally be extended to quantify dissimilarities between any two m -dimensional random vectors. With a slight abuse of notation, for any two measurable spaces $(\Omega_1, \mathcal{A}_1)$ and $(\Omega_2, \mathcal{A}_2)$ the EMD $\Delta : \mathcal{V}^m(\Omega_1, \mathcal{A}_1) \times \mathcal{V}^m(\Omega_2, \mathcal{A}_2) \rightarrow \bar{\mathbb{R}}_+$ will be given by

$$\begin{aligned} \Delta([\mathbb{P}_1, \boldsymbol{\xi}_1], [\mathbb{P}_2, \boldsymbol{\xi}_2]) &= \Delta(\text{law}[\mathbb{P}_1, \boldsymbol{\xi}_1], \text{law}[\mathbb{P}_2, \boldsymbol{\xi}_2]) \\ &= \inf \left\{ \int_{\mathbb{R}^m \times \mathbb{R}^m} \delta(\mathbf{x}_1, \mathbf{x}_2) \mathbb{P}^*(d(\mathbf{x}_1, \mathbf{x}_2)) : \mathbb{P}^* \in \Pi(\text{law}[\mathbb{P}_1, \boldsymbol{\xi}_1], \text{law}[\mathbb{P}_2, \boldsymbol{\xi}_2]) \right\}. \end{aligned} \quad (2)$$

We aim to incorporate distributional uncertainty into decision problems via EMD balls centered on a nominal random vector $[\mathbb{P}, \boldsymbol{\xi}] \in \mathcal{V}^m(\Omega, \mathcal{A})$. Let us first assume that random vectors can take their values from a continuum of possible realizations, and there is ambiguity both in the probability measure and in the realizations. To model such cases, we can construct an EMD ball on the standard probability space, and refer to it as a *continuous EMD ball*. This ball will represent all possible m -dimensional distributions within κ distance from the nominal one:

$$\mathcal{B}_{\delta, \kappa}^{\mathbb{P}}(\boldsymbol{\xi}) = \{\boldsymbol{\zeta} \in \mathcal{L}^m((0, 1), \mathcal{A}_B) : \Delta([\mathbb{P}, \boldsymbol{\xi}], [\mathbb{B}, \boldsymbol{\zeta}]) \leq \kappa\}. \quad (\text{BALL-C})$$

On the other hand, if the realizations of random vectors always belong to some discrete set (e.g., if they are binary), it is not meaningful to consider small variations in realizations. For such cases a natural approach is to construct the EMD ball on the native measurable space of the nominal random vector by allowing the probability measure to change while keeping the realization mapping ξ fixed. We will refer these balls given below as *discrete EMD balls*.

$$\mathcal{B}_{\delta, \kappa}^{\xi}(\mathbb{P}) = \{\mathbb{Q} \in \mathcal{P}(\Omega, \mathcal{A}) : \Delta([\mathbb{P}, \xi], [\mathbb{Q}, \xi]) \leq \kappa\}. \quad (\text{BALL-D})$$

A similar approach is seen, for example, in Pflug and Wozabal (2007), as well as in Pflug and Pichler (2011), where the Wasserstein distance to a reference distribution is minimized among probability distributions with a fixed finite support.

REMARK 2.1 We introduced the definition (BALL-C) instead of the perhaps more natural

$$\mathcal{B}_{\delta, \kappa}([\mathbb{P}, \xi]) = \{\mathbb{Q} \in \mathcal{P}(\mathbb{R}^m, \mathcal{A}_B^m) : \Delta(\text{law}[\mathbb{P}, \xi], \mathbb{Q}) \leq \kappa\}.$$

The two definitions are essentially equivalent, as it is easy to see that $\zeta \in \mathcal{B}_{\delta, \kappa}^{\mathbb{P}}(\xi)$ holds if and only if we have $\text{law}[\mathbb{B}, \zeta] \in \mathcal{B}_{\delta, \kappa}([\mathbb{P}, \xi])$. The definition (BALL-C) was chosen both for notational convenience, and to emphasize that distributions in continuous spaces can be specified via varying outcome mappings (as opposed to varying probability measures). This approach is taken by Pflug et al. (2012) to constructively prove the crucially important Proposition 4.2, which underlies our development in Section 4.

The EMD balls defined in (BALL-C) and (BALL-D) are non-empty for any $\kappa \geq 0$, since due to the reflexivity of δ they always contain the nominal distribution. We also note that the domain of the EMD Δ implicitly depends on the construction used: In (BALL-C) we have $\Delta : \mathcal{V}^m(\Omega, \mathcal{A}) \times \mathcal{V}^m((0, 1), \mathcal{A}_B) \rightarrow \bar{\mathbb{R}}_+$, while in (BALL-D) we have $\Delta : \mathcal{V}^m(\Omega, \mathcal{A}) \times \mathcal{V}^m(\Omega, \mathcal{A}) \rightarrow \bar{\mathbb{R}}_+$. Unless specified otherwise, outside of this preliminary section we will always assume that the sample space Ω is finite, with $\mathcal{A} = 2^{\Omega}$.

The family of EMDs includes widely used metrics such as the *total variation distance*, which (see, e.g., Lindvall, 1992, Theorem 5.2) is the EMD induced by the discrete metric

$$\delta(\mathbf{x}_1, \mathbf{x}_2) = \begin{cases} 0 & \text{if } \mathbf{x}_1 = \mathbf{x}_2 \\ 1 & \text{if } \mathbf{x}_1 \neq \mathbf{x}_2. \end{cases} \quad (3)$$

Wasserstein metrics are also closely related to EMDs. For $p \in [1, \infty)$ the Wasserstein- p metric $W_p : \mathcal{V}^m(\Omega_1, \mathcal{A}_1) \times \mathcal{V}^m(\Omega_2, \mathcal{A}_2) \rightarrow \bar{\mathbb{R}}_+$ is defined as

$$W_p([\mathbb{P}_1, \xi_1], [\mathbb{P}_2, \xi_2]) = \inf \left\{ \left(\int_{\Omega_1 \times \Omega_2} \|\xi_1(\omega_1) - \xi_2(\omega_2)\|_p^p \mathbb{P}^*(d\omega_1, d\omega_2) \right)^{1/p} : \mathbb{P}^* \in \Pi(\text{law}[\mathbb{P}_1, \xi_1], \text{law}[\mathbb{P}_2, \xi_2]) \right\}.$$

It is easy to see that the Wasserstein-1 metric is the EMD induced by the 1-norm distance $\delta(\mathbf{x}_1, \mathbf{x}_2) = \|\mathbf{x}_1 - \mathbf{x}_2\|_1$. More generally, for any $p \in [1, \infty)$ we have $W_p([\mathbb{P}_1, \xi_1], [\mathbb{P}_2, \xi_2]) = \Delta^p([\mathbb{P}_1, \xi_1], [\mathbb{P}_2, \xi_2])^{\frac{1}{p}}$, where Δ^p is the EMD induced by $\delta^p(\mathbf{x}_1, \mathbf{x}_2) = \|\mathbf{x}_1 - \mathbf{x}_2\|_p^p$. It follows that a Wasserstein- p ball of radius κ is identical to the corresponding EMD ball with the same center, and a radius of κ^p .

2.2 Risk measures Unless specified otherwise, the definitions and results in this section are presented for risk measures that are natively defined on a standard L_p space. Any such risk measure $\rho : L_p^S \rightarrow \mathbb{R}$ can be naturally extended to p -integrable random variables defined on an arbitrary probability space $(\Omega, \mathcal{A}, \mathbb{P})$ via *inverse transform sampling* as follows: It is well-known that if $X : \Omega \rightarrow \mathbb{R}$ is a random variable, then its *generalized inverse CDF* $F_X^{(-1)} : (0, 1) \rightarrow \mathbb{R}$, given by $F_X^{(-1)}(\alpha) = \inf\{x \in \mathbb{R} : F_X(x) \geq \alpha\}$ and viewed as a random variable on the standard space

$((0, 1), \mathcal{A}_B, \mathbb{B})$, has the same CDF as X itself. Consequently, X is p -integrable if and only if we have $F_X^{(-1)} \in L_p^S$, in which case with a slight abuse of notation we will write $\rho([\mathbb{P}, X]) = \rho(X) = \rho(F_X^{(-1)})$.

Risk measures are functionals that represent the risk associated with a random variable by a scalar value, and their desirable properties, such as law invariance and coherence, are axiomatized in Artzner et al. (1999). Throughout this paper we limit our attention to law invariant coherent risk measures. We say that a mapping $\rho : L_p^S \rightarrow \mathbb{R}$ is a *coherent risk measure* if ρ has the following properties (for all $V, V_1, V_2 \in L_p^S$):

- *Monotone*: $V_1 \leq V_2 \Rightarrow \rho(V_1) \leq \rho(V_2)$.
- *Convexity*: $\rho(\lambda V_1 + (1 - \lambda)V_2) \leq \lambda\rho(V_1) + (1 - \lambda)\rho(V_2)$ for all $\lambda \in [0, 1]$.
- *Translation equivariant*: $\rho(V + \lambda) = \rho(V) + \lambda$ for all $\lambda \in \mathbb{R}$.
- *Positive homogeneous*: $\rho(\lambda V) = \lambda\rho(V)$ for all $\lambda \geq 0$.

The more general class of convex risk measures is obtained by dropping positive homogeneity (Föllmer and Schied, 2002). For a more general discussion on quantifying risk we refer to Müller and Stoyan (2002), Pflug and Römisch (2007), and Shapiro et al. (2009). We now introduce an important family of coherent risk measures. The *conditional value-at-risk* at confidence level $\alpha \in [0, 1)$ for a random variable Z is defined (Rockafellar and Uryasev, 2000) as

$$\text{CVaR}_\alpha(Z) = \min \left\{ \eta + \frac{1}{1 - \alpha} \mathbb{E}([Z - \eta]_+) : \eta \in \mathbb{R} \right\}. \quad (4)$$

The minimum in (4) is attained at the α -quantile, which is known as the *value-at-risk* (VaR) at confidence level α : $\text{VaR}_\alpha(Z) = \min\{\eta \in \mathbb{R} : P(Z \leq \eta) \geq \alpha\}$. For risk-averse decision makers typical choices for the confidence level are large values such as $\alpha = 0.9$.

Suppose that Z is a discrete random variable with realizations z^1, \dots, z^n , and corresponding probabilities p^1, \dots, p^n . Then $\text{VaR}_\alpha(Z) = z^j$ holds for at least one $j \in [n]$, which implies

$$\text{CVaR}_\alpha(Z) = \min_{j \in [n]} z^j + \frac{1}{1 - \alpha} \sum_{i \in [n]} p^i [z^i - z^j]_+. \quad (5)$$

It is also well known that the optimization problem in (4) can equivalently be formulated as the following linear program:

$$\min \left\{ \eta + \frac{1}{1 - \alpha} \sum_{i \in [n]} p^i v^i : v^i \geq z^i - \eta \quad \forall i \in [n], \quad \mathbf{v} \in \mathbb{R}_+^n, \quad \eta \in \mathbb{R} \right\}. \quad (6)$$

CVaR has been widely used in decision-making problems under uncertainty due to a number of useful properties. It captures a wide range of risk preferences, including risk-neutral (for $\alpha = 0$) and pessimistic worst-case (for sufficiently large values of α , when $\alpha \rightarrow 1$) preferences. It is also a spectral risk measure (Acerbi, 2002) and thus can be viewed as a weighted sum of the least favorable outcomes as illustrated by the following dual representations of CVaR_α :

$$\max \left\{ \frac{1}{1 - \alpha} \sum_{i \in [n]} \beta^i z^i : \sum_{i \in [n]} \beta^i = 1 - \alpha, \quad 0 \leq \beta^i \leq p^i \quad \forall i \in [n] \right\} = \frac{1}{1 - \alpha} \int_\alpha^1 \text{VaR}_a(Z) da. \quad (7)$$

The knapsack-type maximization problem in (7) is equivalent to the linear programming dual of (6), and to the well-known *risk envelope-based dual representation* of CVaR (see, e.g., Rockafellar, 2007). Due to the last equality, CVaR is also known in the literature as *average value-at-risk* and *tail value-at-risk*.

CVaR is of particular importance as it serves as a fundamental building block for other coherent risk

measures (Kusuoka, 2001). It was shown in Noyan and Rudolf (2015) that the class of risk measures that can be obtained by extending a law invariant coherent risk measure from L_p^S via inverse transform sampling coincide with the class of operators with so-called *Kusuoka representations* of the form

$$\rho(X) = \sup_{\mu \in \mathcal{M}} \int_0^1 \text{CVaR}_\alpha(X) \mu(d\alpha) \quad \text{for all } X \in L_p(\Omega, \mathcal{A}, \mathbb{P}),$$

where \mathcal{M} is a family of probability measures on $(0, 1)$. When this family consists of finitely many finitely supported measures, we say that ρ is *finitely representable* (we note that such risk measures are dense among coherent ones, see Noyan and Rudolf, 2013). If the family \mathcal{M} consist only of a *single* such measure, i.e., if ρ is a convex combination of finitely many CVaRs, then ρ is called a *mixed CVaR measure*. Finally, we note that for finite probability spaces the class of mixed CVaR measures coincides with the class of spectral risk measures (Noyan and Rudolf, 2015).

3. Distributionally Robust Optimization Models We are now ready to introduce the main focus of the present work, a class of distributionally robust stochastic optimization problems with decision-dependent ambiguity sets. To begin, let us consider a simple stochastic optimization problem: The decision maker aims to minimize the expected value of an outcome $G(\mathbf{x}, \boldsymbol{\xi})$, where \mathbf{x} is a decision belonging to some feasible set \mathcal{X} , and the outcome, given by the mapping $G : \mathcal{X} \times \mathbb{R}^m \rightarrow \mathbb{R}$, depends on an m -dimensional random vector $\boldsymbol{\xi}$. In particular, we are interested in problems with *endogenous uncertainty*, where the distribution of the parameter vector $\boldsymbol{\xi}$ can depend on the decision \mathbf{x} . More precisely, given mappings $\mathbb{P} : \mathcal{X} \rightarrow \mathcal{P}(\Omega, \mathcal{A})$ and $\boldsymbol{\xi} : \mathcal{X} \rightarrow \mathcal{L}^m(\Omega, \mathcal{A})$ the problem takes the form

$$\min_{\mathbf{x} \in \mathcal{X}} \mathbb{E}_{\mathbb{P}(\mathbf{x})} (G(\mathbf{x}, \boldsymbol{\xi}(\mathbf{x}))). \quad (8)$$

The next step is to account for uncertainty about the distribution of the parameters. To this end, we introduce as our ambiguity set an EMD ball, either of type (BALL-C) or of type (BALL-D), centered on the *nominal* random parameter vector $[\mathbb{P}(\mathbf{x}), \boldsymbol{\xi}(\mathbf{x})] \in \mathcal{V}^m(\Omega, \mathcal{A})$. This leads to the following DRO variants of the (risk-neutral) *underlying problem* (8):

$$\min_{\mathbf{x} \in \mathcal{X}} \sup_{\zeta \in \mathcal{B}_{\delta, \kappa}^{\mathbb{P}(\mathbf{x})}(\boldsymbol{\xi}(\mathbf{x}))} \mathbb{E}_{\mathbb{B}} (G(\mathbf{x}, \zeta)) \quad (\text{DRO-RNC})$$

$$\min_{\mathbf{x} \in \mathcal{X}} \sup_{\mathbb{Q} \in \mathcal{B}_{\delta, \kappa}^{\boldsymbol{\xi}(\mathbf{x})}(\mathbb{P}(\mathbf{x}))} \mathbb{E}_{\mathbb{Q}} (G(\mathbf{x}, \boldsymbol{\xi}(\mathbf{x}))). \quad (\text{DRO-RND})$$

Recalling our notation from Section 2.1, here κ is the radius of the ball, and δ is the underlying distance or dissimilarity measure on \mathbb{R}^m . In an applied context the appropriate choices of κ and δ , as well as the choice between the models (DRO-RNC) and (DRO-RND) will be driven both by the specifics of the base problem and by tractability concerns.

Aiming to minimize the expected value of an outcome represents a risk-neutral attitude. To incorporate risk-aversion into our decision problems we can replace the expected value operator in (DRO-RNC) and (DRO-RND) with an appropriately chosen risk measure ρ , leading to the problems

$$\min_{\mathbf{x} \in \mathcal{X}} \sup_{\zeta \in \mathcal{B}_{\delta, \kappa}^{\mathbb{P}(\mathbf{x})}(\boldsymbol{\xi}(\mathbf{x}))} \rho(G(\mathbf{x}, \zeta)), \quad (\text{DRO-RAC})$$

$$\min_{\mathbf{x} \in \mathcal{X}} \sup_{\mathbb{Q} \in \mathcal{B}_{\delta, \kappa}^{\boldsymbol{\xi}(\mathbf{x})}(\mathbb{P}(\mathbf{x}))} \rho([\mathbb{Q}, G(\mathbf{x}, \boldsymbol{\xi}(\mathbf{x}))]). \quad (\text{DRO-RAD})$$

REMARK 3.1 *While our focus in this paper is on the decision-dependent nominal distribution of the parameter vector, our framework could also allow for the radius κ of the ambiguity set to be decision-*

dependent (as considered by Luo and Mehrotra, 2018). One possible approach is to make κ itself a decision variable and add to the objective function a term that penalizes low values of κ , effectively introducing a cost of robustness (analogous to the cost associated with the reliability level in chance-constrained optimization, see, e.g., Lejeune and Shen, 2016).

3.1 Specifying the nominal distribution One of the main distinguishing features of our approach is that the nominal distribution at the center of the ambiguity set is decision-dependent; in this section we briefly discuss possible ways to describe this dependence.

The case when parameter realizations are decision-dependent, but the probabilities of underlying events are not, is fairly straightforward, as it is sufficient to specify the mappings $\mathbf{x} \mapsto \xi^i(\mathbf{x})$ for each scenario $i \in [n]$. In Section 7.2 we present two representative examples of such mappings in the context of machine scheduling problems, where the uncertain parameters are the processing times of jobs. The first example introduces *linearly compressible processing times* with continuous control decisions, while the second example—*control with discrete resources*—features binary control decisions.

We next turn our attention to the opposite case, when parameter realizations are fixed, but probabilities are decision-dependent. While this setting formally appears quite similar to the one discussed above, it is typically very challenging to construct scenario probability mappings $\mathbf{x} \mapsto p^i(\mathbf{x})$ that can properly model problems of practical interest while maintaining a reasonable level of tractability. In Section 7.1 we discuss the recently developed technique of *distribution shaping*, which allows one to express multiplicative probabilities via linear constraints for certain problem classes with binary decisions. Another interesting special case is when the random parameter vector is drawn from a population that consists of subpopulations whose proportions are decision-dependent (see, e.g., Dupacova, 2006; Hellemo et al., 2018). For example, in a revenue management context the subpopulations would correspond to various customer types or market segments whose proportions are influenced by marketing or pricing decisions. More precisely, given a fixed outcome mapping $\xi : \Omega \rightarrow \mathbb{R}$ let $\mathbb{P}_1, \dots, \mathbb{P}_S \in \mathcal{P}(\Omega)$ denote the probability measures associated with the S subpopulations, and let $\pi_1(\mathbf{x}), \dots, \pi_S(\mathbf{x})$ denote the corresponding proportions of each subpopulation in the population. Then the nominal parameter vector follows a *mixture distribution* $[\mathbb{P}, \xi]$ with $\mathbb{P} = \sum_{s=1}^S \pi_s(\mathbf{x}) \mathbb{P}_s$. If we have $\mathcal{X} \subset \mathbb{R}^r$ for some $r \in \mathbb{N}$, and the mappings π_s are affine, with $\pi_s(\mathbf{x}) = \pi_s^0 + \pi_s^\top \mathbf{x}$ for some $\pi_1, \dots, \pi_S \in \mathbb{R}^r$, then scenario probabilities can be expressed via the linear constraints $p^i(\mathbf{x}) = \sum_{s=1}^S \pi_s^0 p_s^i + \pi_s^\top \mathbf{x} p_s^i$ for $i \in [n]$.

4. Formulations for continuous Wasserstein balls We now turn our attention to a class of problems where outcome mapping G has a bilinear structure, and the ambiguity set is a continuous Wasserstein- p ball. Our principal tool to obtain potentially tractable formulations for problems in this class will be Proposition 4.2, due to Pflug et al. (2012), which generalizes the following well-known consequence of Hölder’s inequality to a stochastic context.

PROPOSITION 4.1 For any two vectors $\mathbf{v}, \mathbf{y}_0 \in \mathbb{R}^m$ and $\kappa \geq 0$ we have

$$\sup_{\mathbf{y} \in B_\kappa^p(\mathbf{y}_0)} \mathbf{y}^\top \mathbf{v} = \mathbf{y}_0^\top \mathbf{v} + \kappa \|\mathbf{v}\|_q,$$

where $B_\kappa^p(\mathbf{y}_0) = \{\mathbf{y} \in \mathbb{R}^m : \|\mathbf{y} - \mathbf{y}_0\|_p \leq \kappa\}$ is the p -norm ball of radius κ centered on \mathbf{y}_0 .

The above proposition concerns the robustification of a scalar product with respect to one of its factors, using a p -norm ball as the ambiguity set. We next consider a stochastic variant of this problem where we replace the central vector \mathbf{y}_0 with a nominal random vector $[\mathbb{B}, \xi]$, and replace the p -norm ball with a Wasserstein- p ball as the ambiguity set. When working in a risk-averse framework, our focus will be on an

appropriate risk measure of the arising random scalar products. Following along the lines of Pflug et al. (2012) we introduce an important class of risk measures.

DEFINITION 4.1 *Let $\rho : L_p^S \rightarrow \mathbb{R}$ be a law-invariant convex risk measure that admits a representation of the form $\rho(V) = \max \{ \mathbb{E}_{\mathbb{B}}(VZ) - R(Z) : Z \in L_q^S \}$ where $R : L_q^S \rightarrow \bar{\mathbb{R}}$ is a convex functional. When $p > 1$, we say that ρ is well-behaved with factor $C \in \mathbb{R}_+$ if*

$$\|Z\|_{L_q} = C \text{ holds for all } Z \in \bigcup_{V \in L_p^S} \arg \max \{ \mathbb{E}_{\mathbb{B}}(VZ) - R(Z) : Z \in L_q^S \}.$$

When $p = 1$, we say that ρ is well-behaved with factor C if, for the random variables Z specified in the above condition, in addition to $\|Z\|_{L_\infty} = C$ we also have $Z \in \{0, C\}$ almost everywhere.

Before we state the following key result from Pflug et al. (2012), we recall from Section 2.1 that the Wasserstein- p ball of radius κ centered on a random vector $[\mathbb{B}, \boldsymbol{\xi}] \in \mathcal{V}^m((0, 1), \mathcal{A}_B)$ is identical to the EMD ball $\mathcal{B}_{\delta^p, \kappa^p}^{\mathbb{B}}(\boldsymbol{\xi})$ with radius κ^p , where δ^p is the measure of dissimilarity induced by the p -th power of the p -norm.

PROPOSITION 4.2 *Consider a random vector $[\mathbb{B}, \boldsymbol{\xi}] \in \mathcal{V}^m((0, 1), \mathcal{A}_B)$, and assume that the law invariant convex risk measure $\rho : L_p \rightarrow \mathbb{R}$ is well-behaved with factor C . Then for any $\mathbf{v} \in \mathbb{R}^m$ such that $\boldsymbol{\xi}^\top \mathbf{v} \in L_p^S$ we have*

$$\sup_{\boldsymbol{\zeta} \in \mathcal{B}_{\delta^p, \kappa^p}^{\mathbb{B}}(\boldsymbol{\xi})} \rho(\boldsymbol{\zeta}^\top \mathbf{v}) = \rho(\boldsymbol{\xi}^\top \mathbf{v}) + C\kappa \|\mathbf{v}\|_q. \quad (9)$$

Wozabal (2014) applies this result to provide robustified versions of many popular risk measures; here we only mention the following important corollary:

$$\sup_{\boldsymbol{\zeta} \in \mathcal{B}_{\delta^1, \kappa}^{\mathbb{B}}(\boldsymbol{\xi})} \text{CVaR}_\alpha(\boldsymbol{\zeta}^\top \mathbf{v}) = \text{CVaR}_\alpha(\boldsymbol{\xi}^\top \mathbf{v}) + \frac{1}{1-\alpha} \kappa \|\mathbf{v}\|_\infty. \quad (10)$$

REMARK 4.1 *Proposition 4.2 provides a closed-form robustification of convex risks for the case of a bilinear outcome mapping. Zhao and Guan (2018), Gao and Kleywegt (2016), and Esfahani and Kuhn (2018) present similarly-structured results for more general classes of outcome functions. However, these results are restricted to the risk-neutral case, and in place of a closed formula they employ duality to express the robustified value as the optimum of a minimization problem.*

We next examine the implications of Proposition 4.2 on the optimization problems introduced in Section 3, focusing on the case when the outcome mapping has a bilinear structure. More precisely, we assume that the outcome mapping is of the form $G(\mathbf{x}, \boldsymbol{\zeta}) = \boldsymbol{\zeta}^\top \mathbf{v}(\mathbf{x})$ for some vector-valued mapping $\mathbf{v} : \mathcal{X} \rightarrow \mathbb{R}^m$. We first observe that in this case, due to the linearity of expectation, the risk-neutral underlying problem (8) is equivalent to the deterministic problem $\min_{\mathbf{x} \in \mathcal{X}} \bar{\boldsymbol{\xi}}^\top(\mathbf{x}) \mathbf{v}(\mathbf{x})$, where the mapping $\bar{\boldsymbol{\xi}} : \mathcal{X} \rightarrow \mathbb{R}^m$ is given by $\bar{\boldsymbol{\xi}}(\mathbf{x}) = \mathbb{E}(\boldsymbol{\xi}(\mathbf{x}))$. Noting that the expected value operator is trivially well-behaved with factor 1, it is easy to verify that for $\rho = \mathbb{E}$ the formula (9) becomes equivalent to the conclusion of Proposition 4.1 with $\mathbf{y}_0 = \mathbb{E}(\boldsymbol{\xi})$. Therefore the risk-neutral DRO problem (DRO-RNC) with decision-dependent ambiguity set $\mathcal{B}_{\delta^p, \kappa^p}^{\mathbb{B}}(\boldsymbol{\xi})$ can be equivalently reformulated as the following deterministic optimization problem:

$$\min_{\mathbf{x} \in \mathcal{X}} \bar{\boldsymbol{\xi}}^\top(\mathbf{x}) \mathbf{v}(\mathbf{x}) + \kappa \|\mathbf{v}(\mathbf{x})\|_q. \quad (11)$$

The risk-averse variant of the problem, where ρ is an arbitrary law invariant convex risk measure that is

well-behaved with some factor C , can be similarly reformulated using Proposition 4.2, leading to

$$\min_{\mathbf{x} \in \mathcal{X}} \rho(\boldsymbol{\xi}^\top(\mathbf{x})\mathbf{v}(\mathbf{x})) + C\kappa\|\mathbf{v}(\mathbf{x})\|_q. \quad (12)$$

In contrast to the risk-neutral case, this reformulated problem typically remains inherently stochastic.

5. Robustified risk measures in finite spaces In Section 4 we managed to convert the minimax DRO problem (DRO-RAC), which features a continuous EMD ball of type (BALL-C) as its ambiguity set, to a straightforward minimization. Our eventual goal is to similarly convert the problem (DRO-RAD), which arises when the ambiguity set is a discrete EMD ball of type (BALL-D). The primary difficulty lies in the fact that Proposition 4.2, which provided an elegant way to robustify risk measures in a continuous context by replacing the supremum over the ambiguity set with the closed-form formula (9), is no longer valid in a discrete setting, as the following example shows.

EXAMPLE 5.1 Let $\boldsymbol{\xi}$ be a 2-dimensional random vector with possible realizations $(1, 0)^\top$ and $(0, 1)^\top$, and let $\mathbf{x} = (1, 1)^\top$. Then $\mathbb{E}_{\mathbb{Q}}(\mathbf{x}^\top \boldsymbol{\xi}) = 1 < 1 + \kappa\|\mathbf{x}\|_q$ for any probability distribution \mathbb{Q} .

We mention that a one-sided version of Proposition 4.2, analogous to Lemma 1 of Pflug et al. (2012), remains true for discrete EMD balls:

PROPOSITION 5.1 Consider an arbitrary measurable space (Ω, \mathcal{A}) and a random vector $[\mathbb{P}, \boldsymbol{\xi}] \in \mathcal{V}^m(\Omega, \mathcal{A})$. If the law invariant convex risk measure $\rho : L_p \rightarrow \mathbb{R}$ is well-behaved with factor C , then for any $\mathbf{v} \in \mathbb{R}^m$ such that $\boldsymbol{\xi}^\top \mathbf{v} \in L_p^S$ we have

$$\sup_{\mathbb{Q} \in \mathcal{B}_{\delta^p, \kappa^p}^{\boldsymbol{\xi}}(\mathbb{P})} \rho([\mathbb{Q}, \boldsymbol{\zeta}^\top \mathbf{v}]) \leq \rho([\mathbb{P}, \boldsymbol{\xi}^\top \mathbf{v}]) + C\kappa\|\mathbf{v}\|_q. \quad (13)$$

PROOF. The discrete EMD ball $\mathcal{B}_{\delta^p, \kappa^p}^{\boldsymbol{\xi}}(\mathbb{P})$ can be trivially embedded into the continuous ball $\mathcal{B}_{\delta^p, \kappa^p}^{\mathbb{P}}(\boldsymbol{\xi})$ as follows. Consider a probability measure $\mathbb{Q} \in \mathcal{B}_{\delta^p, \kappa^p}^{\boldsymbol{\xi}}(\mathbb{P})$. Using the well-known fact that every finite-dimensional distribution can be realized on a probability space that admits a continuous uniform distribution, there exists a mapping $\boldsymbol{\zeta} \in \mathcal{L}^m([0, 1])$ such that $\text{law}[\mathbb{B}, \boldsymbol{\zeta}] = \text{law}[\mathbb{Q}, \boldsymbol{\xi}]$. As EMDs are defined in a law-invariant fashion, $\boldsymbol{\zeta} \in \mathcal{B}_{\delta^p, \kappa^p}^{\mathbb{P}}(\boldsymbol{\xi})$ immediately follows. Furthermore, since ρ is also law-invariant, we have $\rho([\mathbb{Q}, \boldsymbol{\xi}^\top \mathbf{v}]) = \rho([\mathbb{B}, \boldsymbol{\zeta}^\top \mathbf{v}])$. Therefore the supremum in (13) is taken over a smaller set than the one in (9), which implies our proposition. \square

While we do not have a closed-form analogue to formula (9) for discrete spaces, in this section we develop some mathematical tools to replace the supremum involved in the robustification of certain risk measures with an equivalent minimization. These tools will then be utilized to recast (DRO-RAD) as a conventional optimization problem; in Section 6.1 we examine certain important cases where this approach leads to potentially tractable formulations. Throughout the remainder of this section $\boldsymbol{\xi} : \Omega \rightarrow \mathbb{R}^m$ will denote a fixed mapping from a finite sample space of size n , and we will use the notation $\delta^{ij} = \delta(\boldsymbol{\xi}^i, \boldsymbol{\xi}^j)$ for distances among the realizations of $\boldsymbol{\xi}$, where $i, j \in [n]$.

5.1 A parametric relation between random variables For two scalar-valued random variables $X, Y \in \mathcal{L}^1(\Omega, 2^\Omega)$ the usual ordering relation $X \geq Y$ holds if and only if we have $x^i \geq y^i$ for all $i \in [n]$. A key idea behind the developments of this section is that one can robustify certain risk expressions by replacing the usual ordering with a parametric family of relations, and introducing a corresponding ‘‘penalty term’’.

DEFINITION 5.1 Given a threshold $\tau \geq 0$ we define the relation \succeq_τ as follows. For $X, Y \in \mathcal{L}^1(\Omega, \mathcal{A})$

$$X \succeq_\tau Y \quad \text{holds if and only if we have} \quad x^i \geq y^j - \delta^{ij}\tau \quad \text{for all } i, j \in [n].$$

While \succeq_τ is typically not a preorder among random variables, it is closely related to the usual ordering. The properties below are easily verified:

- (i) The relation $X \succeq_\tau Y$ implies $X \geq Y$, due to the reflexivity of δ .
- (ii) If δ is definite, then for sufficiently high values of τ the relations $X \succeq_\tau Y$ and $X \geq Y$ are equivalent. In particular, the equivalence holds when $\tau \geq \max_{i,j \in [n]} \frac{y^j - x^i}{\delta^{ij}}$.
- (iii) The relation $X \succeq_0 Y$ is equivalent to $X \geq \sup(Y)$.
- (iv) When δ is the discrete metric, the relation $X \succeq_\tau Y$ is equivalent to the conventional inequality $X \geq \max(Y, \sup(Y) - \tau)$.

REMARK 5.1 The relation $X \succeq_\tau Y$ is equivalent to $X(\omega^i) \geq R(\tau, i)$ for all $i \in [n]$, where $R : \mathbb{R}_+ \times [n] \rightarrow \mathbb{R}$ is given by $\sup_{j \in [n]} Y(\omega^j) - \tau \delta^{ij}$. Here $-R$ can be viewed as a regularization operator of $Y : \Omega \rightarrow \mathbb{R}$, providing a discrete analogue to Definition 3 of Gao and Kleywegt (2016) (which in turn generalizes the classical Moreau-Yosida regularization).

We next present and discuss the main results of Section 5, which will then be proved in Section 5.3.

5.2 Robustified risk formulas Let $\mathbb{P} \in \mathcal{P}(\Omega, 2^\Omega)$ be a fixed nominal probability measure. Given a risk measure $\rho : L_p^S \rightarrow \mathbb{R}$ and a radius $\kappa \geq 0$ we define the *robustified risk measure* $\rho^\kappa : \mathcal{L}^1(\Omega, 2^\Omega) \rightarrow \mathbb{R}$ on our finite probability space by

$$\rho^\kappa(Z) = \sup \left\{ \rho(\mathbb{Q}, Z) : \mathbb{Q} \in \mathcal{B}_{\delta, \kappa}^\xi(\mathbb{P}) \right\} \quad \text{for } Z \in \mathcal{L}^1(\Omega, 2^\Omega). \quad (14)$$

We now present the robustified versions of several important risk measures; the corresponding proofs can be found in the next section.

5.2.1 Robustified expectation The following expression closely parallels the trivial formula $\mathbb{E}_\mathbb{P}(Z) = \inf \{ \mathbb{E}_\mathbb{P}(V) : V \geq Z \}$ for the nominal expectation, with the relation \succeq_τ playing a similar role to that of the usual ordering \geq :

$$\mathbb{E}^\kappa(Z) = \inf \{ \mathbb{E}_\mathbb{P}(V) + \kappa\tau : \tau \geq 0, V \succeq_\tau Z \}. \quad (15)$$

The additional “robustification term” $\kappa\tau$, which also appears in the results below, is analogous to the term seen when robustifying the expected value operator in a continuous space (see Section 4).

EXAMPLE 5.2 (TOTAL VARIATION DISTANCE) When the ambiguity set is based on the total variation distance, it is easy to identify the worst-case distribution, as it can be obtained by greedily “transferring probability” from favorable outcomes (starting with the lowest outcome value) to the worst-case (highest) outcome, until either the boundary of the ambiguity set is reached, or all probability is transferred to the worst case. As it has been observed in the literature (Jiang and Guan, 2018, Theorem 1; see also Rahimian et al., 2018, Proposition 3), this implies that the robustified expectation is a convex combination of the worst-case outcome and the nominal CVaR at an appropriate level, and thus a coherent risk measure of the outcome. More precisely, if δ is the discrete metric, then, introducing the notation $z^+ = \sup(Z)$, for $\kappa \in [0, 1]$ we have $\mathbb{E}^\kappa(Z) = \kappa z^+ + (1 - \kappa) \text{CVaR}_\kappa(Z)$. Using the representation (6) for CVaR, we can

then express $\mathbb{E}^\kappa(Z)$ as the optimum of the following LP:

$$\min \quad \kappa z^+ + (1 - \kappa) \left(\eta + \frac{1}{1 - \kappa} \sum_{i \in [n]} p^i \hat{v}^i \right) \quad (16a)$$

$$\text{s.t.} \quad \hat{v}^i \geq z^i - \eta, \quad \forall i \in [n] \quad (16b)$$

$$\hat{v}^i \geq 0, \quad \forall i \in [n] \quad (16c)$$

$$\eta \leq z^+. \quad (16d)$$

Here the redundant constraint (16d) reflects the trivial inequality $\text{VaR}_\kappa(Z) \leq \sup(Z)$. The above formulation turns out to be essentially the same as the LP formulation of (15) given in (26). To see the correspondence between these two LPs, we first note that in accordance with Property (iv) we can rewrite constraints (26b) as

$$v^i \geq z^i, \quad \forall i \in [n]$$

$$v^i \geq z^+ - \tau, \quad \forall i \in [n].$$

Let us introduce the change of variables $\eta = z^+ - \tau$, $\hat{v}^i = v^i + \tau - z^+$ for $i \in [n]$. It is now easy to verify that the formulations (16) and (26) are equivalent. We note that the preceding argument constitutes an alternative proof for Theorem 1 of Jiang and Guan (2018) in our discrete setting. Additionally, it follows that the optimum in (15) can be attained when we have $\tau = \sup(Z) - \text{VaR}_\kappa(Z)$.

5.2.2 Robustified CVaR Recalling the definition of CVaR from (4), for a probability level $\alpha \in [0, 1)$ we have

$$\text{CVaR}_\alpha^\kappa(Z) = \inf \left\{ \eta + \mathbb{E}_\mathbb{P}(S) + \kappa\tau : \eta \in \mathbb{R}, \tau \geq 0, S \succeq_\tau \frac{1}{1 - \alpha} [Z - \eta]_+ \right\}. \quad (17)$$

This robustified expression exhibits a similar structure to (4), again with an additional robustification term. By applying a scaling factor of $(1 - \alpha)$ to S and τ , we can also rewrite (17) as

$$\text{CVaR}_\alpha^\kappa(Z) = \inf \left\{ \eta + \mathbb{E}_\mathbb{P} \left(\frac{1}{1 - \alpha} S \right) + \frac{1}{1 - \alpha} \kappa\tau : \eta \in \mathbb{R}, \tau \geq 0, S \succeq_\tau [Z - \eta]_+ \right\}. \quad (18)$$

This version better highlights the parallels with the corresponding continuous result in (10), where the robustification term for $\text{CVaR}_\alpha(\boldsymbol{\xi}^\top \mathbf{z})$ took the form $\frac{1}{1 - \alpha} \kappa \|\mathbf{z}\|_\infty$. However, in contrast to (17), the formula (18) does not generalize in a straightforward fashion to mixed CVaR measures.

EXAMPLE 5.3 (TOTAL VARIATION DISTANCE) Similarly to the case of robustified expectation, when the ambiguity set is based on the total variation distance, we can express $\text{CVaR}_\alpha^\kappa(Z)$ as a convex combination of the worst-case outcome, and a nominal CVaR of the outcome at an appropriate level. Recalling our notation from Example 5.2, we first observe that if $\kappa \geq 1 - \alpha$ holds, then the ambiguity set contains a distribution where Z takes value z^+ with a probability of at least $1 - \alpha$, which immediately implies $\text{CVaR}_\alpha^\kappa(Z) = z^+$. On the other hand, in the non-trivial case when $\kappa \leq 1 - \alpha$ holds, we have

$$\text{CVaR}_\alpha^\kappa(Z) = \frac{\kappa}{1 - \alpha} z^+ + \frac{1 - \alpha - \kappa}{1 - \alpha} \text{CVaR}_{\alpha + \kappa}(Z). \quad (19)$$

While we are not aware of the above formula appearing elsewhere in the literature, it can be proved analogously to Theorem 1 in Jiang and Guan (2018), because the worst-case distribution is obviously the same as for the case of robustified expectation. To obtain an alternative proof, we can also start from the LP representation (28) of the formula (17), and apply the same change of variables as in Example 5.2 to obtain an LP representation of (19). Like before, this approach also shows that the optimum in (17) can be obtained when we have $\tau = z^+ - \text{VaR}_{\alpha + \kappa}(Z)$.

5.2.3 Robustified mixed CVaR Making explicit the definition from Section 2.2, given a finitely supported probability measure μ on the interval $[0, 1)$, the *mixed* CVaR risk measure $\rho_{\{\mu\}} : L_p^S \rightarrow \mathbb{R}$ is given by

$$\rho_{\{\mu\}}(Z) = \int_0^1 \text{CVaR}_\alpha(Z) \mu(d\alpha) = \sum_{\alpha \in \text{supp}(\mu)} \mu(\{\alpha\}) \text{CVaR}_\alpha(Z). \quad (20)$$

We note that, according to the above expression, the risk measure $\rho_{\{\mu\}}$ can be interpreted as the expected value of CVaR_α when the level α is randomly selected from the interval $[0, 1)$ according to the probability measure μ . More precisely, if we denote the identity function of the interval by $A : [0, 1) \rightarrow [0, 1)$, then we have $\rho_{\{\mu\}}([\mathbb{P}, Z]) = \mathbb{E}_\mu(\text{CVaR}_A([\mathbb{P}, Z]))$. The robustification of $\rho_{\{\mu\}}^\kappa$ is now given by the following generalization of (17):

$$\rho_{\{\mu\}}^\kappa(Z) = \inf \left\{ \mathbb{E}_\mu(H) + \mathbb{E}_\mathbb{P}(S) + \kappa\tau : H \in \mathbb{R}^{[0,1)}, \tau \geq 0, S \succeq_\tau \mathbb{E}_\mu \left(\frac{1}{1-A} [Z - H]_+ \right) \right\}. \quad (21)$$

Here A (“capital alpha”) is viewed as the probability level of CVaR, selected randomly according to μ . Similarly, the random variable H (“capital eta”) plays the role of the VaR value at level A .

5.2.4 Robustified finitely representable risk measures As discussed in Section 2.2, a finite family \mathcal{M} of finitely supported probability measures on $[0, 1)$ defines a finitely representable risk measure $\rho_{\mathcal{M}} : L_p^S \rightarrow \mathbb{R}$ given by

$$\rho_{\mathcal{M}}(Z) = \sup_{\mu \in \mathcal{M}} \rho_{\{\mu\}}(Z). \quad (22)$$

While the motivation behind the next formula is to robustify this important class of risk measures, it remains valid even when the cardinality of the family \mathcal{M} is infinite.

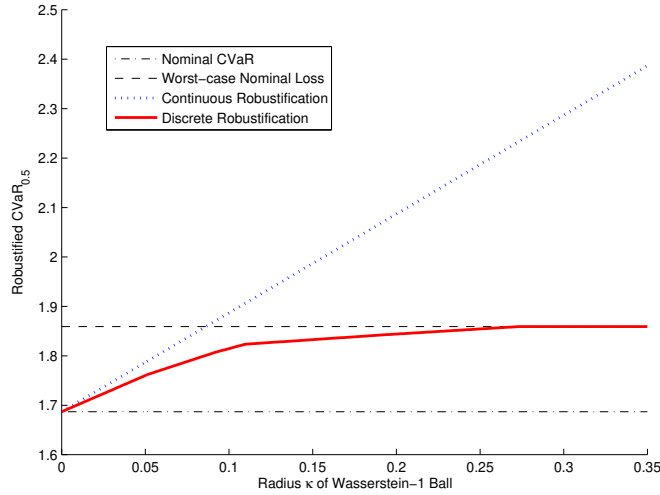
$$\rho_{\mathcal{M}}^\kappa(Z) = \inf \left\{ R \in \mathbb{R} : H \in \mathbb{R}^{[0,1)}, \tau \in \mathbb{R}_+^{\mathcal{M}}, S_\mu \succeq_{\tau_\mu} \mathbb{E}_\mu \left(\frac{1}{1-A} [Z - H]_+ \right), \forall \mu \in \mathcal{M}, R \geq \mathbb{E}_\mu(H) + \mathbb{E}_\mathbb{P}(S_\mu) + \kappa\tau_\mu \right\}. \quad (23)$$

We remark that the domain of the mapping $H : [0, 1) \rightarrow \mathbb{R}$ in the above formulas can be restricted from $[0, 1)$ to the support set $\bigcup_{\mu \in \mathcal{M}} \text{supp}(\mu)$. Similarly to the role of the threshold η in the expected excess-based representation (4) of CVaR, we can view H as representing the VaR functional under the worst-case distribution in the ambiguity set. More precisely, if for $\rho = \rho_{\mathcal{M}}$ the supremum in (14) is attained at $\mathbb{P}^* \in \mathcal{B}_{\delta, \kappa}^\xi(\mathbb{P})$, then the choice $H^*(\alpha) = \text{VaR}_\alpha([\mathbb{P}^*, Z])$ is optimal in (21) and (23).

5.2.5 Robustification in discrete and continuous cases We would like to highlight that the above robustification formulas exhibit fundamentally different qualitative properties than their counterparts in continuous spaces, despite the similar formal structures, due to the boundedness of the distance operator δ . In more detail, Pflug et al. (2012) show that, when taking the supremum in a Wasserstein ball of type (BALL-C), the worst-case distribution can be obtained by starting from the nominal random realization vector, and moving in a fixed direction until we reach the boundary of the ball. This leads to the robustified risk growing linearly in terms of the ball radius, as seen in (9). By contrast, when considering balls of type (BALL-D), the supremum is bounded by the risk attained at the degenerate distribution where all probability is concentrated on the worst-case outcome. Therefore, if the ambiguity ball is large enough to contain this degenerate distribution, further increasing the radius has no impact on the robustified risk. These behaviors are illustrated in Figure 1, which compares the Wasserstein-1 robustifications of $\text{CVaR}_{0.5}$ for an equal-weight three-asset portfolio, where the nominal asset loss realizations have been randomly generated, and are equally likely.

5.3 Proof of robustified risk formulas We will use linear programming duality to derive the formulas of the previous section. To this end, let us begin by establishing a characterization of EMD

Figure 1: Continuous vs. discrete robustification



balls in finite probability spaces via a system of linear inequalities. We mention that this result has been established for the case of p -norm distances in Gao and Kleywegt (2016) via more general strong duality results.

LEMMA 5.1 For two probability measures $\mathbb{P}, \mathbb{Q} \in \mathcal{P}(\Omega, 2^\Omega)$ and a radius $\kappa \geq 0$ we have $\mathbb{Q} \in B_{\delta, \kappa}^\xi(\mathbb{P})$ if and only if the following system of inequalities is feasible.

$$\sum_{j \in [n]} \gamma^{ij} = p^i, \quad \forall i \in [n] \quad (24a)$$

$$\sum_{i \in [n]} \gamma^{ij} = q^j, \quad \forall j \in [n] \quad (24b)$$

$$\sum_{i \in [n]} \sum_{j \in [n]} \delta^{ij} \gamma^{ij} \leq \kappa, \quad (24c)$$

$$\gamma \in \mathbb{R}_+^{n \times n}. \quad (24d)$$

PROOF. Introducing the notation $\hat{\mathbb{P}} = \text{law}[\xi, \mathbb{P}]$ and $\hat{\mathbb{Q}} = \text{law}[\xi, \mathbb{Q}]$, the condition $\mathbb{Q} \in B_{\delta, \kappa}^\xi(\mathbb{P})$ is by definition equivalent to the inequality $\Delta(\hat{\mathbb{P}}, \hat{\mathbb{Q}}) \leq \kappa$. This inequality is in turn equivalent to the feasibility of the following system of inequalities:

$$\sum_{\mathbf{y} \in \text{supp}(\hat{\mathbb{Q}})} \hat{\gamma}(\mathbf{x}, \mathbf{y}) = \hat{\mathbb{P}}(\{\mathbf{x}\}), \quad \forall \mathbf{x} \in \text{supp}(\hat{\mathbb{P}}) \quad (25a)$$

$$\sum_{\mathbf{x} \in \text{supp}(\hat{\mathbb{P}})} \hat{\gamma}(\mathbf{x}, \mathbf{y}) = \hat{\mathbb{Q}}(\{\mathbf{y}\}), \quad \forall \mathbf{y} \in \text{supp}(\hat{\mathbb{Q}}) \quad (25b)$$

$$\sum_{\mathbf{x} \in \text{supp}(\hat{\mathbb{P}})} \sum_{\mathbf{y} \in \text{supp}(\hat{\mathbb{Q}})} \delta(\mathbf{x}, \mathbf{y}) \hat{\gamma}(\mathbf{x}, \mathbf{y}) \leq \kappa, \quad (25c)$$

$$\hat{\gamma} : \text{supp}(\hat{\mathbb{P}}) \times \text{supp}(\hat{\mathbb{Q}}) \rightarrow \mathbb{R}_+. \quad (25d)$$

We can obtain this second equivalence by directly applying the EMD definition (1) to finitely supported measures, with the joint probability measure \mathbb{P}^* supported on $\text{supp}(\hat{\mathbb{P}}) \times \text{supp}(\hat{\mathbb{Q}})$ and given there by $\mathbb{P}^*(\{(\mathbf{x}, \mathbf{y})\}) = \hat{\gamma}(\mathbf{x}, \mathbf{y})$. The lemma then follows immediately from the two observations below:

- Assume that the system (24) has a feasible solution γ . Keeping in mind the trivial equalities

$\hat{\mathbb{P}}(\{\mathbf{x}\}) = \sum_{i: \xi^i = \mathbf{x}} p^i$ and $\hat{\mathbb{Q}}(\{\mathbf{y}\}) = \sum_{j: \xi^j = \mathbf{y}} q^j$, it is easy to verify that the aggregated values $\hat{\gamma}(\mathbf{x}, \mathbf{y}) = \sum_{i: \xi^i = \mathbf{x}} \sum_{j: \xi^j = \mathbf{y}} \gamma^{ij}$ solve the system (25), which implies $\mathbb{Q} \in B_{\delta, \kappa}^{\xi}(\mathbb{P})$.

- If $\mathbb{Q} \in B_{\delta, \kappa}^{\xi}(\mathbb{P})$ holds, then the system (25) has a feasible solution $\hat{\gamma}$. It is again easy to verify that the disaggregated values $\gamma^{ij} = \hat{\gamma}(\xi(\omega^i), \xi(\omega^j)) \frac{p^i}{\hat{\mathbb{P}}(\xi^i)} \frac{q^j}{\hat{\mathbb{Q}}(\xi^j)}$, where $\frac{0}{0}$ is understood as zero, solve the system (24).

□

5.3.1 Robustified expectation We first point out that the formula (15) follows directly from applying the CVaR formula (17) with $\alpha = 0$. Here we also present a short stand-alone proof, which will serve as a template for our later more complex arguments. By Lemma 5.1 we can express the robustified expectation $\mathbb{E}^{\kappa}(Z)$ as the optimum value of the following LP:

$$\max \left\{ \sum_{j \in [n]} z^j q^j : (24a)-(24d) \right\}.$$

We can somewhat simplify this LP by replacing each variable q^j with the sum $\sum_{i \in [n]} \gamma^{ij}$, and removing the now redundant defining constraints (24b). By taking the dual of the simplified LP we can express $\mathbb{E}^{\kappa}(Z)$ via linear minimization as

$$\min \sum_{i \in [n]} p^i v^i + \kappa \tau \quad (26a)$$

$$\text{s.t. } v^i \geq z^i - \delta^{ij} \tau, \quad \forall i, j \in [n] \quad (26b)$$

$$\tau \geq 0. \quad (26c)$$

Noting that $\sum_{i \in [n]} p^i v^i = \mathbb{E}_{\mathbb{P}}(V)$, and that the constraints (26b) are equivalent to the relation $V \succeq_{\tau} Z$, the desired formula (15) follows.

5.3.2 Robustified CVaR Following the same logic as before, we can combine Lemma 5.1 with the dual representation of CVaR given in (7) to obtain the robustified CVaR value $\text{CVaR}_{\alpha}^{\kappa}(Z)$ as the optimum value of the LP

$$\max \frac{1}{1-\alpha} \sum_{j \in [n]} z^j \beta^j \quad (27a)$$

$$\text{s.t. } (24a)-(24c), \quad (27b)$$

$$\beta^j \leq q^j, \quad \forall j \in [n] \quad (27c)$$

$$\sum_{j \in [n]} \beta^j = 1 - \alpha, \quad (27d)$$

$$\gamma \in \mathbb{R}_+^{n \times n}, \quad \beta \in \mathbb{R}_+^n. \quad (27e)$$

We can again simplify the LP formulation by eliminating the q^j variables, and take the dual afterwards. Applying a scaling factor of $1-\alpha$ to each dual variable, we arrive at the following expression of $\text{CVaR}_{\alpha}^{\kappa}(Z)$:

$$\min \eta + \frac{1}{1-\alpha} \sum_{i \in [n]} p^i v^i + \frac{1}{1-\alpha} \kappa \tau \quad (28a)$$

$$\text{s.t. } v^i \geq z^i - \eta - \delta^{ij} \tau, \quad \forall i, j \in [n] \quad (28b)$$

$$\mathbf{v} \in \mathbb{R}_+^n, \quad (28c)$$

$$\tau \geq 0. \tag{28d}$$

The constraints (28b) are clearly equivalent to the relation $V \succeq_{\tau} Z - \eta$, and the non-negativity of V immediately implies $V \succeq_{\tau} 0$. Combining these two relations we obtain $V \succeq_{\tau} [Z - \eta]_+$, and the formula (18), which is trivially equivalent to the desired (17), follows. We mention that, in addition to its role in proving our concise formulas, the LP formulation (28) will also prove valuable as a tool to explicitly incorporate robustified risk into mathematical programming formulations.

5.3.3 Robustified mixed CVaR Linear formulations involving CVaR can be extended to mixed CVaR measures by introducing duplicate variables and constraints corresponding to each probability level in the (finite) support of the mixing measure (see Noyan and Rudolf 2013 or Noyan and Rudolf 2018 for more detailed discussion and examples). The desired formula (21) follows from these extended linear formulations via LP duality in exactly the same fashion as before, so for the sake of conciseness we omit the lengthy details.

5.3.4 Robustified finitely representable risk measures Finally, bypassing a direct LP duality argument, the formula (23) follows directly from (21). Combining the observation that we have

$$\rho_{\mathcal{M}}^{\kappa}(Z) = \sup_{\mathcal{Q}} \sup_{\mu} \rho_{\{\mu\}}([\mathbb{P}, Z]) = \sup_{\mu} \sup_{\mathcal{Q}} \rho_{\{\mu\}}([\mathbb{P}, Z]) = \sup_{\mu} \rho_{\{\mu\}}^{\kappa}(Z)$$

with the trivial formula $\sup A = \inf\{R \in \mathbb{R} : R \geq a \ \forall a \in A\}$ for expressing the supremum of a set $A \subset \mathbb{R}$ we immediately obtain (23), with one slight difference: The formula (23) features a single variable H , while the direct approach we outlined would introduce an indexed family $(H_{\mu})_{\mu \in \mathcal{M}}$, similarly to other duplicated variables. However, as discussed at the end of Section 5.2, it can be assumed without loss of generality that these H_{μ} variables all express the VaR functional under the worst-case distribution, and therefore coincide.

6. Formulations for discrete EMD balls The robustification formula (17) and its LP expression (28) enable us to recast our minimax DRO problem as a conventional minimization problem for the case $\rho = \text{CVaR}_{\alpha}$. Using the system (28) to represent the supremum in (DRO-RAD) we obtain the formulation

$$\min \quad \eta + \frac{1}{1-\alpha} \sum_{i \in [n]} p^i(\mathbf{x}) v^i + \frac{1}{1-\alpha} \kappa \tau \tag{29a}$$

$$\text{s.t.} \quad v^i \geq G(\mathbf{x}, \boldsymbol{\xi}^j(\mathbf{x})) - \eta - \delta^{ij} \tau, \quad \forall i, j \in [n] \tag{29b}$$

$$\delta^{ij} = \delta(\boldsymbol{\xi}^i(\mathbf{x}), \boldsymbol{\xi}^j(\mathbf{x})), \quad \forall i, j \in [n] \tag{29c}$$

$$\mathbf{v} \in \mathbb{R}_+^n, \tau \in \mathbb{R}_+, \mathbf{x} \in \mathcal{X}. \tag{29d}$$

The case when we have $\alpha = 0$ and $\rho = \text{CVaR}_0 = \mathbb{E}$ is somewhat simpler, because we can utilize (26) in place of (28) to formulate (DRO-RND) as

$$\min \quad \sum_{i \in [n]} p^i(\mathbf{x}) v^i + \kappa \tau \tag{30a}$$

$$\text{s.t.} \quad v^i \geq G(\mathbf{x}, \boldsymbol{\xi}^j(\mathbf{x})) - \delta^{ij} \tau, \quad \forall i, j \in [n] \tag{30b}$$

$$\delta^{ij} = \delta(\boldsymbol{\xi}^i(\mathbf{x}), \boldsymbol{\xi}^j(\mathbf{x})), \quad \forall i, j \in [n] \tag{30c}$$

$$\mathbf{v} \in \mathbb{R}^n, \tau \in \mathbb{R}_+, \mathbf{x} \in \mathcal{X}. \tag{30d}$$

We saw in Section 4 that the risk-neutral underlying problem (8) is deterministic. However, this is no longer the case for the above DRO variant. As Example 7.4 shows, it is possible that, given two nominal distributions with the same mean, the resulting robustified expectations are different.

REMARK 6.1 *The LP expression of the robustified CVaR formula (17) facilitates a conventional optimization formulation of (DRO-RAD). As discussed in Section 5.3, analogous, although more complex, linear expressions can be obtained for the robustification formulas (21) and (23) for mixed and finitely representable coherent risk measures. Similarly to (29), these linear formulations can then be used to cast (DRO-RAD) as a conventional minimization problem when the risk measure ρ belongs to one of these more general classes. As our primary focus in the remainder of this paper is on problems that feature the canonical risk measure $\rho = \text{CVaR}_\alpha$, the arising extended versions of (29) are omitted for the sake of brevity.*

6.1 Towards tractable formulations We have seen that, under appropriate assumptions, it is possible to reformulate (DRO-RAD) as a (typically non-linear) optimization problem of the form (29). We now turn our attention to the computational challenges involved in solving such problems, and will examine several important problem classes where these challenges can be mitigated.

6.1.1 Decision-independent nominal realizations If the uncertain vector $\xi(\mathbf{x})$ depends on the decision \mathbf{x} in a non-trivial fashion, then this dependence becomes a significant source of non-linearity in (29). However, if the nominal realizations are decision-independent, then we can drop the argument \mathbf{x} from the terms $\xi^i(\mathbf{x})$, $\xi^j(\mathbf{x})$ for all $i, j \in [n]$, and replace them with a common uncertain vector ξ . Consequently, the distance values δ^{ij} can also be viewed as fixed parameters, defined by the equations $\delta^{ij} = \delta(\xi^i, \xi^j)$. If the set \mathcal{X} of feasible decisions is polyhedral, and the mapping $\mathbf{x} \mapsto G(\mathbf{x}, \xi)$ is linear, then (29) becomes a linearly constrained problem (apart from the possible non-linearity implicit in the constraint $\mathbf{x} \in \mathcal{X}$). A more general version of this statement is given precise form in the remark below. Along similar lines, if $\mathbb{P}(\mathbf{x})$ depends on \mathbf{x} in a linear fashion, then the objective function is quadratic.

REMARK 6.2 *Consider a feasible set $\mathcal{X} \subset \mathbb{R}^{r_1}$ for some $r_1 \in \mathbb{N}$, and assume that we can express $G(\mathbf{x}, \xi^j)$ as the minimum of an LP. More precisely, we assume that for each $j \in [n]$ there exist matrices $A_j \in \mathbb{R}^{r_3 \times r_2}$, $B_j \in \mathbb{R}^{r_3 \times r_1}$ and vectors $\mathbf{c}_j \in \mathbb{R}^{r_1}$, $\mathbf{d}_j \in \mathbb{R}^{r_2}$, $\mathbf{b}_j \in \mathbb{R}^{r_3}$ for some $r_2, r_3 \in \mathbb{N}$ such that for every decision $\mathbf{x} \in \mathcal{X}$ the outcome $G(\mathbf{x}, \xi^j)$ is the minimum of the LP*

$$\begin{aligned} \min \quad & \mathbf{c}_j^\top \mathbf{x} + \mathbf{d}_j^\top \mathbf{y} \\ \text{s.t.} \quad & A_j \mathbf{y} \geq B_j \mathbf{x} + \mathbf{b}_j, \\ & \mathbf{y} \in \mathbb{R}^{r_2}. \end{aligned}$$

Then we can formulate (29) as the following linearly constrained program:

$$\begin{aligned} \min \quad & \eta + \frac{1}{1-\alpha} \sum_{i \in [n]} p^i(\mathbf{x}) v^i + \frac{1}{1-\alpha} \kappa \tau \\ \text{s.t.} \quad & v^i \geq \mathbf{c}_j^\top \mathbf{x} + \mathbf{d}_j^\top \mathbf{y}_j - \eta - \delta^{ij} \tau, & \forall i, j \in [n] \\ & A_j \mathbf{y}_j \geq B_j \mathbf{x} + \mathbf{b}_j, & \forall j \in [n] \\ & \mathbf{v} \in \mathbb{R}_+^n, \tau \in \mathbb{R}_+, \mathbf{x} \in \mathcal{X}, \\ & \mathbf{y}_j \in \mathbb{R}^{r_2}, & \forall j \in [n]. \end{aligned}$$

6.1.2 Using the discrete metric Let us assume that δ is the discrete metric given by (3). As discussed in Section 2.1, this choice of δ allows us to use total variation distance-based balls as ambiguity sets. We now present a streamlined formulation of our DRO problem under the additional assumptions that neither the nominal realizations nor the outcome realizations are decision-dependent. Remarkably, while these assumptions appear to be highly restrictive, the resulting problem class still contains highly non-trivial instances of practical interest, such as our formulations for the pre-disaster planning problems

detailed in Section 7.1. Let us again denote the nominal realizations by $\xi^i \in \mathbb{R}^m$, and the corresponding outcome realizations by $G^i \in \mathbb{R}$, for $i \in [n]$. In addition, let $G^+ = \max_{j \in [n]} G^j$. We can then reformulate (29) as follows (matching Property (iv) in Section 5.1):

$$\begin{aligned} \min \quad & \eta + \frac{1}{1-\alpha} \sum_{i \in [n]} p^i(\mathbf{x}) v^i + \frac{1}{1-\alpha} \kappa \tau \\ \text{s.t.} \quad & v^i \geq G^i - \eta, & \forall i \in [n] \\ & v^i \geq G^+ - \eta - \tau, & \forall i \in [n] \\ & \mathbf{v} \in \mathbb{R}_+^n, \tau \in \mathbb{R}_+, \mathbf{x} \in \mathcal{X}. \end{aligned}$$

Analogously to the difference between (29) and (30), when the underlying problem is risk-neutral, i.e., when we have $\alpha = 0$, we can further simplify the above formulation by removing (or setting to zero) the auxiliary variable η , and dropping the non-negativity requirement for the variables \mathbf{v} .

6.1.3 Using the Wasserstein-1 metric When the nominal realizations are decision-dependent, the distances between pairs of realizations are represented by the variables δ^{ij} in (29). Whether the corresponding defining constraints (29c) can be represented in a fashion that is amenable to computations depends on the choice of the reflexive mapping $\delta : \mathbb{R}^m \times \mathbb{R}^m \rightarrow \mathbb{R}_+$. We now examine the important case when δ is the 1-norm distance, i.e., when the ambiguity set is a Wasserstein-1 ball. Let us assume that the decision-dependent parameters are bounded, i.e., that there exists some $M \in \mathbb{R}_+$ such that we have $\|\xi(\mathbf{x})\|_{L_\infty} < \frac{M}{4}$ for all $\mathbf{x} \in \mathcal{X}$.

REMARK 6.3 *The scaling for the constant $\frac{M}{4}$ in the previous condition was chosen in order to simplify the notation in our optimization formulations. It is easy to see that the condition is satisfied when \mathcal{X} is compact and the mapping $\mathbf{x} \mapsto \xi(\mathbf{x})$ is continuous. In the general case the boundedness condition can be replaced by the following weaker requirement: We assume that the range of each coordinate of the parameter vector is bounded by a decision-independent constant, i.e., that there exists $M \in \mathbb{R}_+$ such that $|\xi_k^i(\mathbf{x}) - \xi_k^j(\mathbf{x})| < \frac{M}{2}$ holds for all $i, j \in [n]$, $k \in [m]$, and $\mathbf{x} \in \mathcal{X}$.*

Noting that the equations in (29c) will take the form

$$\delta^{ij} = \delta(\xi^i(\mathbf{x}), \xi^j(\mathbf{x})) = \|\xi^i(\mathbf{x}) - \xi^j(\mathbf{x})\|_1 = \sum_{k \in [m]} \left| \xi_k^i(\mathbf{x}) - \xi_k^j(\mathbf{x}) \right|, \quad (32)$$

let us introduce the auxiliary variables ν_k^{ij} to represent the values $|\xi_k^i(\mathbf{x}) - \xi_k^j(\mathbf{x})|$ for all $i, j \in [n]$ and $k \in [m]$. We can then equivalently reformulate our problem as

$$\min \quad \eta + \frac{1}{1-\alpha} \sum_{i \in [n]} p^i(\mathbf{x}) v^i + \frac{1}{1-\alpha} \kappa \tau \quad (33a)$$

$$\text{s.t.} \quad v^i \geq G(\mathbf{x}, \xi^j(\mathbf{x})) - \eta - \sum_{k \in [m]} \nu_k^{ij} \tau, \quad \forall i \in [n], j \in [n] \quad (33b)$$

$$\nu_k^{ij} \leq \xi_k^i(\mathbf{x}) - \xi_k^j(\mathbf{x}) + M \lambda_k^{ij}, \quad \forall i \in [n], j \in [n], k \in [m] \quad (33c)$$

$$\nu_k^{ij} \leq -\xi_k^i(\mathbf{x}) + \xi_k^j(\mathbf{x}) + M(1 - \lambda_k^{ij}), \quad \forall i \in [n], j \in [n], k \in [m] \quad (33d)$$

$$\lambda \in \{0, 1\}^{n \times n \times m}, \quad \nu \in \mathbb{R}_+^{n \times n \times m}, \quad (33e)$$

$$\mathbf{v} \in \mathbb{R}_+^n, \tau \in \mathbb{R}_+, \mathbf{x} \in \mathcal{X}. \quad (33f)$$

We note that the constraints (33c)–(33e) are equivalent to the inequalities $\nu_k^{ij} \leq |\xi_k^i(\mathbf{x}) - \xi_k^j(\mathbf{x})|$ for all $i, j \in [n]$, and $k \in [m]$. It is possible to ensure (without changing the optimum of the problem) that the opposite inequalities $\nu_k^{ij} \geq |\xi_k^i(\mathbf{x}) - \xi_k^j(\mathbf{x})|$ also hold, by adding the corresponding redundant constraints

$\nu_k^{ij} \geq \xi_k^i(\mathbf{x}) - \xi_k^j(\mathbf{x})$ and $\nu_k^{ij} \geq -\xi_k^i(\mathbf{x}) + \xi_k^j(\mathbf{x})$ to (33).

6.1.4 Utilizing a comonotone structure The formulation (33) features the auxiliary variables λ_k^{ij} , along with the corresponding constraints (33c)–(33e), which represent the potentially non-convex relations $\nu_k^{ij} \leq |\xi_k^i(\mathbf{x}) - \xi_k^j(\mathbf{x})|$. The introduction of binary variables and big-M constraints often leads to significant computational challenges. However, this issue can be avoided when the mappings $i \mapsto \xi_k^i(\mathbf{x}_1)$ and $i \mapsto \xi_k^i(\mathbf{x}_2)$ are comonotone for any $\mathbf{x}_1, \mathbf{x}_2 \in \mathcal{X}$ and $k \in [m]$. If this condition is satisfied, then for any $i, j \in [n]$ and $k \in [m]$ there are two possibilities: Either $\xi_k^i(\mathbf{x}) \geq \xi_k^j(\mathbf{x})$ holds for all $\mathbf{x} \in \mathcal{X}$, in which case we can set $\nu_k^{ij} = \xi_k^i(\mathbf{x}) - \xi_k^j(\mathbf{x})$, or $\xi_k^i(\mathbf{x}) \leq \xi_k^j(\mathbf{x})$ holds for all $\mathbf{x} \in \mathcal{X}$, in which case we can set $\nu_k^{ij} = -\xi_k^i(\mathbf{x}) + \xi_k^j(\mathbf{x})$. Since these new equality constraints ensure that we have $\nu_k^{ij} = |\xi_k^i(\mathbf{x}) - \xi_k^j(\mathbf{x})|$ for all $i, j \in [n]$ and $k \in [m]$, the auxiliary λ_k^{ij} variables can be dropped from the formulation along with the constraints (33c)–(33e).

While the above comonotonicity condition is naturally satisfied for certain applications, including some of the machine scheduling problems we discuss in Section 7.2, it is somewhat restrictive. However, it is possible to take partial advantage of the above improvements even if the condition is violated. Let us partition the index set $I = [n] \times [n] \times [m]$ into sets $I_+ = \{(i, j, k) \in I \mid \xi_k^i(\mathbf{x}) \geq \xi_k^j(\mathbf{x}) \forall \mathbf{x} \in \mathcal{X}\}$, $I_- = \{(i, j, k) \in I \mid \xi_k^i(\mathbf{x}) \leq \xi_k^j(\mathbf{x}) \forall \mathbf{x} \in \mathcal{X}\}$, and $I_0 = I \setminus (I_+ \cup I_-)$. Then constraints (33c)–(33e) can be replaced with their simpler variants outlined above for indices in $I_+ \cup I_-$, and kept in their original form for indices in I_0 .

6.1.5 A parametric programming approach We again consider the general setting where nominal realizations are decision-dependent, and note that the non-convex quadratic terms $\delta^{ij}\tau$ in the constraints (29b) constitute a significant potential obstacle when working toward a tractable approach to solving the problem (29). Fortunately, all of these terms feature the variable τ as a common factor. Therefore, if we fix the value of τ , all of the quadratic terms in question become linear. In certain cases this leads to an optimization problem that belongs to a more tractable class than the original. For example, if the mapping $\mathbf{x} \mapsto G(\mathbf{x}, \boldsymbol{\xi}(\mathbf{x}))$ was linear, then fixing the value of τ would change quadratic constraints into linear ones. We can therefore attempt to solve (29) by performing a single-parameter search over the possible values of τ .

This approach is closely related to the field of parametric programming. In this context, calculating the optimum of (29b) for a fixed value of τ can be seen as evaluating the *optimum value function* (OVF) of a parametric non-linear program (see, e.g., Kyparisis and Fiacco, 1987, both for a quick introduction to the subject, and for a precise statement of the convexity results discussed below). If the OVF has certain favorable properties, such as convexity or unimodality, then the aforementioned single-parameter search can potentially lead to a viable solution strategy (e.g., by using the bisection or golden section search method) with performance guarantees. While there are a variety of results that prove generalized convexity properties for OVFs, they typically require objective and constraining functions to be jointly convex in all variables. It appears that establishing joint convexity for general problems in the classes that we study is highly non-trivial, except under very restrictive assumptions (such as requiring all probabilities p^i and realization distances δ^{ij} to be decision-independent). However, it still seems plausible that this approach can be leveraged for problems with additional underlying structure. Along similar lines, it can be relatively straightforward to obtain a Lipschitz constant for the OVF in specific problem instances. While the algorithmic consequences are less dramatic than those of, say, unimodality, efficient derivative-free global optimization methods exist in the literature for minimizing univariate Lipschitz-continuous functions (see, e.g., Hansen et al., 1992).

7. Applications We next provide some examples of how our results can be utilized to provide tractable formulations for specific applied problems. In Section 7.1 we briefly describe a class of network problems that feature decision-dependent nominal probabilities and decision-independent nominal realizations. Noting that problems with this type of endogenous uncertainty have received significant attention in the recent literature, we apply our DRO approach to them, using the natural choice of a total variation-based ball as the ambiguity set. We use the ideas of Sections 6.1.1 and 6.1.2, in combination with some recent developments from the literature, to obtain mixed-integer linear programming (MILP) formulations; some technical details and additional discussions are relegated to the Appendix. Then in Section 7.2 we turn our attention to a novel distributionally robust machine scheduling problem with discrete control options, featuring decision-independent nominal probabilities and decision-dependent nominal realizations. We again obtain MILP formulations, this time making use of the developments in Sections 6.1.3 and 6.1.4. These formulations will also be the focus of the computational study in Section 8.

7.1 Network models with independent component failures We first examine an important general class of optimization problems with decision-dependent uncertainty, which includes stochastic network reliability and network interdiction problems. Problems in this class, as discussed in Haus et al. (2017), typically feature an underlying system represented by a graph whose edges and/or nodes are subject to random failures. Accordingly, possible states of the system can be represented by binary vectors, where each coordinate corresponds to a network component. By convention, a value of 1 signifies the survival of the component, while a value of 0 indicates failure. These binary vectors can then be viewed as scenarios, with corresponding scenario probabilities determined by the (independent) survival probabilities of system components. Decision-dependent uncertainty arises when design decisions can be made to affect these survival probabilities, with the aim of improving the post-failure performance of the system. Performance is usually evaluated using an outcome function that quantifies network properties such as connectivity or shortest path lengths. An important assumption is that the outcome can be expressed solely as a function of the state of the system. Under this assumption, the objective function of a decision maker who is interested in minimizing the expected outcome (or, more generally, a risk measure of the outcome) will only depend on decisions through their effect on scenario probabilities. Before introducing a formal description of models in this problem class, we present two important applications, which were the primary motivation behind the developments in this section.

Stochastic pre-disaster investment planning problem (SPIPP). As discussed in the introduction, this problem (originally proposed by Peeta et al., 2010) has been receiving significant attention in the recent literature. The problem models a transportation network as an undirected graph, where the edges correspond to highway links, and the edge lengths correspond to traversal costs. In the event of a disaster (such as an earthquake), the links are subject to random failures, which are assumed to be independent. The failure probability of each link can be reduced by making an investment to strengthen it, and the goal is to use a limited budget to improve the post-disaster connectivity of the network. In the simplest case connectivity is quantified as the length of the shortest path between an origin-destination (O-D) pair. To obtain a more nuanced measure of connectivity, one can also consider an appropriately weighted sum of shortest path lengths between multiple O-D pairs.

Stochastic network interdiction problem (SNIP). In this problem, introduced by Cormican et al. (1998), a defender attempts to block arcs in a capacitated network, using a limited budget, in order to diminish an attacker’s ability to perform a task such as the distribution of nuclear weapons or illegal drugs. Blocking attempts are assumed to randomly succeed or fail in a binary fashion, independently for each arc. The attacker’s goal is then to maximize a flow through the remaining network.

7.1.1 Description of the base model Let us consider a system that consists of L components, which—keeping in mind the motivating pre-disaster planning problem—we will refer to as ‘links’. Each link $\ell \in [L]$ has a baseline survival probability of $\sigma_\ell^0 \in [0, 1]$, which increases to $\sigma_\ell^1 \in [\sigma_\ell^0, 1]$ if the link is strengthened (or “reinforced”). We introduce the binary decision variables $\mathbf{x} \in \{0, 1\}^L$, where x_ℓ takes value 1 if and only if an investment is made to strengthen link $\ell \in [L]$. For the sake of our discussions we assume that the set $\mathcal{X} \subset \{0, 1\}^L$ of feasible decisions is defined by a set of linear inequalities (in practice typically by a single budget constraint), although our formulations remain valid for the general case.

The post-failure state of the system can be described via a binary vector of length L , whose ℓ th component takes value 1 if and only if link $\ell \in [L]$ survives. In order to be consistent with the formalism of the preceding sections, we introduce the following notation. Observing that there are $n = 2^L$ possible system states (“scenarios”), let $\{\boldsymbol{\xi}^1, \dots, \boldsymbol{\xi}^n\} = \{0, 1\}^L$ be a list of all L -dimensional binary vectors. For $i \in [n]$, $\ell \in [L]$ we will interpret the equality $\xi_\ell^i = 1$ as “link ℓ survives in scenario i ”, and conversely interpret the equality $\xi_\ell^i = 0$ as “link ℓ fails in scenario i ”. Given the reinforcement decision $x_\ell \in \{0, 1\}$, the survival probability of link $\ell \in [L]$ is $(1 - x_\ell)\sigma_\ell^0 + x_\ell\sigma_\ell^1$, while its failure probability is $(1 - x_\ell)(1 - \sigma_\ell^0) + x_\ell(1 - \sigma_\ell^1)$. Therefore, under the assumption that link failures are independent, the decision-dependent probability distribution $\mathbb{P}(\mathbf{x})$ is given by the following formula for the probability of scenario $i \in [n]$:

$$p^i(\mathbf{x}) = \prod_{\ell \in [L]: \xi_\ell^i = 1} [(1 - x_\ell)\sigma_\ell^0 + x_\ell\sigma_\ell^1] \prod_{\ell \in [L]: \xi_\ell^i = 0} [(1 - x_\ell)(1 - \sigma_\ell^0) + x_\ell(1 - \sigma_\ell^1)]. \quad (34)$$

Finally, recalling that under our assumption the outcome function is decision-independent, let us denote its value in scenario $i \in [n]$ by G^i . We can then formulate the risk-neutral version of our base problem as

$$\min_{\mathbf{x} \in \mathcal{X}} \sum_{i \in [n]} p^i(\mathbf{x}) G^i. \quad (35)$$

In Appendices A.1 and A.2 we describe in detail how to specify the components of the above general model for our motivating applications, SPIPP and SNIP.

The base problem (35) presents two major challenges. First, the probability expression (34) is a highly non-linear function of the decision variables. Second, the number of scenarios is exponential in terms of the number of links, which leads to prohibitively large formulations. In the next sections we discuss two recently developed and closely related techniques that help tackle these issues.

7.1.2 Distribution shaping Peeta et al. (2010) have solved an approximate version of SPIPP, which is obtained by replacing the highly polynomial objective of (35) with a multilinear function. Significant effort has been made in the recent literature to develop efficient solution methods that improve on this rough approximation (Flach and Poggi, 2010; Laumanns et al., 2014; Schichl and Sellmann, 2015; Haus et al., 2017). Here, with the goal of a straightforward and intuitive presentation in mind, we closely follow the work of Laumanns et al. (2014), who provide efficient exact methods that will also remain applicable in our DRO context. A key element of their approach is the technique of *distribution shaping*, which enables one to characterize the decision-dependent scenario probabilities via a set of linear constraints.

Given a decision vector $\mathbf{x} \in \mathcal{X}$, we introduce for all $\ell \in [L]$ the truncated vector $\check{\mathbf{x}}^\ell$ given by $\check{x}_j^\ell = x_j$ if $1 \leq j \leq \ell$, and by $\check{x}_j^\ell = 0$ if $\ell < j \leq L$. Let us denote the corresponding scenario probabilities by $\pi_\ell^i = p^i(\check{\mathbf{x}}^\ell)$, and note that the trivial equality $\mathbf{x} = \check{\mathbf{x}}^L$ implies $p^i(\mathbf{x}) = \pi_L^i$ for all $i \in [n]$. The key observation behind scenario shaping is that, in accordance with Bayes’ rule, the probability measures defined by successive truncations $\check{\mathbf{x}}^{\ell-1}$ and $\check{\mathbf{x}}^\ell$ have a linearly expressible relationship (see Laumanns et al.,

2014, for the exact details). We can therefore formulate (35) as the MILP

$$\min \sum_{i \in [n]} \pi_L^i G^i \tag{36a}$$

$$\text{s.t. } \pi_\ell^i \leq \frac{\sigma_\ell^1}{\sigma_\ell^0} \pi_{\ell-1}^i + 1 - x_\ell, \quad \forall \ell \in [L], i \in [n] : \xi_\ell^i = 1 \tag{36b}$$

$$\pi_\ell^i \leq \frac{1 - \sigma_\ell^1}{1 - \sigma_\ell^0} \pi_{\ell-1}^i + 1 - x_\ell, \quad \forall \ell \in [L], i \in [n] : \xi_\ell^i = 0 \tag{36c}$$

$$\pi_\ell^i \leq \pi_{\ell-1}^i + x_\ell, \quad \forall \ell \in [L], i \in [n] \tag{36d}$$

$$\sum_{i \in [n]} \pi_\ell^i = 1, \quad \ell \in [L] \tag{36e}$$

$$\boldsymbol{\pi} \in [0, 1]^{n \times L}, \tag{36f}$$

$$\mathbf{x} \in \mathcal{X}, \tag{36g}$$

where $\pi_0^i = \prod_{\ell: \xi_\ell^i=1} \sigma_\ell^0 \prod_{\ell: \xi_\ell^i=0} (1 - \sigma_\ell^0)$ denotes the baseline probability of scenario $i \in [n]$. It is straightforward to verify that, due to the distribution shaping relations (36b)–(36f), the defining equalities $\pi_\ell^i = p^i(\tilde{\mathbf{x}}^\ell)$ will hold for all $\ell \in [L], i \in [n]$.

7.1.3 Scenario bundling We now turn our attention to the crucial problem of reducing the number of scenarios. Sampling methods are often used for this purpose; we will briefly discuss their applicability in our setting at the end of Appendix A.7, but for the moment we keep our focus on exact methods. Taking advantage of the crucial assumption that outcomes are decision-independent, it is possible to group together scenarios that lead to the same outcome. Two important goals need to be kept in mind when attempting to implement this idea. First, suitable scenario groups should be identified without having to perform an excessive number of outcome function evaluations. Second, the highly effective distribution shaping approach should remain applicable. To accomplish these goals, one can follow the scenario bundling technique described in Laumanns et al. (2014), which was demonstrated to work well on both SPIPP and SNIP. Appendix A.5 explains in detail how to apply this technique in our setting.

7.1.4 DRO problem variants We now proceed to introduce DRO variants of the underlying problem (35). To account for uncertainty about the scenario probabilities, we will use a discrete EMD ball of type (BALL-D) with radius κ around the decision-dependent nominal distribution $\mathbb{P}(\mathbf{x})$ as our ambiguity set. Noting that the assumptions of Section 6.1.1 are satisfied, we can formulate the arising instance of (DRO-RND) as the following simple variant of (30):

$$\min \sum_{i \in [n]} p^i(\mathbf{x}) v^i + \kappa \tau \tag{37a}$$

$$\text{s.t. } v^i \geq G^i - \delta^{ij} \tau, \quad \forall i, j \in [n] \tag{37b}$$

$$\mathbf{v} \in \mathbb{R}^n, \tau \in \mathbb{R}_+, \mathbf{x} \in \mathcal{X}. \tag{37c}$$

The probabilities $p^i(\mathbf{x})$ in the objective are given by the highly non-linear formula (34). As seen in Section 7.1.2, distribution shaping allows us to eliminate this non-linearity. To this end, we can simply introduce the auxiliary variables π_ℓ^i for $\ell \in [L], i \in [n]$, add the corresponding defining constraints (36b)–(36f) to the problem (37), and replace $p^i(\mathbf{x})$ in (37a) with π_L^i . While the resulting formulation will be valid for any choice of the “scenario distance” δ , from now on we will restrict ourselves to the special case where δ is the discrete metric. Accordingly, ambiguity sets will be based on the total variation distance, which has the downside of ignoring potentially meaningful information about degrees of similarity between various scenarios. However, as outlined in Appendix A.6, this choice of δ will eventually enable us to use scenario

bundling methods for our DRO problems in a straightforward fashion.

7.1.5 MILP formulations for the total variation ball When δ is the discrete metric, we can combine distribution shaping with the ideas of Section 6.1.2 to reformulate (37) as

$$\min \sum_{i \in [n]} \pi_L^i v^i + \kappa \tau \quad (38a)$$

$$\text{s.t. } v^i \geq G^i, \quad \forall i \in [n] \quad (38b)$$

$$v^i \geq G^+ - \tau, \quad \forall i \in [n] \quad (38c)$$

$$(36b)-(36f), \quad (38d)$$

$$\mathbf{v} \in \mathbb{R}^n, \tau \in \mathbb{R}_+, \mathbf{x} \in \mathcal{X}. \quad (38e)$$

The constraints of the above mixed-integer program are linear, but the summation in the objective function features non-convex quadratic terms. Our next goal is to linearize these terms. According to Example 5.2, as we are using a total variation-based ambiguity set, the optimum of (38) is a convex combination of the worst-case outcome G^+ and the nominal $\text{CVaR}_\kappa(G)$. Moreover, this optimum is attained when we have $\tau = G^+ - \text{VaR}_\kappa(G)$. Keeping in mind the representation (5) of CVaR as a minimum taken over the finite set of realizations (interpreted as the possible VaR values), we can now solve our DRO problem as follows. If we fix the value of the variable τ in (38), then we can also fix the values v^i by setting $v^i = \max\{G^i, G^+ - \tau\}$ for all $i \in [n]$, and the problem becomes an MILP. Let us then separately solve the n different MILPs obtained from (38) by fixing $\tau = G^+ - G^j$ for some $j \in [n]$. The MILP with the smallest optimum value will also provide the solution for (38). While the solution of the n MILPs can naturally be parallelized, in Appendix A.3 we also present an alternative disjunction-based approach, which combines these subproblems into a single MILP formulation that can be significantly strengthened via the reformulation-linearization technique (Sherali and Adams, 1994; Sherali et al., 1998).

Throughout this section we considered the DRO version of the risk-neutral underlying problem (35), which was observed to be equivalent to a (non-DRO) risk-averse variant of the underlying problem. Notably, we did not consider the DRO version of a risk-averse underlying problem. The reason behind this omission is the following: Let us assume that $\kappa \leq 1 - \alpha$ holds. When δ is the discrete metric, we know from Section 5.2.2 that the κ -robustification of CVaR_α is a convex combination of the worst case outcome and $\text{CVaR}_{\alpha+\kappa}$. When outcomes are scenario-independent, the worst-case outcome is constant. Therefore in this case minimizing the robustified risk measure $\text{CVaR}_\alpha^\kappa$ is equivalent to minimizing $\text{CVaR}_{\alpha+\kappa}$, which is in turn equivalent to minimizing the robustified expectation $\text{CVaR}_0^{\alpha+\kappa} = \mathbb{E}^{\alpha+\kappa}$. On a related note, while the MILP formulations in Appendix A.3 are derived from the starting point of robustifying a risk-neutral underlying problem, the above arguments show that equivalent formulations can be obtained starting from a non-DRO CVaR -minimization problem. Appendix A.4 explores this perspective in more detail.

All of the formulations we encountered so far rely on a full (exponential sized) scenario set, which usually makes it impossible to solve problem instances of practical interest. In Appendix A.6 we outline how the concept of scenario bundling (which was introduced in Section 7.1.3 and further detailed in Appendix A.5) can be applied in a DRO context to ameliorate this situation by reducing problem sizes. This is followed in Appendix A.7 by a discussion of the limitations to the bundling approach.

7.2 Stochastic Single-Machine Scheduling We now introduce a new risk-averse machine scheduling problem with controllable stochastic processing times. Risk-averse developments in the scheduling literature are relatively recent. We refer to Pinedo (2008) for an overview of conventional

stochastic machine scheduling, and to Atakan et al. (2017) for a review on risk-averse approaches to single-machine scheduling (including traditional robust optimization). The studies by Hradovich et al. (2018) and Chang et al. (2017), along with the aforementioned Atakan et al. (2017), are among the very limited number of studies in the field that involve the use of a risk measure, and therefore are of particular relevance to our work. There has also been some recent interest in adopting a DRO approach to hedge against distributional uncertainty in processing times (see, e.g., Chang et al., 2017; Zhang et al., 2018; Niu et al., 2019). However, to the best of our knowledge the present work is the first attempt to apply these risk-averse approaches to scheduling problems with endogenous uncertainty.

Let us consider a simple single-machine scheduling problem featuring L jobs, with processing times ξ_l and importance weights w_l for $l \in [L]$. Schedules will be evaluated based on the *total weighted completion time* (TWCT) of the jobs, which is a widely used performance measure (see, e.g., Pinedo, 2008). It will be helpful to assume that the TWCT is interpreted on a monetary scale; this can be accomplished by appropriately scaling the weights w_l .

We are primarily interested in the case where the processing times are stochastic, and can be affected by control decisions. Accordingly, let $(\Omega, \mathcal{A}, \mathbb{P})$ be an arbitrary (not necessarily finite) probability space, and let us introduce the mapping $\xi : \mathcal{U} \rightarrow \mathcal{L}^L(\Omega, \mathcal{A})$. Here \mathcal{U} is the set of feasible control decisions, and $\xi_l(\mathbf{u}) \in \mathcal{L}^1(\Omega, \mathcal{A})$ is the random processing time of job $l \in [L]$ given decision \mathbf{u} . In addition, we denote the cost associated with decision \mathbf{u} by $h(\mathbf{u})$; the cost mapping $h : \mathcal{U} \rightarrow \mathbb{R}$ is often chosen to be linear.

In the deterministic scheduling literature a wide variety of schemes have been proposed to control processing times, see, e.g., Shabtay and Steiner (2007). We will now adapt two important models of control to our stochastic setting.

- *Linearly compressible processing times* take the form $\xi_l(\mathbf{u}) = \hat{\xi}_l - \hat{a}_l u_l$, where $\hat{\xi}_l \in \mathcal{L}^1(\Omega, \mathcal{A})$ is the *baseline* random processing time of job $l \in [L]$, and $\hat{a}_l \in \mathcal{L}^1(\Omega, \mathcal{A})$ is the corresponding stochastic compression rate. Feasible control decisions will then constitute a set

$$\mathcal{U} \subset \left\{ \mathbf{u} \in \mathbb{R}^L : 0 \leq u_l \leq \text{ess inf} \frac{\hat{\xi}_l}{\hat{a}_l} \quad \forall l \in [L] \right\}.$$

EXAMPLE 7.1 *In the case $\hat{a}_l = \hat{\xi}_l$ processing times are given by $\xi_l(\mathbf{u}) = (1 - u_l)\hat{\xi}_l$, and the decision $u_l \in [0, 1]$ can be interpreted as a proportional decrease in the processing time of job l .*

- *Control with discrete resources*: A finite set of T control options is available for every job, and selecting option $t \in [T]$ for job $l \in [L]$ leads to a random processing time of $\hat{\xi}_{tl}$. Let us introduce the binary decision variables u_{tl} for $t \in [T]$, $l \in [L]$, that take value 1 if and only if control option t is selected for job l . Then the processing time of job l is given by $\xi_l(\mathbf{u}) = \sum_{t \in [T]} u_{tl} \hat{\xi}_{tl}$ for $l \in [L]$, and the feasible control decisions constitute a set

$$\mathcal{U} \subset \left\{ \mathbf{u} \in \{0, 1\}^{T \times L} : \sum_{t \in [T]} u_{tl} = 1 \quad \forall l \in [L] \right\}. \quad (39)$$

EXAMPLE 7.2 *Assume that for each job the decision maker can choose to apply one of T possible linear compression rates, given by $\hat{a}_{tl} \in [0, 1]$ for $t \in [T]$, $l \in [L]$, and let us denote the corresponding speedup factors by $a_{tl} = 1 - \hat{a}_{tl}$. The controllable processing times then take the form $\xi_l^i(\mathbf{u}) = \hat{\xi}_l^i \left(1 - \sum_{t \in [T]} \hat{a}_{tl} u_{tl} \right) = \hat{\xi}_l^i \sum_{t \in [T]} a_{tl} u_{tl}$, where $\hat{\xi}_l^i$ again denotes the baseline random processing time.*

It is easy to verify that the comonotonicity condition discussed in Section 6.1.4 holds both for Example 7.1 and for Example 7.2.

We next describe the sequencing aspect of our scheduling problems using the well-known *linear ordering formulation*. Let us introduce the binary decision variables θ_{kl} for $k, l \in [L]$ that take value 1 if job k precedes job l in the processing sequence, and take value 0 otherwise. Then the set \mathcal{T} of feasible scheduling decisions consists of the binary matrices $\boldsymbol{\theta} \in \{0, 1\}^{L \times L}$ that satisfy the system

$$\theta_{ll} = 1, \quad \forall l \in [L] \quad (40a)$$

$$\theta_{kl} + \theta_{lk} = 1, \quad \forall k, l \in [L] : k < l \quad (40b)$$

$$\theta_{kl} + \theta_{lh} + \theta_{hk} \leq 2, \quad \forall k, l, h \in [L] : k < l < h. \quad (40c)$$

Here constraints (40a) express the convention that each job is considered to precede itself, constraints (40b) ensure that no job simultaneously precedes and succeeds a different job, while constraints (40c) prevent cyclic subsequences of length three.

If we assume zero release dates for all jobs, then the completion time of job $l \in [L]$ is given by $\sum_{k \in [L]} \xi_k(\mathbf{u})\theta_{kl}$. Introducing the matrix $\Theta = (\theta_{kl})_{k, l \in [L]}$, we can express the TWCT objective as

$$\sum_{l \in [L]} w_l \sum_{k \in [L]} \xi_k(\mathbf{u})\theta_{kl} = \sum_{k \in [L]} \sum_{l \in [L]} \xi_k(\mathbf{u})\theta_{kl}w_l = \boldsymbol{\xi}^\top(\mathbf{u})\Theta\mathbf{w}.$$

The risk-averse version of our stochastic single-machine scheduling problem can now be formulated as

$$\min_{(\boldsymbol{\theta}, \mathbf{u}) \in \mathcal{T} \times \mathcal{U}} h(\mathbf{u}) + \rho(\boldsymbol{\xi}^\top(\mathbf{u})\Theta\mathbf{w}), \quad (41)$$

where ρ is a law-invariant coherent risk measure. We next proceed to examine DRO variants of this underlying problem.

7.2.1 Continuous Wasserstein balls Let us first consider the case when processing times can take their values from a continuous spectrum and are subject to ambiguity, with a continuous Wasserstein- p ball of radius κ as the ambiguity set. As outlined in Section 4, the DRO variant of the underlying risk-averse problem (41) then takes the form

$$\min_{(\boldsymbol{\theta}, \mathbf{u}) \in \mathcal{T} \times \mathcal{U}} h(\mathbf{u}) + \sup_{\boldsymbol{\zeta} \in \mathcal{B}_{\delta^p, \kappa, p}^{\mathbb{P}}(\boldsymbol{\xi}(\mathbf{u}))} \rho(\boldsymbol{\zeta}^\top \Theta \mathbf{w}). \quad (42)$$

If the risk measure ρ is well-behaved with some factor C , then it immediately follows from Proposition 4.2 that the problem (42) can be equivalently reformulated as

$$\min_{(\boldsymbol{\theta}, \mathbf{u}) \in \mathcal{T} \times \mathcal{U}} h(\mathbf{u}) + \rho(\boldsymbol{\xi}^\top(\mathbf{u})\Theta\mathbf{w}) + C\kappa\|\Theta\mathbf{w}\|_q. \quad (43)$$

The only difference between this formulation and the underlying problem (41) is the additional robustification term $C\kappa\|\Theta\mathbf{w}\|_q$, which, due to the convexity of the q -norm, is a convex function of the sequencing variables θ_{kl} . The example below shows that this term can affect the optimal schedule, even when the underlying scheduling problem is deterministic with no compression decisions.

EXAMPLE 7.3 Consider the following deterministic instance of the scheduling problem introduced in Section 7.2. There are two jobs (Job 1 and Job 2) with respective weights 2 and 3, and respective non-compressible processing times 21 and 32. Scheduling Job 1 before Job 2 (“schedule $1 \prec 2$ ”) leads to a TWCT of 201, which is superior to the TWCT of 202 for schedule $2 \prec 1$. However, in the DRO version of the problem where the ambiguity set for the processing time vector is the 2-norm ball $B_4^2((21, 32)^\top)$ of radius 4 around the nominal values, the robustified TWCT for schedule $1 \prec 2$ becomes (approximately)

224.32, which is inferior to the robustified TWCT of 223.54 for schedule $2 \prec 1$. We note that, in accordance with our observations at the end of Section 4, the same results will hold for any risk-neutral stochastic version of the problem with expected nominal processing times 21 and 32, and a continuous Wasserstein-2 ball of radius 4 as the ambiguity set.

However, in certain settings the underlying problem (41) and the robustified problem (43) are guaranteed to have the same solution. In the case $p = 1$ it is easy to verify that the robustification term is always equal to the constant $C\kappa \sum_{l \in [L]} w_l$, and thus has no impact on the optimal solution. Along similar lines, if we replace the total weighted completion time in the objective function by the (unweighted) total completion time, i.e., if we set $w_l = 1$ for all $l \in [L]$, then the robustification term becomes $C\kappa \left(\sum_{l=1}^L l^q \right)^{\frac{1}{q}}$, which again does not depend on the decision variables.

To finish this subsection, we briefly discuss two further cases when (42), our DRO problem with decision-dependent uncertainty, reduces to a more familiar type of problem.

OBSERVATION 7.1 *Assume that the processing times are compressed in a scenario-independent fashion, i.e., that we have $\xi(\mathbf{u}) = \hat{\xi} - \hat{\mathbf{a}}(\mathbf{u})$ for some baseline random processing time vector $\hat{\xi} \in \mathcal{L}^m(\Omega, \mathcal{A})$, and a deterministic compression mapping $\hat{\mathbf{a}} : \mathcal{U} \rightarrow \mathbb{R}^L$. We can then rewrite (42) as a “traditional” DRO problem without decision-dependent uncertainty. More precisely, it is easy to verify that a random vector $[\mathbb{B}, \zeta]$ belongs to the decision-dependent ambiguity set $\mathcal{B}_{\delta, \kappa}^{\mathbb{P}}(\xi(\mathbf{u}))$ if and only if it is of the form $\zeta = \hat{\zeta} - \hat{\mathbf{a}}(\mathbf{u})$, where $\hat{\zeta}$ belongs to the decision-independent ambiguity set $\mathcal{B}_{\delta, \kappa}^{\mathbb{P}}(\hat{\xi})$.*

OBSERVATION 7.2 *Let us examine the risk-neutral case, where we have $\rho = \mathbb{E}$. In Section 4 we established that, for a general class of problems, the arising DRO instance (DRO-RNC) is equivalent to the deterministic problem (11). In our scheduling context this result leads to the following deterministic reformulation of (42):*

$$\min_{(\boldsymbol{\theta}, \mathbf{u}) \in \mathcal{T} \times \mathcal{U}} h(\mathbf{u}) + \bar{\xi}^\top(\mathbf{u})\Theta\mathbf{w} + C\kappa\|\Theta\mathbf{w}\|_q,$$

where the operator $\bar{\xi} : \mathcal{U} \rightarrow \mathbb{R}^L$ gives the decision-dependent expected processing time vector $\bar{\xi}(\mathbf{u}) = \mathbb{E}(\xi(\mathbf{u}))$ for $\mathbf{u} \in \mathcal{U}$.

7.2.2 Discrete EMD balls We now consider the case when processing times can take their values from some discrete set, and accordingly the ambiguity set is a discrete EMD ball of type (BALL-D). We begin with a simple example that illustrates the impact of the DRO approach on optimal scheduling decisions, and also shows that equivalent underlying problems can have non-equivalent robustifications.

EXAMPLE 7.4 *We consider an instance of the stochastic scheduling problem introduced in Section 7.2 with two scenarios, which in the nominal distribution \mathbb{P} both have probability 0.5. There are two jobs, with non-compressible nominal processing times as follows: $\xi_1^1 = 2$, $\xi_1^2 = 4$ for Job 1, and $\xi_2^1 = \xi_2^2 = 6$ for Job 2. We take a risk-neutral approach, and aim to minimize the expected TWCT when the respective weights of Jobs 1 and 2 are 20 and 39. Scheduling Job 1 before Job 2 (“schedule $1 \prec 2$ ”) leads to an expected TWCT of 411, which is superior to the expected TWCT of 414 for schedule $2 \prec 1$. On the other hand, in the DRO variant of the problem where the ambiguity set for the random processing time vector is the discrete Wasserstein-1 ball $\mathcal{B}_{\delta^{\xi}, 0.1}^{\xi}(\mathbb{P})$ of radius 0.1, the robustified expected TWCT for schedule $1 \prec 2$ is 416.9, which is inferior to the robustified expected TWCT of 416 for schedule $2 \prec 1$.*

To obtain the deterministic counterparts of these problems, we need to replace $[\mathbb{P}, \xi]$ with the trivial distribution where the processing time of Job 1 is changed to its expected value of 3 in both scenarios. As discussed at the end of Section 4, in the risk-neutral case the underlying (i.e., non-robustified) stochastic

problem is equivalent to its deterministic counterpart. However, DRO variants of these two equivalent problems are no longer equivalent. More precisely, since any EMD ball of type (BALL-D) around a deterministic nominal vector is trivial (i.e., it contains only its center), all DRO variants of the deterministic problem are equivalent to the underlying non-robustified one. However, as we have just seen, robustifying the original stochastic underlying problem can affect the optimal schedule.

In this section we study the following DRO variant of the risk-averse scheduling problem (41):

$$\min_{(\boldsymbol{\theta}, \mathbf{u}) \in \mathcal{T} \times \mathcal{U}} h(\mathbf{u}) + \sup_{\mathbb{Q} \in \mathcal{B}_{\delta, \kappa}^{\xi(\mathbf{u})}(\mathbb{P})} \rho(\mathbf{w}^\top \Theta^\top \boldsymbol{\xi}(\mathbf{u})). \quad (44)$$

For the case $\rho = \text{CVaR}_\alpha$ we can adapt the general formulation (29) to equivalently express our problem (44) as

$$\min \quad h(\mathbf{u}) + \eta + \frac{1}{1-\alpha} \sum_{i \in [n]} p^i v^i + \frac{1}{1-\alpha} \kappa \tau \quad (45a)$$

$$\text{s.t.} \quad v^i \geq \sum_{l \in [L]} w_l \sum_{k \in [L]} \xi_k^j(\mathbf{u}) \theta_{kl} - \eta - \delta^{ij} \tau, \quad \forall i, j \in [n] \quad (45b)$$

$$\delta^{ij} = \delta(\mathbf{w}^\top \Theta^\top \boldsymbol{\xi}^i(\mathbf{u}), \mathbf{w}^\top \Theta^\top \boldsymbol{\xi}^j(\mathbf{u})), \quad \forall i, j \in [n] \quad (45c)$$

$$(\boldsymbol{\theta}, \mathbf{u}) \in \mathcal{T} \times \mathcal{U}, \quad \mathbf{v} \in \mathbb{R}_+^n, \quad \tau \geq 0. \quad (45d)$$

REMARK 7.1 *In order to keep the presentation simple, we implicitly assumed that the costs associated with our decisions are deterministic. For risk-neutral problems this assumption is without loss of generality, because stochastic costs can be equivalently replaced with their expected values. While this is no longer the case in a risk-averse context, we can easily adapt our formulations to a setting with stochastic costs. Denoting the cost of decision $\mathbf{u} \in \mathcal{U}$ under scenario $i \in [n]$ by $h^i(\mathbf{u})$, we can simply remove $h(\mathbf{u})$ from the objective function in (45a), and instead incorporate the costs into the random outcome mapping by adding the term $h^i(\mathbf{u})$ to the right-hand side of constraint (45b).*

The formulation (45) is generally a very challenging mixed-integer non-linear program (MINLP), due in part to the non-convex quadratic terms in constraints (45b), and in part to the possible non-linearity in constraints (45c). We next provide potentially tractable forms of this problem for the case of control with discrete resources, when using a Wasserstein-1 ambiguity set.

Let us assume that the processing time of job $l \in [L]$ is $\xi_l(\mathbf{u}) = \sum_{t \in [T]} u_{tl} \hat{\xi}_{tl}$ for $l \in [L]$, where the set \mathcal{U} of feasible control decisions is given as in (39). Then in constraints (45b) we can rewrite $\xi_k^j(\mathbf{u}) \theta_{kl}$ as $\sum_{t \in [T]} \hat{\xi}_{tl}^j u_{tk} \theta_{kl}$, and use McCormick envelopes (McCormick, 1976) to linearize the arising quadratic terms $u_{tk} \theta_{kl}$. In addition, when δ is the 1-norm distance, we can express the δ^{ij} values as in (32), and incorporate them into our optimization problem via mixed-integer big-M constraints as in (33). The problem (45) then takes the following form:

$$\min \quad h(\mathbf{u}) + \eta + \frac{1}{1-\alpha} \sum_{i \in [n]} p^i v^i + \frac{1}{1-\alpha} \kappa \tau \quad (46a)$$

$$\text{s.t.} \quad v^i \geq \sum_{l \in [L]} \sum_{k \in [L]} \sum_{t \in [T]} w_l \hat{\xi}_{tk}^j z_{tkl} - \eta - \sum_{l \in [L]} \nu_l^{ij} \tau, \quad \forall i, j \in [n] \quad (46b)$$

$$z_{tkl} \leq u_{tk}, \quad \forall t \in [T], k, l \in [L] \quad (46c)$$

$$z_{tkl} \leq \theta_{kl}, \quad \forall t \in [T], k, l \in [L] \quad (46d)$$

$$z_{tkl} \geq u_{tk} + \theta_{kl} - 1, \quad \forall t \in [T], k, l \in [L] \quad (46e)$$

$$\nu_l^{ij} \leq \xi_l^i(\mathbf{u}) - \xi_l^j(\mathbf{u}) + M\lambda_l^{ij}, \quad \forall i, j \in [n], l \in [L] \quad (46f)$$

$$\nu_l^{ij} \leq -\xi_l^i(\mathbf{u}) + \xi_l^j(\mathbf{u}) + M(1 - \lambda_l^{ij}), \quad \forall i, j \in [n], l \in [L] \quad (46g)$$

$$\boldsymbol{\lambda} \in \{0, 1\}^{n \times n \times L}, \quad \boldsymbol{\nu} \in \mathbb{R}_+^{n \times n \times L}, \quad (46h)$$

$$(\boldsymbol{\theta}, \mathbf{u}) \in \mathcal{T} \times \mathcal{U}, \quad \mathbf{v} \in \mathbb{R}_+^n, \quad \tau \geq 0, \quad \mathbf{z} \in [0, 1]^{T \times L \times L}. \quad (46i)$$

Keeping in mind that the auxiliary variables \mathbf{z} are non-negative, the constraints (46c)–(46e) ensure that $z_{tkl} = u_{tk}\theta_{kl}$ holds for all $t \in [T]$, $k, l \in [L]$. Also, since the boundedness condition established at the start of Section 6.1.3 trivially holds in the case of discrete control decisions, for sufficiently high values of the parameter M the constraints (46f)–(46h) are equivalent to the inequalities $\nu_k^{ij} \leq |\xi_l^i(\mathbf{u}) - \xi_l^j(\mathbf{u})|$ for all $i, j \in [n]$, $l \in [L]$.

REMARK 7.2 *Due to the symmetry of the absolute value function we can assume that $\nu_k^{ij} = \nu_k^{ji}$ and $\lambda_k^{ij} = \lambda_k^{ji}$ hold for all $i, j \in [n]$ and $l \in [L]$. Taking into account the trivial observation $\nu_l^{ii} = |\xi_l^i(\mathbf{u}) - \xi_l^i(\mathbf{u})| = 0$, it is clearly sufficient to only define the variables ν_k^{ij} and λ_k^{ij} for indices that satisfy $i < j$. Our MILP formulations detailed in Section 8.3.1 all take advantage of the resulting reduction in the number of auxiliary variables.*

We mention that the reformulation-linearization technique yields the valid inequalities $\sum_{t \in [T]} z_{tkl} = \theta_{kl}$ for all $k, l \in [L]$, which can be added to strengthen the formulation (46). Recalling our discussions from Section 6.1.5, we also point out that, while the constraints still feature continuous non-convex quadratic terms, these terms all involve the common scalar variable τ . Therefore, if we fix the value of τ , then the above problem becomes a linearly constrained mixed-integer program.

We briefly outline another approach to eliminating these quadratic terms from our problem under certain linearity assumptions; a detailed discussion of the arising tractable MILP formulations can be found in Section 8.3.1. It is possible to linearize terms of the form $\nu_l^{ij}\tau$ by introducing auxiliary variables to represent them, and employing McCormick envelopes. Alternatively, as we have seen in Section 6.1.4, under appropriate comonotonicity assumptions it is possible to significantly simplify (46) by eliminating the auxiliary variables $\boldsymbol{\nu}$ and $\boldsymbol{\lambda}$. We can then again apply McCormick envelopes to obtain streamlined MILP formulations.

8. Computational study The purpose of this section, where we focus on solving certain variants of the risk-averse scheduling problem, is twofold. First, we are interested in the impact of our modeling approach and parameter choices (in particular, the degree of robustness and the degree of risk aversion) on optimal decisions and objectives. Second, we aim to examine the tractability of our optimization problems.

8.1 Problem variants and data generation We generate instances of the risk-averse scheduling problem introduced in Section 7.2, keeping baseline number ranges largely in line with those presented in Atakan et al. (2017). All randomly generated parameter values are drawn from independent (integer or continuous) uniform distributions. Importance weights of jobs are chosen to be integers drawn from $[1, 3]$. Nominal probabilities of all scenarios are taken to be equal, and baseline realizations of processing times are integer values from the range $[10, 90]$, drawn separately for each job and each scenario.

We assume that two control options are available for each job. The first option, which leaves baseline processing times unchanged, has zero cost. The second option, which decreases the processing times in each scenario, has a cost directly proportional to the nominal expectation of the decrease. Our primary focus is on problems that exhibit the comonotone structure outlined in Section 6.1.4, and therefore allow us to utilize streamlined MILP formulations. Let us recall the notation established in Example 7.2, and

set all factors a_{1l} to a value of 1. The factors a_{2l} , which describe the proportional time reduction effected by control decisions, are drawn from the range $[0.4, 0.6]$ for each job $l \in [L]$. We refer to the arising problems as *comonotone instances*. For the sake of comparison we also want to generate some problem instances without comonotone structure. To accomplish this, we allow the control decisions to reduce processing times by scenario-dependent factors drawn separately from $[0.2, 0.8]$ for all $l \in [L]$, $i \in [n]$, and refer to the arising problems as *non-comonotone instances*.

In scheduling problems with control decisions one typically aims to achieve good performance (represented in our case by a low TWCT) while simultaneously maintaining low control costs. To model this multi-objective perspective, we consider two formulation variants. In the *budget-constrained* variant control decisions are subject to a budget constraint, and the objective is to minimize TWCT. In the *cost minimizing* variant of our problem the control decisions are unconstrained, and the objective is to minimize a weighted combination of TWCT and control costs. When taking this approach, control decisions can become trivial if either component of the objective function dominates the other. Keeping in mind that TWCT typically exhibits quadratic growth as the number of jobs increases (in contrast to the linear growth of potential costs), if we are to generate meaningful decision problems, we need to appropriately calibrate the ratio between control costs and expected decreases in nominal processing times. In our examples we set this ratio as 10% of the expected sum of the randomly generated importance weights (taken over all L jobs), to ensure that the two components of the objective function are of comparable magnitude.

8.2 Analysis and interpretation of results In this section we study a randomly generated comonotone problem instance with $L = 15$ jobs and $n = 100$ scenarios, to evaluate the impact of key modeling parameters on optimal objective values.

We have observed in Section 5.2.5 that when the radius of the ambiguity set is sufficiently high, the DRO approach becomes equivalent to the simpler traditional robust model that optimizes worst-case performance across all scenarios. This limiting behavior already occurs at the radius $\kappa^{\max} = \sum_{l \in [L]} \max_{t \in [T]} (\hat{\xi}_{lt}^{\max} - \bar{\xi}_{lt})$, where $\hat{\xi}_{lt}^{\max} = \max_{i \in [n]} \hat{\xi}_{lt}^i$ denotes the maximal possible processing time of job $l \in [L]$ under control option $t \in [T]$, while $\bar{\xi}_{lt}$ denotes its nominal expected processing time. Accordingly, in our tables and figures we represent the radius κ on a 0–1 scale, as a proportion of κ^{\max} . Unless noted otherwise, we use certain base settings for model parameters throughout Section 8. The base radius of the ambiguity set is $\kappa = 0.01\kappa^{\max}$, which is high enough to make a non-trivial impact on decisions, while typically remaining below the threshold for the limiting behavior.

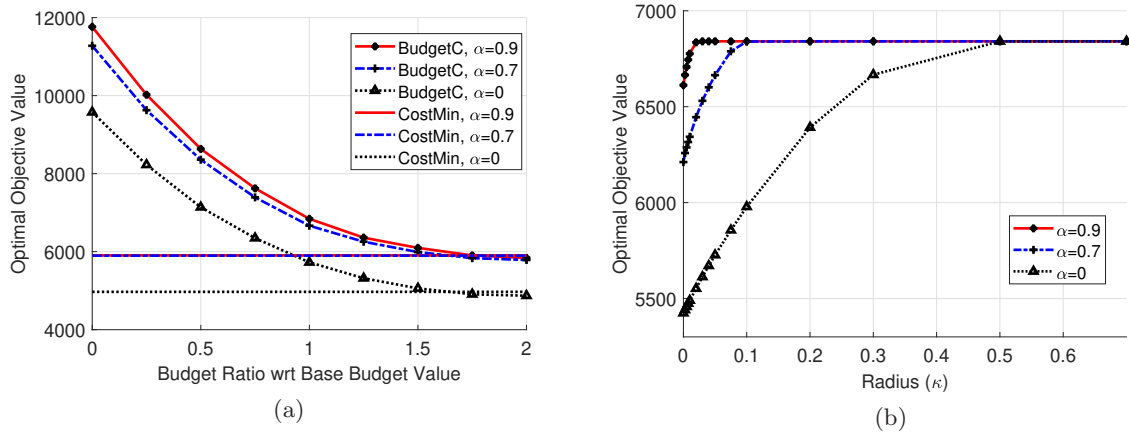
Keeping in mind that in risk-averse decision making practice high CVaR confidence levels are commonly employed, we set the base value at $\alpha = 90\%$. We consider our base budget to be half of the sum of all potential control costs. The reason behind this choice is that setting the budget closer to either the low or the high end of the spectrum could potentially reduce the combinatorial complexity of control decisions. When the available budget is varied, we represent it via a *budget ratio* (BR for short), as a proportion of the base budget value. Figure 2a shows how increasing the budget lowers the optimal objective value for the budget-constrained problem, eventually outperforming the cost-minimization variant. This behavior is expected, because with a sufficiently high budget it becomes possible to reduce the processing time of every job, for zero cost in the objective.

Figure 2b shows that if the CVaR confidence level α is higher, then the limiting (worst-case robust) behavior is reached at a lower value of κ . This pattern, which will be illustrated in more detail by Figure 3b of the next section, reflects the following simple observation.

OBSERVATION 8.1 *Assume that the CVaR confidence level in the objective function is α . If the ambiguity*

set contains a distribution in which the worst-case scenario has probability at least $1 - \alpha$, then the robustified CVaR_α of the outcome equals the worst-case outcome. The supremum in the DRO problem (44) is therefore equal to the worst-case outcome, and it is attained at this “bad” distribution. Consequently, if α is high, then even an EMD ball with a small radius κ will contain a “bad” distribution, and thus lead to the worst-case robust optimal solution.

Figure 2: Optimal objective function value (robustified CVaR_α of TWCT) for varying radius and budget



For any fixed scheduling and control decisions, the objective function is typically a non-constant monotone increasing function of both the radius κ and the CVaR confidence level α . It follows that the changes we observe in optimal objective values do not automatically correspond to changes in optimal decisions. Therefore, in order to better understand the decision-making impact of our modeling approach, we will now more closely examine a specific smaller problem instance.

8.2.1 An illustrative example We consider a budget-constrained scheduling problem of comonotone type with five jobs, and ten, nominally equally likely scenarios. Baseline processing times of jobs 1–5 are given in Table 1a; the respective importance weights are 6, 6, 1, 9, and 6. Control decisions can be made to reduce the processing time of up to two chosen jobs by 50% (uniformly in every scenario).

Solving the problem for a range of κ and α values produces six distinct optimal decisions, listed in Table 1b. For example, in the risk-neutral optimal Solution A one makes the control decision to speed up the processing of jobs 1 and 4, and employs a sequence that starts with job 4 and ends with job 3. The stochastic optimal solutions we find are typically different from the deterministic ones: while Solution B is also the deterministic optimal solution for scenarios 1 and 8, none of the other solutions in Table 1b are scenario-optimal. In particular, the worst-case robust optimal Solution F is not the optimal deterministic solution for the worst-case Scenario 3. Figure 3 shows optimal objective values for various parameter values, and indicates the corresponding optimal decisions.

The similarity between Figures 2b and 3b, indicates that the behavior of objective values is consistent with what we have seen for larger instances. We now turn our attention to the decisions. Certain scheduling and control choices appear to be intuitively obvious. For example, Job 3 has both a low importance weight, and a low nominal expected processing time. Accordingly, it is the last job in every optimal schedule, and optimal control decisions never reduce its processing time. However, optimal solutions do not always conform to naive expectations. Job 4 has the highest weight by a significant margin, and its processing time is reduced in every optimal control decision. However, while it is scheduled first in both the risk-averse Solution A and the worst-case robust Solution F, somewhat surprisingly this

is not the case for Solutions B and D. This highlights the fact that our complex decision criteria lead to non-trivial variation in optimal decisions between the risk-neutral and maximally risk-averse ends of the spectrum.

From a managerial point of view both κ and α can be seen to express degrees of conservatism. For the case of total variation-based ambiguity sets this connection was made explicit in formula (19), which showed that it is possible to equivalently convert between these two parameters. In contrast, our current results indicate that when we use a Wasserstein-1 ball, κ and α play qualitatively different roles. For example, Solution D is only optimal for high values of κ and low values of α , while, conversely, solution C is only optimal for lower κ and higher α values. Our model therefore has the capacity to express a range of risk-averse preferences that would not be possible to capture by either a “purely robust” or a “purely CVaR-based” approach.

Table 1: Illustrative example

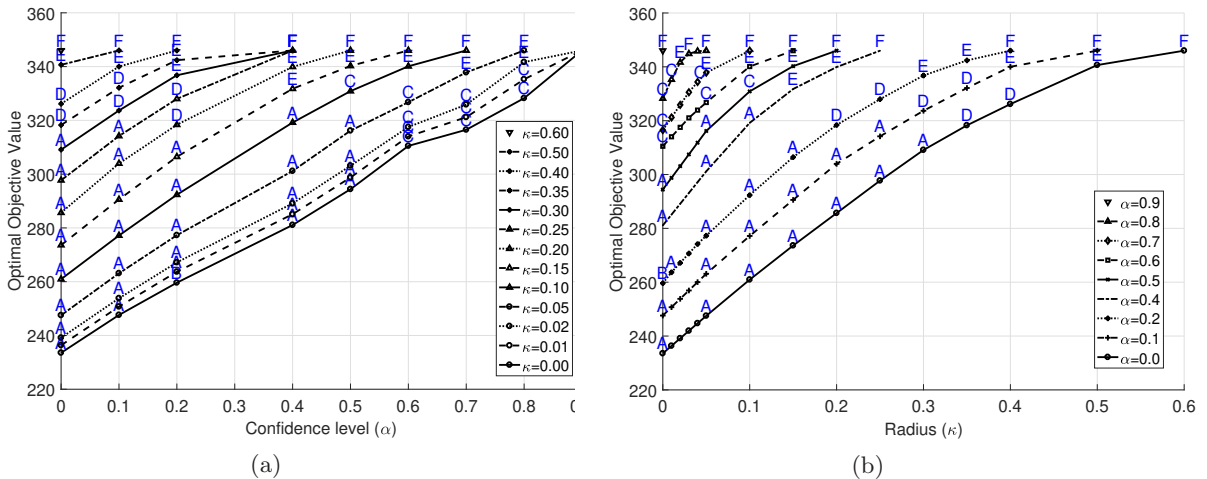
	Job 1	Job 2	Job 3	Job 4	Job 5
Sce. 1	5	5	5	9	9
Sce. 2	3	10	6	10	10
Sce. 3	10	2	4	8	4
Sce. 4	8	9	5	4	6
Sce. 5	8	1	8	5	9
Sce. 6	6	1	9	1	2
Sce. 7	4	3	7	4	1
Sce. 8	3	3	7	8	4
Sce. 9	2	2	4	7	6
Sce. 10	2	3	1	10	1

(a) Baseline processing times

Solution	Job sequence	Control dec.
A	4 1 2 5 3	1 0 0 1 0
B	2 4 1 5 3	0 1 0 1 0
C	4 2 1 5 3	0 1 0 1 0
D	1 4 2 5 3	1 0 0 1 0
E	4 5 1 2 3	0 0 0 1 1
F	4 1 5 2 3	0 0 0 1 1

(b) Optimal solutions

Figure 3: Optimal objective function values and solutions for illustrative example (A-F: Optimal decisions)



8.3 Computational performance Optimization problems were modeled with the AMPL mathematical programming language, and solved using CPLEX 12.6 with its default set of options and parameters. Experiments were carried out on four threads of a Lenovo® workstation with two Intel® Xeon® 2.30 GHz CE5-2630 CPUs and 64 GB memory running on Microsoft Windows Server 8.1 Pro x64 Edition. All data points in Tables 2–4 have been averaged over two instances. Reported times are the elapsed

solution times of the branch-and-bound algorithm, and the time limit is set to 7200 seconds. If optimality is not proven within the time limit, we record both the best lower bound on the optimal objective value (retrieved from CPLEX and denoted by LB) and the best available objective value (denoted by UB). Relative optimality gap values $\text{ROG} = (\text{UB} - \text{LB}) / \text{UB}$ are then reported.

8.3.1 Overview of solved MILP formulation variants In this section we derive some equivalent MILP formulations for our DRO machine scheduling problem. Our starting point is (46), with the functional forms $h(\mathbf{u}) = \sum_{t \in [T]} \sum_{l \in [L]} h_{tl} u_{tl}$ and $\xi_l^i(\mathbf{u}) = \sum_{t \in [T]} \hat{\xi}_{tl}^i u_{tl}$ substituted into it, and incorporating the symmetry considerations from Remark 7.2. Introducing the reduced index sets $I = \{(i, j, l) \in [n] \times [n] \times [L] : i < j\}$ and $J = \{(i, j \in [n] \times [n] : i < j\}$, we have:

$$\min \sum_{t \in [T]} \sum_{l \in [L]} h_{tl} u_{tl} + \eta + \frac{1}{1 - \alpha} \sum_{i \in [n]} p^i v^i + \frac{1}{1 - \alpha} \kappa \tau \quad (47a)$$

s.t. (46c)–(46e),

$$v^i \geq \sum_{l \in [L]} \sum_{k \in [L]} \sum_{t \in [T]} w_l \hat{\xi}_{tk}^j z_{tkl} - \eta, \quad \forall i \in [n] \quad (47b)$$

$$v^i \geq \sum_{l \in [L]} \sum_{k \in [L]} \sum_{t \in [T]} w_l \hat{\xi}_{tk}^j z_{tkl} - \eta - \sum_{l \in [L]} \nu_l^{ij} \tau, \quad \forall (i, j) \in J \quad (47c)$$

$$v^j \geq \sum_{l \in [L]} \sum_{k \in [L]} \sum_{t \in [T]} w_l \hat{\xi}_{tk}^i z_{tkl} - \eta - \sum_{l \in [L]} \nu_l^{ij} \tau, \quad \forall (i, j) \in J \quad (47d)$$

$$\nu_l^{ij} \leq \sum_{t \in [T]} \hat{\xi}_{tl}^i u_{tl} - \sum_{t \in [T]} \hat{\xi}_{tl}^j u_{tl} + M_l^{ij} \lambda_l^{ij}, \quad \forall (i, j, l) \in I \quad (47e)$$

$$\nu_l^{ij} \leq - \sum_{t \in [T]} \hat{\xi}_{tl}^i u_{tl} + \sum_{t \in [T]} \hat{\xi}_{tl}^j u_{tl} + M_l^{ij} (1 - \lambda_l^{ij}), \quad \forall (i, j, l) \in I \quad (47f)$$

$$(\boldsymbol{\theta}, \mathbf{u}) \in \mathcal{T} \times \mathcal{U}, \mathbf{v} \in \mathbb{R}_+^n, \tau \geq 0, \mathbf{z} \in [0, 1]^{T \times L \times L}, \boldsymbol{\lambda} \in \{0, 1\}^I, \boldsymbol{\nu} \in \mathbb{R}_+^I. \quad (47g)$$

Note that we have replaced the generic big-M constant with constraint-specific ones, which can be calculated as $M_l^{ij} = 2 \max_{t \in [T]} |\hat{\xi}_{tl}^i - \hat{\xi}_{tl}^j|$. While this change does not affect the feasible set of the problem, it has been observed in the literature that using tighter big-M bounds often leads to significantly improved computational performance. The above problem is an MINLP, with a non-convex continuous relaxation. We will obtain our first, basic MILP reformulation in three steps. First, we introduce non-negative auxiliary variables $\mathbf{q} \in \mathbb{R}_+^I$, where q_l^{ij} represents the quantity $\nu_l^{ij} \tau$. We can then replace constraints (47c)–(47f) with the following set of inequalities:

$$v^i \geq \sum_{l \in [L]} \sum_{k \in [L]} \sum_{t \in [T]} w_l \hat{\xi}_{tk}^j z_{tkl} - \eta - \sum_{l \in [L]} q_l^{ij}, \quad \forall (i, j) \in J \quad (48a)$$

$$v^j \geq \sum_{l \in [L]} \sum_{k \in [L]} \sum_{t \in [T]} w_l \hat{\xi}_{tk}^i z_{tkl} - \eta - \sum_{l \in [L]} q_l^{ij}, \quad \forall (i, j) \in J \quad (48b)$$

$$q_l^{ij} \leq \left(\sum_{t \in [T]} \hat{\xi}_{tl}^i u_{tl} - \sum_{t \in [T]} \hat{\xi}_{tl}^j u_{tl} + M_l^{ij} \lambda_l^{ij} \right) \tau, \quad \forall (i, j, l) \in I \quad (48c)$$

$$q_l^{ij} \leq \left(- \sum_{t \in [T]} \hat{\xi}_{tl}^i u_{tl} + \sum_{t \in [T]} \hat{\xi}_{tl}^j u_{tl} + M_l^{ij} (1 - \lambda_l^{ij}) \right) \tau, \quad \forall (i, j, l) \in I. \quad (48d)$$

Second, we introduce auxiliary variables $\mathbf{y} \in \mathbb{R}_+^{T \times L}$ and $\mathbf{r} \in \mathbb{R}_+^I$ to linearize terms of the form $u_{tl} \tau$ and $\lambda_l^{ij} \tau$, respectively, via two sets of McCormick inequalities:

$$y_{tl} \leq M_\tau u_{tk}, \quad \forall t \in [T], l \in [L] \quad (49a)$$

$$y_{tl} \leq \tau, \quad \forall t \in [T], l \in [L] \quad (49b)$$

$$y_{tl} \geq M_\tau u_{tk} + \tau - M_\tau, \quad \forall t \in [T], l \in [L] \quad (49c)$$

$$r_l^{ij} \leq \tau, \quad \forall (i, j, l) \in I \quad (49d)$$

$$r_l^{ij} \leq M_\tau \lambda_l^{ij}, \quad \forall (i, j, l) \in I \quad (49e)$$

$$r_l^{ij} \geq M_\tau \lambda_l^{ij} + \tau - M_\tau, \quad \forall (i, j, l) \in I, \quad (49f)$$

where M_τ is an upper bound for τ . We can calculate a suitable big-M bound as

$$M_\tau = \left(\sum_{l \in [L]} w_l \right) \left(\sum_{l \in [L]} \max_{t \in [T]} \hat{\xi}_{tl}^{\max} \right) / \min_{(i,j) \in J} \left(\sum_{l \in [L]} \min_{t \in [T]} |\hat{\xi}_{tl}^i - \hat{\xi}_{tl}^j| \right).$$

Third, we substitute the auxiliary variables \mathbf{y} and \mathbf{r} into (48c)–(48d). The resulting equivalent basic MILP reformulation of the MINLP (47) then reads

$$\min \sum_{t \in [T]} \sum_{l \in [L]} h_{tl} u_{tl} + \eta + \frac{1}{1-\alpha} \sum_{i \in [n]} p^i v^i + \frac{1}{1-\alpha} \kappa \tau$$

s.t. (46c)–(46e); (48a)–(48b); (49a)–(49f),

$$q_l^{ij} \leq \sum_{t \in [T]} \hat{\xi}_{tl}^i y_{tl} - \sum_{t \in [T]} \hat{\xi}_{tl}^j y_{tl} + M r_l^{ij}, \quad \forall (i, j, l) \in I \quad (50a)$$

$$q_l^{ij} \leq - \sum_{t \in [T]} \hat{\xi}_{tl}^i y_{tl} + \sum_{t \in [T]} \hat{\xi}_{tl}^j y_{tl} + M(\tau - r_l^{ij}), \quad \forall (i, j, l) \in I \quad (50b)$$

$$(\boldsymbol{\theta}, \mathbf{u}) \in \mathcal{T} \times \mathcal{U}, \quad \mathbf{v} \in \mathbb{R}_+^n, \quad \tau \geq 0, \quad \mathbf{z} \in [0, 1]^{T \times L \times L}, \quad \boldsymbol{\lambda} \in \{0, 1\}^I, \quad \boldsymbol{\nu} \in \mathbb{R}_+^I, \quad (50c)$$

$$\mathbf{q} \in \mathbb{R}_+^I, \quad \mathbf{y} \in \mathbb{R}_+^{T \times L}, \quad \mathbf{r} \in \mathbb{R}_+^I. \quad (50d)$$

As this formulation does not assume, or make use of, any comonotone structure in the data, we refer to it as the *non-comonotone* MILP, or **NCM** for short.

We also aim to take advantage of any comonotone structure that might be present in our dataset. Let us recall our notation from Section 6.1.4. Keeping in mind our symmetry considerations, we introduce the index sets $I_+ = \{(i, j, l) \in I \mid \hat{\xi}_{tl}^i \geq \hat{\xi}_{tl}^j \forall t \in [T]\}$, $I_- = \{(i, j, l) \in I \mid \hat{\xi}_{tl}^i \leq \hat{\xi}_{tl}^j \forall t \in [T]\}$, and $I_0 = I \setminus (I_+ \cup I_-)$. For index triplets in $(i, j, l) \in I_+ \cup I_-$ we can eliminate the variables q_l^{ji} from constraints (48a)–(48b), and instead for all $(i, j) \in J$ write

$$v^i \geq \sum_{l \in [L]} \sum_{k \in [L]} \sum_{t \in [T]} w_l \hat{\xi}_{tk}^j z_{tkl} - \eta - \sum_{l: (i,j,l) \in I_0} q_l^{ij} - \sum_{l: (i,j,l) \in I_+} \sum_{t \in [T]} (\hat{\xi}_{tl}^i - \hat{\xi}_{tl}^j) y_{tl} - \sum_{l: (i,j,l) \in I_-} \sum_{t \in [T]} (-\hat{\xi}_{tl}^i + \hat{\xi}_{tl}^j) y_{tl},$$

$$v^j \geq \sum_{l \in [L]} \sum_{k \in [L]} \sum_{t \in [T]} w_l \hat{\xi}_{tk}^i z_{tkl} - \eta - \sum_{l: (i,j,l) \in I_0} q_l^{ij} - \sum_{l: (i,j,l) \in I_+} \sum_{t \in [T]} (\hat{\xi}_{tl}^i - \hat{\xi}_{tl}^j) y_{tl} - \sum_{l: (i,j,l) \in I_-} \sum_{t \in [T]} (-\hat{\xi}_{tl}^i + \hat{\xi}_{tl}^j) y_{tl}.$$

The variables λ_l^{ji} , q_l^{ji} , r_l^{ji} , and the corresponding constraints (49d)–(49f), (50a)–(50b) then only need to be defined for indices $(i, j, l) \in I_0$. We will refer to the arising reduced formulation, which takes advantage of the comonotone structure present in the dataset, as the *partially comonotone* MILP, or **PCM** for short.

Finally, as noted at the end of Section 7.2.2, using the reformulation-linearization technique we can add the valid inequalities $\sum_{t \in [T]} y_{tl} = \tau$ and $\sum_{t \in [T]} z_{tkl} = \theta_{kl}$ for all $k, l \in [L]$ to our partially comonotone MILP, to obtain the formulation **PCM-RLT**. We also point out that under the strong comonotonicity assumption of Section 6.1.4 we have $I_0 = \emptyset$, so in this case the auxiliary variables λ_l^{ji} , q_l^{ji} , r_l^{ji} , along with the corresponding constraints, can be completely dropped from **PCM** and **PCM-RLT**. We refer to the resulting streamlined formulations as *completely comonotone* MILPs (**CCM** and **CCM-RLT** for short).

REMARK 8.1 *All of the formulations in this section were presented for the cost minimizing variant of the*

problem. We can obtain the budget-constrained variants by dropping the cost term $\sum_{t \in [T]} \sum_{l \in [L]} h_{tl} u_{tl}$ from the objective, and adding a constraint of the form $\sum_{t \in [T]} \sum_{l \in [L]} h_{tl} u_{tl} \leq B$.

8.3.2 Solution times First, in order to evaluate the performance of the completely comonotone formulation variants introduced in Section 8.3.1, we have considered comonotone datasets of various sizes (changing the number n of scenarios between 50 and 500, and the number L of jobs between 10 and 50). As Table 2 shows, the cost minimizing and the budget-constrained models are of comparable difficulty. This is not surprising, because adding a budget constraint reduces the space of feasible control decisions at the cost of introducing additional combinatorial complexity, and the impact of this tradeoff can vary significantly between individual instances. We also see that (with the exception of the smallest, 10-job scheduling problems) the reformulation-linearization technique consistently and dramatically improves computational performance. Notably, several large problem instances which could not be solved within the time limit using the basic comonotone formulation **CCM** are rendered tractable by employing **CCM-RLT** instead.

We next examine the performance impact of modeling parameters, using the overall best-performing **CCM-RLT** formulation on problems with comonotone data. Table 3 shows that, somewhat contrary to expectations (cf. Noyan and Rudolf, 2013, Section 7.2.1), problems with higher CVaR confidence levels are not always easier to solve than their counterparts with lower α values. In line with our observations in Section 8.2, lowering the budget to 30% of its base value reduces the space of feasible control decisions, and thus typically leads to lower running times.

We have seen that with completely comonotone formulations it is often possible to solve moderate-sized comonotone problems within reasonable time. However, non-comonotone instances prove to be significantly more challenging. The relevant performance data can be found in Table 4. Notably, using the **PCM** formulation, which exploits the partial comonotone structure present in the data, drastically improves computational efficiency even when the entire data set is not comonotone. As before, for instances with 15 or more jobs the performance is further improved by applying the reformulation-linearization technique, which makes **PCM-RLT** overall the most efficient of our formulations for non-comonotone instances. However, larger instances typically cannot be solved to optimality within the two-hour time limit. We therefore also report the results when we add an alternative stopping criterion of a 2% optimality gap, and note that it is consistently possible to find near-optimal solutions for instances that feature up to 30 jobs and 100 scenarios.

REMARK 8.2 To quantify how far the datasets used in the construction of Table 4 are from being comonotone, for each instance we calculated the Kendall rank correlation coefficients between baseline and compressed processing times for each job. The average of these coefficients (taken over all jobs) will be equal to 1 if and only if our comonotonicity condition is satisfied, and will be strictly less than 1 otherwise. For our randomly generated non-comonotone data sets typical values of these averages were around 0.61, with a standard deviation of approximately 0.01.

The performance results in Tables 2–4 reflect the increasing complexity of our MILP formulations for larger instances. In non-comonotone formulations the number of binary variables, continuous variables, and constraints all grow as $O(n^2L + L^2)$. Completely comonotone formulations reduce the number of binary variables to $O(L^2)$; in the partially comonotone case the extent of the reduction is typically smaller, depending on the structure of the particular data set. Therefore, consistently with the patterns generally seen in the deterministic scheduling literature, considering a larger number L of jobs leads to a rapid increase in difficulty in all cases, due to the higher combinatorial complexity. As often observed

in stochastic optimization, increasing the number n of scenarios also presents a major computational bottleneck; the effect is especially pronounced for non-comonotone data sets. Since our problems combine the challenges of two already computationally demanding fields, they can only be solved efficiently when one takes advantage of additional structure in the data.

Table 2: Computational performance – comonotone data.

		Time [ROG (%)]			
		Cost Minimizing		Budget-Constrained	
L	n	CCM	CCM-RLT	CCM	CCM-RLT
10	50	8.4	7.4	8.9	7.4
	100	52.9	46.6	61.6	23.5
	150	271.3	137.0	226.2	175.5
	200	567.2	539.9	615.3	359.4
	250	1210.0	775.6	1626.2	451.8
	300	1002.6	1023.4	1251.2	2077.5
	400	1991.6	3155.5	1963.7	1742.2
	500	5160.7	[†] 6046.3[0.13]	4496.1	[†] 5783.11[0.45]
15	50	113.1	43.0	172.2	34.5
	100	603.3	190.0	627.8	155.3
	150	2513.8	555.4	2460.8	376.1
	200	5240.7	1221.9	5891.9	510.6
	250	[†] 6737.8[3.29]	2659.1	^{††} [4.47]	2135.5
	300	^{††} [9.54]	3581.7	^{††} [12.92]	2524.2
	400	–	^{††} [1.34]	–	5013.1
	500	–	^{††} [1.89]	–	^{††} [10.91]
20	50	2318.6	62.4	3357.8	71.6
	100	^{††} [7.89]	462.6	^{††} [18.87]	523.0
	150	^{††} [16.87]	1326.7	^{††} [29.79]	1044.5
	200	–	2768.2	–	2697.5
	250	–	5411.8	–	5035.7
	300	–	^{††} [1]	–	6057.3
	400	–	^{††} [4.71]	–	^{††} [3.23]
	25	50	^{††} [14.45]	157.0	^{††} [24.83]
100		^{††} [32.41]	1353.8	^{††} [46.12]	1514.4
150		–	5274.7	–	1611.7
200		–	^{††} [0.42]	–	5390.6
250		–	^{††} [1.22]	–	[†] 6985[0.1]
300		–	^{††} [1.58]	–	^{††} [1.19]
30	50	^{††} [28.39]	284.4	^{††} [43.13]	383.0
	100	^{††} [47.15]	1942.9	^{††} [55.13]	3551.4
	150	–	[†] 5590.2[0.06]	–	^{††} [0.46]
50	50	^{††} [60.97]	5614.17	^{††} [74.09]	2402.6
	100	^{††} [69.19]	^{††} [1.21]	^{††} [85.5]	^{††} [0.23]

[†]: Each dagger sign indicates one instance hitting the time limit with an integer feasible solution.

9. Further avenues of research One of the main distinguishing features of our approach is that the nominal distribution at the center of the ambiguity set is decision-dependent. In this regard, it would be essential to investigate possible ways to describe this dependence, and in particular to develop meaningful and tractable characterizations of decision-dependent nominal parameter realizations and/or scenario probabilities (akin to the distribution shaping equations) for practical applications.

While scenario bundling is a very effective method of reducing problem sizes, most EMDs are not

Table 3: Impact of modeling parameters on performance of **CCM-RLT**

		Time [ROG (%)]				
		Cost Minimizing (α)		Budget-Constrained (BR, α)		
L	n	0.7	0.9	(1, 0.7)	(1, 0.9)	(0.3, 0.9)
15	50	37.8	43.0	31.7	34.5	17.3
	100	179.1	190.0	100.4	155.3	79.9
	150	526.0	555.4	367.8	376.1	274.4
	200	902.2	1221.9	814.6	510.6	540.5
	250	3155.5	2659.1	1909.5	2135.5	1132.5
	300	2921.2	3581.7	2640.6	2524.2	2815.5
	400	$\dagger\dagger[0.95]$	$\dagger\dagger[1.35]$	$\dagger 5469.5[1.78]$	5013.1	3457.6
20	50	80.5	62.4	62.1	71.6	49.4
	100	367.7	462.6	337.3	523.0	260.6
	150	972.1	1326.7	833.2	1044.5	592.2
	200	3146.0	2768.2	1742.4	2697.5	1226.6
	250	5459.7	5411.8	3172.8	5035.7	2408.3
	300	4820.1	$\dagger\dagger[1]$	$\dagger 5973.3[1.28]$	6057.3	$\dagger 5080.345[0.53]$
	400					

\dagger : Each dagger sign indicates one instance hitting the time limit with an integer feasible solution.

Table 4: Computational performance on non-comonotone instances

		Time [ROG (%)]			
L	n	NCM	PCM	PCM-RLT	PCM-RLT CPLEX mipgap=2%
10	50	3036.7	151.6	149.5	138.8[0.84]
	100	$\dagger\dagger[13.87]$	1361.3	4050.3	1618.8[0.93]
	150	$\dagger\dagger[38.88]$	$\dagger\dagger[0.78]$	$\dagger\dagger[3.14]$	–
	200	–	$\dagger\dagger[23.15]$	$\dagger\dagger[10.41]$	–
	250	–	$\dagger\dagger[10.10]$	$\dagger\dagger[4.03]$	–
15	50	$\dagger\dagger[0.59]$	1375.8	398.5	40.9[1.23]
	100	$\dagger\dagger[25.25]$	$\dagger\dagger[4.14]$	$\dagger\dagger[0.42]$	1403.4[1.9]
	150	$\dagger\dagger[76.41]$	$\dagger\dagger[22.66]$	$\dagger\dagger[2.27]$	$\dagger 4538.4[2.48]$
	200	–	–	$\dagger\dagger[2.60]$	–
20	50	$\dagger\dagger[18.9]$	4570.5	1405.3	105.5[0.68]
	100	$\dagger\dagger[65.81]$	$\dagger\dagger[23.54]$	$\dagger\dagger[1.28]$	2598.2[1.71]
	150	–	$\dagger\dagger[43.5]$	$\dagger\dagger[2.21]$	$\dagger 5140.6[2.28]$
25	50	$\dagger\dagger[42.27]$	$\dagger\dagger[19.04]$	1670.8	161.7[1.95]
	100	$\dagger\dagger[79.21]$	$\dagger\dagger[43.98]$	$\dagger\dagger[0.52]$	1156.6[1.11]
	150	–	–	$\dagger\dagger[2.69]$	–
30	50	$\dagger\dagger[57.81]$	$\dagger\dagger[36.44]$	4501.7	234[0.54]
	100	–	–	$\dagger\dagger[1.37]$	3220.5[1.78]

\dagger : Each dagger sign indicates one instance hitting the time limit with an integer feasible solution.

compatible with this approach. As we have seen, the total variation metric is a notable exception. In Remark A.2 we briefly touched on another, potentially more informative, class of outcome-based scenario distances, which give rise to EMDs that can be used in conjunction with bundling. However, as many of our developments strongly depend on the structural properties of the variation metric, incorporating this new class of distances into our formulations would require additional work.

For problems of practical interest where bundling methods are not applicable, one might instead consider sampling methods to reduce the number of scenarios. The use of sampling, however, comes with two important caveats. First, sampling typically sacrifices exact solutions in exchange for tractability. Second, even in cases where a particular sampling method (such as importance sampling) is known to work well for the underlying problem, this does not automatically translate to a performance guarantee for the DRO variant. For example, when using an ambiguity set based on the discrete metric, the worst-case distribution will be highly sensitive to the worst scenario included in a sample. However, studies such as Bardou et al. (2009) indicate that sampling approaches could be better suited for CVaR-based and other risk-averse formulations which do not explicitly feature the ambiguity set. Along similar lines, the use of scenario reduction techniques could also be explored.

In line with the majority of the DRO literature, we have adopted a risk-averse pessimistic viewpoint, focusing on the worst outcome in the ambiguity set. In contrast, optimistic robust optimization has recently been suggested to be meaningful and relevant for certain application areas, including machine learning (see, e.g., Norton et al., 2017). Taking an optimistic view in our problems would lead us to replace maximization over the ambiguity set with minimization, typically making the resulting problems significantly more straightforward. Therefore, this seems to be a worthwhile avenue to explore whenever an optimistic view is warranted by a particular application.

References

- Acerbi, C. (2002). Spectral measures of risk: a coherent representation of subjective risk aversion. *Journal of Banking and Finance*, 26(7):1505–1518.
- Artzner, P., Delbaen, F., Eber, J., and Heath, D. (1999). Coherent measures of risk. *Mathematical Finance*, 9(3):203–228.
- Atakan, S., Bülbül, K., and Noyan, N. (2017). Minimizing value-at-risk in single-machine scheduling. *Annals of Operations Research*, 248(1):25–73.
- Bardou, O., Frikha, N., and Pages, G. (2009). Computing var and cvar using stochastic approximation and adaptive unconstrained importance sampling. *Monte Carlo Methods and Applications*, 15(3):173–210.
- Bayraksan, G. and Love, D. (2015). *TutORials in Operations Research*, chapter Data-Driven Stochastic Programming using Phi-Divergences, pages 1–15. INFORMS.
- Bertsimas, D. and Vayanos, P. (2014). Data-driven learning in dynamic pricing using adaptive optimization. Available at Optimization Online: http://www.optimization-online.org/DB_FILE/2014/10/4595.pdf.
- Blanchet, J., Kang, Y., Zhang, F., He, F., and Hu, Z. (2017). Doubly robust data-driven distributionally robust optimization. Technical report. Technical Report 1705.07168, ArXiv.
- Calafiore, G. C. (2007). Ambiguous risk measures and optimal robust portfolios. *SIAM Journal on Optimization*, 18(3):853–877.
- Chang, Z., Song, S., Zhang, Y., Ding, J.-Y., Zhang, R., and Chiong, R. (2017). Distributionally robust single machine scheduling with risk aversion. *European Journal of Operational Research*, 256(1):261 – 274.

- Cormican, K. J., Morton, D. P., and Wood, R. K. (1998). Stochastic network interdiction. *Operations Research*, 46(2):184–197.
- Delage, E. and Ye, Y. (2010). Distributionally robust optimization under moment uncertainty with application to data-driven problems. *Operations Research*, 58(3):595–612.
- Dentcheva, D. and Ruszczyński, A. (2018). Risk forms: Representation, disintegration, and application to partially observable two-stage systems. Technical report. Technical Report 807.02300, ArXiv.
- Dupacova, J. (2006). Optimization under exogenous and endogenous uncertainty. University of West Bohemia in Pilsen. Unpublished, <http://www.karlin.mff.cuni.cz/~dupacova/papers/MME06sty.pdf>.
- Erdoğ̃an, E. and Iyengar, G. (2006). Ambiguous chance constrained problems and robust optimization. *Mathematical Programming*, 107(1-2):37–61.
- Esfahani, P. and Kuhn, D. (2018). Data-driven distributionally robust optimization using the wasserstein metric: performance guarantees and tractable reformulations. *Mathematical Programming*. Online first <https://link.springer.com/article/10.1007/s10107-017-1172-1>.
- Flach, B. and Poggi, M. (2010). On a class of stochastic programs with endogenous uncertainty: theory, algorithm and application. Technical report. Monografias em Ciencia da Computacao No 05/10, PUC, Rio.
- Föllmer, H. and Schied, A. (2002). Convex measures of risk and trading constraints. *Finance and Stochastics*, 6(4):429–447.
- Gao, R., Chen, X., and Kleywegt, A. J. (2017). Wasserstein distributional robustness and regularization in statistical learning. *CoRR*, abs/1712.06050.
- Gao, R. and Kleywegt, A. (2016). Distributionally robust stochastic optimization with Wasserstein distance. Technical report. Technical Report 1604.02199, ArXiv.
- Gao, R. and Kleywegt, A. (2017). Distributionally robust stochastic optimization with dependence structure. Available at Optimization Online: http://www.optimization-online.org/DB_HTML/2017/01/5817.html.
- Goel, V. and Grossmann, I. E. (2004). A stochastic programming approach to planning of offshore gas field developments under uncertainty in reserves. *Computers & Chemical Engineering*, 28(8):1409 – 1429.
- Goel, V. and Grossmann, I. E. (2006). A class of stochastic programs with decision dependent uncertainty. *Mathematical Programming*, 108(2):355–394.
- Goh, J. and Sim, M. (2010). Distributionally robust optimization and its tractable approximations. *Operations Research*, 58(4-part-1):902–917.
- Gupte, A., Ahmed, S., Dey, S. S., and Cheon, M. S. (2017). Relaxations and discretizations for the pooling problem. *Journal of Global Optimization*, 67(3):631–669.
- Hansen, P., Jaumard, B., and Lu, S.-H. (1992). Global optimization of univariate lipschitz functions: Ii. new algorithms and computational comparison. *Mathematical Programming*, 55(1):273–292.
- Haus, U.-U., Michini, C., and Laumanns, M. (2017). Scenario aggregation using binary decision diagrams for stochastic programs with endogenous uncertainty. Unpublished, <https://arxiv.org/pdf/1701.04055.pdf>.
- Hellemo, L., Barton, P., and Tomasgard, A. (2018). Decision-dependent probabilities in stochastic programs with recourse. *Computational Management Science*, 15(3):369–395.

- Hradovich, M., Kasperski, A., and Zieliński, P. (2018). Risk averse scheduling with scenarios. In Kliewer, N., Ehmke, J. F., and Borndörfer, R., editors, *Operations Research Proceedings 2017*, pages 435–441, Cham. Springer International Publishing.
- Hu, Z. and Hong, L. J. (2012). Kullback-Leibler divergence constrained distributionally robust optimization. Available at Optimization Online: http://www.optimization-online.org/DB_FILE/2012/11/3677.pdf.
- Ji, R. and Lejeune, M. (2017). Data-driven optimization of ambiguous reward-risk ratio measures. Working Paper, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2707122.
- Ji, R. and Lejeune, M. (2018). Data-driven distributionally robust chance-constrained optimization with wasserstein metric. Working Paper, http://www.optimization-online.org/DB_HTML/2018/07/6697.html.
- Jiang, R. and Guan, Y. (2015). Data-driven chance constrained stochastic program. *Mathematical Programming*, 158:291–327.
- Jiang, R. and Guan, Y. (2018). Risk-averse two-stage stochastic program with distributional ambiguity. *Operations Research*. Online first <https://pubsonline.informs.org/doi/10.1287/opre.2018.1729>.
- Jonsbråten, T., Wets, R.-B., and Woodruff, D. (1998). A class of stochastic programs with decision dependent random elements. *Annals of Operations Research*, 82(0):83–106.
- Kantorovich, L. V. and Rubinshtein, G. S. (1958). On a space of totally additive functions. *Vestnik Leningradskogo Universiteta*, 13, pp. 52–59.
- Khaligh, F. H. and MirHassani, S. (2016). A mathematical model for vehicle routing problem under endogenous uncertainty. *International Journal of Production Research*, 54(2):579–590.
- Kusuoka, S. (2001). On law invariant coherent risk measures. *Advances in Mathematical Economics*, 3:83–95.
- Kyparisis, J. and Fiacco, A. V. (1987). Generalized convexity and concavity of the optimal value function in nonlinear programming. *Mathematical Programming*, 39(3):285–304.
- Lam, H. (2016). Recovering best statistical guarantees via the empirical divergence-based distributionally. Technical report. Technical Report 1605.09349, ArXiv.
- Lappas, N. H. and Gounaris, C. E. (2018). Robust optimization for decision-making under endogenous uncertainty. *Computers & Chemical Engineering*. Online first <https://www.sciencedirect.com/science/article/pii/S0098135418300152>.
- Laumanns, M., Prestwich, S., and Kawas, B. (2014). Distribution shaping and scenario bundling for stochastic programs with endogenous uncertainty. Unpublished, <https://edoc.hu-berlin.de/bitstream/handle/18452/9095/5.pdf>.
- Lejeune, M. A. and Shen, S. (2016). Multi-objective probabilistically constrained programs with variable risk: Models for multi-portfolio financial optimization. *European Journal of Operational Research*, 252(2):522 – 539.
- Lindvall, T. (1992). *Lectures on the Coupling Method*. Wiley Series in Probability and Statistics - Applied Probability and Statistics Section. Wiley.
- Liu, X., Kucukyavuz, S., and Noyan, N. (2017). Robust multicriteria risk-averse stochastic programming models. *Annals of Operations Research*, 259(1):259–294.
- Luo, F. and Mehrotra, S. (2017). Decomposition algorithms for distributionally robust optimization using wasserstein metric. http://www.optimization-online.org/DB_HTML/2017/04/5946.html.

- Luo, F. and Mehrotra, S. (2018). Distributionally robust optimization with decision dependent ambiguity sets. Technical report. Technical Report 1806.09215, ArXiv.
- McCormick, G. (1976). Computability of global solutions to factorable nonconvex programs: Part I – convex underestimating problems. *Mathematical Programming*, 10(1):147–175.
- Müller, A. and Stoyan, D. (2002). *Comparison Methods for Stochastic Models and Risks*. John Wiley & Sons, Chichester.
- Niu, S., Song, S., Ding, J.-Y., Zhang, Y., and Chiong, R. (2019). Distributionally robust single machine scheduling with the total tardiness criterion. *Computers & Operations Research*, 101:13 – 28.
- Nohadani, O. and Sharma, K. (2018). Optimization under decision-dependent uncertainty. *SIAM Journal on Optimization*, 2:1773–1795.
- Norton, M., Takeda, A., and Mafusalov, A. (2017). Optimistic Robust Optimization With Applications To Machine Learning. *ArXiv e-prints*.
- Noyan, N. and Rudolf, G. (2013). Optimization with multivariate conditional value-at-risk-constraints. *Operations Research*, 61(4):990–1013.
- Noyan, N. and Rudolf, G. (2015). Kusuoka representations of coherent risk measures in general probability spaces. *Annals OR*, 229(1):591–605.
- Noyan, N. and Rudolf, G. (2018). Optimization with stochastic preferences based on a general class of scalarization functions. *Operations Research*, 66(2):463–486.
- Noyan, N., Rudolf, G., and Lejeune, M. (2018). Distributionally robust optimization with decision-dependent ambiguity set. Available at Optimization Online: http://www.optimization-online.org/DB_FILE/2018/09/6821.pdf.
- Peeta, S., Salman, F., Gunneç, D., and Viswanath, K. (2010). Pre-disaster investment decisions for strengthening a highway network. *Computers and Operations Research*, 37:1708–1719.
- Pflug, G. and Wozabal, D. (2007). Ambiguity in portfolio selection. *Quantitative Finance*, 7(4):435–442.
- Pflug, G. C. and Pichler, A. (2011). *Approximations for Probability Distributions and Stochastic Optimization Problems*, pages 343–387. Springer New York, New York, NY.
- Pflug, G. C., Pichler, A., and Wozabal, D. (2012). The 1/N investment strategy is optimal under high model ambiguity. *Journal of Banking & Finance*, 36(2):410–417.
- Pflug, G. C. and Römisich, W. (2007). *Modelling, managing and measuring risk*. World Scientific publishing, Singapore.
- Pinedo, M. (2008). *Scheduling: Theory, Algorithms, and Systems*. Springer, 3rd edition.
- Postek, K., Den Hertog, D., and Melenberg, B. (2016). Computationally tractable counterparts of distributionally robust constraints on risk measures. *SIAM Review*, 58(4):603–650.
- Rahimian, H., Bayraksan, G., and Homem-de Mello, T. (2018). Identifying effective scenarios in distributionally robust stochastic programs with total variation distance. *Mathematical Programming*.
- Rockafellar, R. (2007). Coherent approaches to risk in optimization under uncertainty. In *Tutorials in Operations Research*, pages 38–61. INFORMS.
- Rockafellar, R. and Uryasev, S. (2000). Optimization of conditional value-at-risk. *The Journal of Risk*, 2(3):21–41.
- Royset, J. O. and Wets, R. J.-B. (2017). Variational theory for optimization under stochastic ambiguity. *SIAM Journal on Optimization*, 27(2):1118–1149.

- Rubner, Y., Tomasi, C., and Guibas, L. J. (1998). A metric for distributions with applications to image databases. In *Sixth International Conference on Computer Vision (IEEE Cat. No.98CH36271)*, pages 59–66.
- Schichl, H. and Sellmann, M. (2015). Predisaster preparation of transportation networks. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence, AAAI'15*, pages 709–715. AAAI Press.
- Shabtay, D. and Steiner, G. (2007). A survey of scheduling with controllable processing times. *Discrete Applied Mathematics*, 155(13):1643 – 1666.
- Shapiro, A., Dentcheva, D., and Ruszczyński, A. (2009). *Lectures on stochastic programming: modeling and theory*. The society for industrial and applied mathematics and the mathematical programming society, Philadelphia, USA.
- Sherali, H. D. and Adams, W. P. (1994). A hierarchy of relaxations and convex hull representations for mixed-integer zero-one programming problems. *Discrete Appl. Math.*, 52(1):83–106.
- Sherali, H. D., Adams, W. P., and Driscoll, P. J. (1998). Exploiting special structures in constructing a hierarchy of relaxations for 0-1 mixed integer problems. *Operations Research*, 46(3):396–405.
- Van Parys, B., Estafahani, M., and Kuhn, D. (2017). From data to decisions: Distributionally robust optimization is optimal. Technical report. Technical Report 1704.04118, ArXiv.
- Wang, Z., Glynn, P. W., and Ye, Y. (2016). Likelihood robust optimization for data-driven problems. *Computational Management Science*, 13(2):241–261.
- Wiesemann, W., Kuhn, D., and Sim, M. (2014). Distributionally robust convex optimization. *Operations Research*, 62(6):1358–1376.
- Wozabal, D. (2014). Robustifying convex risk measures for linear portfolios: A nonparametric approach. *Operations Research*, 62(6):1302–1315.
- Zhang, J., Xu, H., and Zhang, L. (2016). Quantitative stability analysis for distributionally robust optimization with moment constraints. *SIAM Journal on Optimization*, 26(3):1855–1882.
- Zhang, Y., Shen, Z.-J. M., and Song, S. (2018). Exact algorithms for distributionally β -robust machine scheduling with uncertain processing times. *INFORMS Journal on Computing*, 30(4):662–676.
- Zhao, C. and Guan, Y. (2015). Data-driven risk-averse two-stage stochastic program with ζ -structure probability metrics. Available at Optimization Online: http://www.optimization-online.org/DB_HTML/2015/07/5014.html.
- Zhao, C. and Guan, Y. (2018). Data-driven risk-averse stochastic optimization with Wasserstein metric. *Operations Research Letters*, 46(2):262 – 267.
- Zymler, S., Kuhn, D., and Rustem, B. (2013). Worst-case value at risk of nonlinear portfolios. *Management Science*, 59(1):172–188.

Appendix A. Stochastic network models This appendix contains supplementary materials and discussions for Section 7.1. In A.1 and A.2 we detail how the components of the general model (35), and its DRO variant (37), can be specified for our motivating applications. We then discuss disjunction-based MILP formulations in A.3, and explore a CVaR-based alternative perspective in A.4. The remaining sections of the Appendix are dedicated to describing scenario bundling (A.5), its application to our DRO setting (A.6), and the limitations of the bundling approach (A.7).

A.1 Model specifics for stochastic pre-disaster investment planning problem (SPIPP)

Let us assume that a highway network is modeled by the undirected graph (V, E) on vertex set V with edge set $E = \{e_1, \dots, e_L\}$, and edge lengths represent traversal costs. The interpretation of the reinforcement decision variables \mathbf{x} and the survival probabilities $\sigma_\ell^0, \sigma_\ell^1$ is self-explanatory. If the cost of reinforcing the highway link $\ell \in [L]$ is c_ℓ , and the available budget is C , then the set of feasible decisions is given by $\mathcal{X} = \{\mathbf{x} \in \{0, 1\}^L : \sum_{\ell \in [L]} c_\ell x_\ell \leq C\}$. The post-disaster state of the network in scenario $i \in [n]$ can be modeled by the graph (V, E^i) , where $E^i = \{e_\ell : \xi_\ell^i = 1\}$ is the set of edges that survive in scenario i . In the simplest case, when we consider only single O-D pair, the outcome G^i in scenario $i \in [n]$ is given by the length of the shortest O-D path in the surviving graph (V, E^i) . However, in practice this value is modified: if the shortest path length exceeds a certain acceptability threshold, or if no O-D path exists in (V, E^i) , then the value of G^i is set to a fixed penalty parameter. This parameter typically represents the traversal costs associated with an alternate mode of transportation (such as helicopter), which relief organizations can resort to when no acceptable routes survive in the network. When multiple O-D pairs are given, along with corresponding importance weights, the outcome G^i naturally becomes the weighted sum of the (modified) shortest O-D path lengths in (V, E^i) .

A.2 Model specifics for stochastic network interdiction problem (SNIP)

Specifying the model components in accordance with the conventions we established is not always as straightforward as for SPIPP. Let us now assume that a transportation network for SNIP is modeled by a digraph (V, A) with arc set $A = \{a_1, \dots, a_L\}$, specified source and sink nodes, and given arc capacities. We recall that under our assumptions the ordering of survival probabilities is fixed for every link: we have $\sigma_\ell^0 \leq \sigma_\ell^1$ for all $\ell \in [L]$. Since adherence to this convention will be necessary when applying the so-called *distribution shaping* methods discussed in the next section, we need to denote the baseline survival probability of arc a_ℓ by σ_ℓ^1 . If an attempt is made to block this arc, its survival probability is reduced to σ_ℓ^0 . Accordingly, our decision variables will have a somewhat counterintuitive interpretation, wherein $x_\ell = 1$ holds if and only if the defender does **not** attempt to block arc a_ℓ . If the cost of attempting to block arc a_ℓ is c_ℓ , and the defender's budget is C , then the set of feasible decisions is given by $\mathcal{X} = \{\mathbf{x} \in \{0, 1\}^L : \sum_{\ell \in [L]} c_\ell (1 - x_\ell) \leq C\}$. Analogously to the case of SPIPP, the surviving digraph in scenario $i \in [n]$ is (V, A^i) , with $A^i = \{a_\ell : \xi_\ell^i = 1\}$. The capacities of surviving arcs remain unchanged, and the outcome G^i in scenario $i \in [n]$ is given by the value of the maximum source-sink flow on (V, A^i) .

A.3 A disjunctive approach

Our goal is to formulate the DRO problem (38) as a single MILP. To this end, let us introduce for all $j \in [n]$ a binary variable β^j that takes value 1 if and only if the nominal κ -level Value-at-Risk of G equals the realization G^j . Then, recalling the condition $\tau = G^+ - \text{VaR}_\kappa(G)$, at an optimum solution of (38) we have $v^i = \max\{G^i, G^+ - \tau\} = \max\{G^i, \sum_{j \in [n]} G^j \beta^j\} = \sum_{j \in [n]} \max\{G^i, G^j\} \beta^j$. We can now use McCormick envelopes to linearize the quadratic terms $\pi_L^i v^i$ in the objective (38a) by introducing the auxiliary variables $z^{ij} = \pi_L^i \beta^j$, which leads to the following MILP formulation of (38):

$$\min \sum_{i \in [n]} \sum_{j \in [n]} \max\{G^i, G^j\} z^{ij} + \kappa \left(G^+ - \sum_{j \in [n]} G^j \beta^j \right) \quad (52a)$$

$$\text{s.t.} \quad (36b)-(36f), \quad (52b)$$

$$z^{ij} \leq \pi_L^i, \quad \forall i, j \in [n] \quad (52c)$$

$$z^{ij} \leq \beta^j, \quad \forall i, j \in [n] \quad (52d)$$

$$z^{ij} \geq \pi_L^i + \beta^j - 1, \quad \forall i, j \in [n] \quad (52e)$$

$$\sum_{j \in [n]} \beta^j = 1, \quad (52f)$$

$$\boldsymbol{\beta} \in \{0, 1\}^n, \quad \mathbf{z} \in [0, 1]^{n \times n}. \quad (52g)$$

Keeping in mind that the auxiliary variables \mathbf{z} are non-negative, the constraints (52c)–(52e) will ensure that the defining equations $z^{ij} = \pi_L^i \beta^j$ hold for all $i, j \in [n]$.

We can significantly strengthen the above MILP. Introducing the notation $\mathbf{1} = (1, \dots, 1) \in \mathbb{R}^n$, in any feasible solution of (52) the triple $(\mathbf{z}, \boldsymbol{\beta}, \boldsymbol{\pi}_L)$ belongs to the mixed-integer set

$$T = \{(\mathbf{z}, \boldsymbol{\beta}, \boldsymbol{\pi}_L) \in [0, 1]^{n \times n} \times \{0, 1\}^n \times [0, 1]^n : \mathbf{z} = \boldsymbol{\pi}_L \boldsymbol{\beta}^\top, \quad \boldsymbol{\beta}^\top \mathbf{1} = 1, \quad \mathbf{1}^\top \boldsymbol{\pi}_L = 1\}.$$

Sets of similar structure appear in so-called *pooling problems*, as discussed, for example, by Gupte et al. (2017), who use the reformulation-linearization technique to describe the convex hull (see also Sherali et al., 1998). Adapting their result to our setting, we obtain

$$\text{conv}(T) = \{(\mathbf{z}, \boldsymbol{\beta}, \boldsymbol{\pi}_L) \in [0, 1]^{n \times n} \times [0, 1]^n \times [0, 1]^n : \mathbf{1}^\top \mathbf{z} = \boldsymbol{\beta}^\top, \quad \mathbf{z} \mathbf{1} = \boldsymbol{\pi}_L, \quad \mathbf{1}^\top \boldsymbol{\pi}_L = 1\}.$$

Incorporating the arising valid inequalities into an MILP can dramatically improve computational performance, as observed for example by Liu et al. (2017) in a different context. We then arrive at the following formulation:

$$\min \sum_{i \in [n]} \sum_{j \in [n]} \max\{G^i, G^j\} z^{ij} + \kappa \left(G^+ - \sum_{j \in [n]} G^j \beta^j \right) \quad (53a)$$

$$\text{s.t.} \quad (36b)–(36f), \quad (53b)$$

$$\sum_{i \in [n]} z^{ij} = \beta^j, \quad \forall j \in [n] \quad (53c)$$

$$\sum_{j \in [n]} z^{ij} = \pi_L^i, \quad \forall i \in [n] \quad (53d)$$

$$\sum_{j \in [n]} \beta^j = 1, \quad (53e)$$

$$\boldsymbol{\beta} \in \{0, 1\}^n, \quad \mathbf{z} \in [0, 1]^{n \times n}. \quad (53f)$$

A.4 Non-DRO risk-averse optimization perspective In this appendix we provide an alternative view of the developments in Section 7.1.5 and Appendix A.3, under the same assumptions, and keeping our previous notation. Let us recall from Example 5.2 that, when we use a total variation-based ball of radius $\kappa \in [0, 1]$, the robustified expectation is a convex combination of the worst-case outcome (with weight κ) and the κ -level CVaR of the outcome. Accordingly, under our assumptions (DRO-RND) is equivalent to the following problem:

$$\min_{\mathbf{x} \in \mathcal{X}} \quad \kappa G^+ + (1 - \kappa) \text{CVaR}_\kappa([\mathbb{P}(\mathbf{x}), G]). \quad (54)$$

We remark that (54) is equivalent to a non-DRO CVaR minimization problem, because the worst-case outcome G^+ is not decision-dependent. In settings without decision-dependent uncertainty it is possible to obtain a linear formulation for CVaR minimization problems by using representation (5). A direct implementation of this approach in the presence of decision-dependent probabilities gives rise to a highly non-linear model. As we saw in Section 7.1.2, distribution shaping can be used to get rid of the nonlinearity in probability expressions of the form (34), and reformulate (54) as follows:

$$\min \quad \kappa G^+ + (1 - \kappa) \sum_{j \in [n]} \left(G^j + \frac{1}{1 - \kappa} \sum_{i \in [n]} \pi_L^i [G^i - G^j]_+ \right) \beta^j \quad (55a)$$

$$\text{s.t.} \quad (36b)–(36g), \quad (55b)$$

$$\sum_{j \in [n]} \beta^j = 1, \quad (55c)$$

$$\beta \in \{0, 1\}^n. \quad (55d)$$

It is well-known that in settings without decision-dependent uncertainty CVaR minimization can be expressed in a linear fashion. In contrast, the formulation (55) still has a non-convex quadratic objective. Analogously to the case of problem (38), one way to obtain an optimal solution of the problem (55) is to separately solve the n MILP formulations that arise when we fix $\beta^j = 1$ for $j \in [n]$. Alternatively, we can again use McCormick envelopes to linearize the objective function, introducing the auxiliary variables $z^{ij} = \pi_L^i \beta^j$. We omit the details here, as the arising formulation will be essentially equivalent to (52). To see that this equivalence holds, it is sufficient to verify that the objective functions (52a) and (55a) are equal, which in turn follows by straightforward calculation from the simple observation that we have $[G^i - G^j]_+ = \max\{G^i, G^j\} - G^j$ for $i, j \in [n]$.

It is also possible to obtain an MILP formulation of (54) without linearizing the quadratic terms $\pi_L^i \beta^j$. Let us introduce the auxiliary variable μ to represent $\text{CVaR}_\kappa([\mathbb{P}(\mathbf{x}), G])$. Then, using a disjunction-based representation of the finite minimum in the CVaR representation (5), we arrive at

$$\min \quad \kappa G^+ + (1 - \kappa)\mu \quad (56a)$$

$$\text{s.t.} \quad (36b)-(36g), \quad (56b)$$

$$\mu \leq G^j + \frac{1}{1 - \kappa} \sum_{i \in [n]} \pi_L^i [G^i - G^j]_+, \quad \forall j \in [n] \quad (56c)$$

$$\mu \geq G^j + \frac{1}{1 - \kappa} \sum_{i \in [n]} \pi_L^i [G^i - G^j]_+ - (1 - \beta^j)M, \quad \forall j \in [n] \quad (56d)$$

$$(55c) - (55d). \quad (56e)$$

Here, if $M \in \mathbb{R}_+$ is a suitably large constant, the constraints (56c)–(56e) ensure the validity of the defining equation

$$\mu = \text{CVaR}_\kappa([\mathbb{P}(\mathbf{x}), \xi]) = \min_{j \in [n]} G^j + \frac{1}{1 - \kappa} \sum_{i \in [n]} \pi_L^i [G^i - G^j]_+. \quad (57)$$

We could also have arrived at an equivalent form of the MILP (56) by applying a similar disjunctive approach, but with the formulation (38) as our starting point. As discussed in Section 7.1.5, the optimum in (38) can be expressed as a finite minimum. More precisely, the optimum is attained when we have $\tau = G^+ - G^j$ for some $j \in [n]$, with corresponding values $v^i = \max\{G^i, G^j\}$. Recalling that we have $\max\{G^i, G^j\} = [G^i - G^j]_+ + G^j$ for $i, j \in [n]$, we can now express the optimum value of (38) as the minimum of the following expression, taken over all (\mathbf{x}, π) satisfying (36b)–(36g):

$$\min_{j \in [n]} \sum_{i \in [n]} \pi_L^i \max\{G^i, G^j\} + \kappa(G^+ - G^j) \quad (58a)$$

$$= \kappa G^+ + \min_{j \in [n]} \sum_{i \in [n]} \pi_L^i \max\{G^i, G^j\} - \kappa G^j \quad (58b)$$

$$= \kappa G^+ + \min_{j \in [n]} \sum_{i \in [n]} \pi_L^i ([G^i - G^j]_+ + G^j) - \kappa G^j \quad (58c)$$

$$= \kappa G^+ + \min_{j \in [n]} \sum_{i \in [n]} \pi_L^i [G^i - G^j]_+ + \sum_{i \in [n]} \pi_L^i G^j - \kappa G^j \quad (58d)$$

$$= \kappa G^+ + \min_{j \in [n]} \sum_{i \in [n]} \pi_L^i [G^i - G^j]_+ + (1 - \kappa)G^j \quad (58e)$$

$$= \kappa G^+ + (1 - \kappa) \min_{j \in [n]} G^j + \frac{1}{1 - \kappa} \sum_{i \in [n]} \pi_L^i [G^i - G^j]_+ \quad (58f)$$

$$= \kappa G^+ + (1 - \kappa) \mu, \quad (58g)$$

where μ is given by the expression in (57).

A.5 Scenario bundling We introduce the symbol $*$ to indicate that status of a link is unknown, and say that a ternary vector $\mathbf{s} \in \{0, 1, *\}^L$ represents the scenario set $B_{\mathbf{s}} = \{i \in [n] : \xi_{\ell}^i = s_{\ell} \vee s_{\ell} = * \forall \ell \in [L]\}$. If we have $G^i = G^j$ for all $i, j \in B_{\mathbf{s}}$, that is, if the vector \mathbf{s} represents a set of scenarios that all result in the same outcome, then we call \mathbf{s} a *scenario bundle*, and denote this common outcome by $G^{\mathbf{s}}$. Given a decision $\mathbf{x} \in \mathcal{X}$ we define the probability of bundle \mathbf{s} as $\bar{p}^{\mathbf{s}}(\mathbf{x}) = \sum_{i \in B_{\mathbf{s}}} p^i(\mathbf{x})$. If the elements of a family S of scenario bundles represent a partition $\bigcup_{\mathbf{s} \in S} B_{\mathbf{s}} = [n]$ of the scenario set, then we call S a *bundling*. We note if S is a bundling, then $\sum_{\mathbf{s} \in S} \bar{p}^{\mathbf{s}}(\mathbf{x}) = 1$ holds for any decision $\mathbf{x} \in \mathcal{X}$.

The distribution shaping method extends naturally to scenario bundlings. Analogously to Section 7.1.2, we introduce the notation $\phi_{\ell}^{\mathbf{s}} = \bar{p}^{\mathbf{s}}(\bar{\mathbf{x}}^{\ell})$ for bundle probabilities corresponding to truncated decisions, and observe that we have $\phi_L^{\mathbf{s}} = \bar{p}^{\mathbf{s}}(\mathbf{x})$. If S is a bundling, then with slight modification of the distribution shaping relations (36b)–(36f) we can now reformulate (36) as

$$\min \sum_{i \in [n]} \phi_L^{\mathbf{s}} G^{\mathbf{s}} \quad (59a)$$

$$\text{s.t. } \phi_{\ell}^{\mathbf{s}} \leq \frac{\sigma_{\ell}^1}{\sigma_{\ell}^0} \phi_{\ell-1}^{\mathbf{s}} + 1 - x_{\ell}, \quad \forall \ell \in [L], s \in S : s_{\ell} = 1 \quad (59b)$$

$$\phi_{\ell}^{\mathbf{s}} \leq \frac{1 - \sigma_{\ell}^1}{1 - \sigma_{\ell}^0} \phi_{\ell-1}^{\mathbf{s}} + 1 - x_{\ell}, \quad \forall \ell \in [L], s \in S : s_{\ell} = 0 \quad (59c)$$

$$\phi_{\ell}^{\mathbf{s}} \leq \phi_{\ell-1}^{\mathbf{s}}, \quad \forall \ell \in [L], s \in S : s_{\ell} = * \quad (59d)$$

$$\phi_{\ell}^{\mathbf{s}} \leq \phi_{\ell-1}^{\mathbf{s}} + x_{\ell}, \quad \forall \ell \in [L], s \in S \quad (59e)$$

$$\sum_{s \in S} \phi_{\ell}^{\mathbf{s}} = 1, \quad \forall \ell \in [L] \quad (59f)$$

$$\phi \in [0, 1]^{S \times L}, \quad (59g)$$

$$\mathbf{x} \in \mathcal{X}, \quad (59h)$$

where $\phi_0^{\mathbf{s}} = \prod_{l: s_l=1} \sigma_l^0 \prod_{l: s_l=0} (1 - \sigma_l^0)$ denotes the baseline probability of bundle $\mathbf{s} \in S$.

If we can find a small bundling S , then the above formulation allows us to significantly reduce the problem size. Laumanns et al. (2014) show that, for the case of SPIPP with a single O-D pair, it is indeed possible to efficiently construct such small bundlings, which on a real-world instance can reduce 2^{30} scenarios to a few hundred bundles. We mention that, while for the base problem these advances translate readily to the case of multiple O-D pairs, this is no longer the case for DRO variants; see Section A.7 for a more detailed discussion.

REMARK A.1 *A more refined version of the scenario bundling technique, shown to be effective both for SPIPP and SNIP, is developed in Haus et al. (2017). The key difference is that, instead of defining a bundle by specifying the statuses of some links, this scenario aggregation method utilizes binary decision diagrams. The methodology and results outlined in Appendix A.6 can be naturally adapted to this more general scheme. However, to keep our presentation straightforward, we do not discuss details of this approach.*

A.6 DRO and scenario bundling Our goal in this section is to incorporate scenario bundling into the formulations (38), (52), and (53), in order to reduce problem sizes. The first natural step—recalling the definitions and notation from Appendix A.5—is to everywhere replace the full scenario set $[n]$ with a bundling S , and the scenario-based indexing $i \in [n]$ with bundle-based indexing $\mathbf{s} \in S$. Accordingly, we also replace the scenario probabilities π with the bundle probabilities ϕ , the distribution shaping constraints (36b)–(36f) with their bundle-based counterparts (59b)–(59g), and the scenario outcomes G^i with bundle outcomes $G^{\mathbf{s}}$. The last step is to replace the scenario distances δ^{ij} with suitably defined *bundle distances*.

For a general distance δ this last step is highly non-trivial, and typically leads to undesired consequences, such as formulations whose optimal solution depends on the particular choice of bundling. In more detail: Let S be a bundling, as defined in Appendix A.5. Any probability distribution \mathbb{P} on the scenario set then induces a distribution $\bar{\mathbb{P}}$ on the bundles, given by $\bar{p}^{\mathbf{s}} = \sum_{i \in B_{\mathbf{s}}} p^i$ for $\mathbf{s} \in S$. Let us now consider an EMD ball around $\bar{\mathbb{P}}$, based on some distance δ . In order for the scenario bundling approach to work in a DRO context, we would need to characterize the family of probability distributions on S that are induced by the elements of this EMD ball. Furthermore, if we aim to achieve this goal via a straightforward adaptation of our previous MILP formulations, the characterization should be in the form of an EMD ball around the induced distribution $\bar{\mathbb{P}}$, with respect to some distance $\bar{\delta}$ among bundles. Unfortunately, there is no general scheme to define such a bundle distance. In particular, applying natural schemes (such as a Hausdorff distance-like maximin approach, or defining the distance between two bundles as the smallest distance between any two of their respective scenarios) to our problems can lead to results that are not only inexact, but depend strongly on the particular choice of bundling. However, as the next lemma shows, adopting the total variation distance (i.e., setting δ as the discrete metric) neatly sidesteps such issues.

LEMMA A.1 *Let us denote the discrete metric on $\{0, 1\}^L$ by δ , and the discrete metric on a bundling S by $\bar{\delta}$. Then, denoting the identity map of S by $\bar{\xi}$, for any radius $\kappa > 0$, we have*

$$\mathcal{B}_{\bar{\delta}, \kappa}^{\bar{\xi}}(\bar{\mathbb{P}}) = \left\{ \bar{\mathbb{Q}} : \mathbb{Q} \in \mathcal{B}_{\delta, \kappa}^{\xi}(\mathbb{P}) \right\}. \quad (60)$$

PROOF. Let us first consider a distribution $\mathbb{Q} \in \mathcal{B}_{\delta, \kappa}^{\xi}(\mathbb{P})$. According to Lemma 5.1, there exists a corresponding solution $\gamma \in \mathbb{R}_+^{n \times n}$ to system (24). Let us now define the aggregated solution $\bar{\gamma} \in \mathbb{R}_+^{S \times S}$ by $\bar{\gamma}^{\mathbf{st}} = \sum_{i \in B_{\mathbf{s}}} \sum_{j \in B_{\mathbf{t}}} \gamma^{ij}$ for $\mathbf{s}, \mathbf{t} \in S$, and observe that for all $i \in B_{\mathbf{s}}, j \in B_{\mathbf{t}}$ we have $\bar{\delta}^{\mathbf{st}} \leq \delta^{ij}$. Then, using Lemma 5.1 again, $\bar{\mathbb{Q}} \in \mathcal{B}_{\bar{\delta}, \kappa}^{\bar{\xi}}(\bar{\mathbb{P}})$ follows from the next set of inequalities:

$$\begin{aligned} \sum_{\mathbf{t} \in S} \bar{\gamma}^{\mathbf{st}} &= \sum_{\mathbf{t} \in S} \sum_{i \in B_{\mathbf{s}}} \sum_{j \in B_{\mathbf{t}}} \gamma^{ij} = \sum_{i \in B_{\mathbf{s}}} \sum_{j \in [n]} \gamma^{ij} = \sum_{i \in B_{\mathbf{s}}} p^i = \bar{p}^{\mathbf{s}}, & \forall \mathbf{s} \in S \\ \sum_{\mathbf{s} \in S} \bar{\gamma}^{\mathbf{st}} &= \sum_{\mathbf{s} \in S} \sum_{i \in B_{\mathbf{s}}} \sum_{j \in B_{\mathbf{t}}} \gamma^{ij} = \sum_{j \in B_{\mathbf{t}}} \sum_{i \in [n]} \gamma^{ij} = \sum_{j \in B_{\mathbf{t}}} q^j = \bar{q}^{\mathbf{t}}, & \forall \mathbf{t} \in S \\ \sum_{\mathbf{s} \in S} \sum_{\mathbf{t} \in S} \bar{\delta}^{\mathbf{st}} \bar{\gamma}^{\mathbf{st}} &= \sum_{\mathbf{s} \in S} \sum_{\mathbf{t} \in S} \sum_{i \in B_{\mathbf{s}}} \sum_{j \in B_{\mathbf{t}}} \delta^{ij} \gamma^{ij} \leq \sum_{\mathbf{s} \in S} \sum_{\mathbf{t} \in S} \sum_{i \in B_{\mathbf{s}}} \sum_{j \in B_{\mathbf{t}}} \delta^{ij} \gamma^{ij} = \sum_{i \in [n]} \sum_{j \in [n]} \delta^{ij} \gamma^{ij} \leq \kappa. \end{aligned}$$

Now let us consider a distribution $\mathbb{T} \in \mathcal{B}_{\bar{\delta}, \kappa}^{\bar{\xi}}(\bar{\mathbb{P}})$, along with the corresponding $\bar{\gamma} \in \mathbb{R}_+^{S \times S}$ guaranteed by Lemma 5.1. As the bundling S represents a partition of $[n]$, for every $i \in [n]$ there exists a unique bundle $\mathbf{s}(i) \in S$ such that $i \in B_{\mathbf{s}(i)}$ holds. Using this notation, let us define the disaggregation $\gamma \in \mathbb{R}_+^{n \times n}$ by

$$\gamma^{ij} = \begin{cases} \frac{p^i p^j}{\bar{p}^{\mathbf{s}(i)} \bar{p}^{\mathbf{s}(j)}} \bar{\gamma}^{\mathbf{s}(i) \mathbf{s}(j)} & \text{if } \mathbf{s}(i) \neq \mathbf{s}(j), \\ \frac{p^i}{\bar{p}^{\mathbf{s}(i)}} \bar{\gamma}^{\mathbf{s}(i) \mathbf{s}(i)} & \text{if } i = j, \\ 0 & \text{otherwise.} \end{cases}$$

If we define the probability distribution \mathbb{Q} on the scenarios by $q^j = \sum_{i \in [n]} \gamma^{ij}$, then it is easy to verify that $\mathbb{T} = \bar{\mathbb{Q}}$. Observing that $\gamma^{ij} \neq 0$ implies $\delta^{ij} = \bar{\delta}^{\mathbf{s}^{(i)}\mathbf{s}^{(j)}}$, it is also straightforward to verify that γ is a solution of the system (24). It follows that we have $\mathbb{Q} \in \mathcal{B}_{\delta, \kappa}^{\xi}(\mathbb{P})$, which completes the proof. \square

REMARK A.2 *The above lemma remains valid if we replace the discrete metric with a reflexive scenario distance δ that depends only on the outcomes. More precisely, let us assume that for all $i, j \in [n]$ we have $\delta^{ij} = d(G^i, G^j)$, for some mapping $d : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}_+$ that satisfies $d(G, G) = 0$ for all $G \in \mathbb{R}$. Then, given two scenario bundles $\mathbf{s}, \mathbf{t} \in S$, for any $i_1, i_2 \in B_{\mathbf{s}}$ and $j_1, j_2 \in B_{\mathbf{t}}$ we have $\delta^{i_1 j_1} = d(G^{i_1}, G^{j_1}) = d(G^{i_2}, G^{j_2}) = \delta^{i_2 j_2}$. Let us denote this common value by $\bar{\delta}^{\mathbf{s}\mathbf{t}}$ to define a reflexive bundle distance $\bar{\delta} : S \times S \rightarrow \mathbb{R}_+$. The equality (60) then follows by using the same proof as before.*

A.7 Limitations of scenario bundling The SPIPP problem originally introduced in Peeta et al. (2010) features multiple O-D pairs with corresponding weights. The formulations (37), (38), (52), and (53) can all be naturally adapted to this case by having the random outcome G represent the weighted sum of shortest path lengths between these pairs. However, bundling methods can no longer be directly applied here, as links that are irrelevant to the length of a shortest path for one O-D pair will typically not be irrelevant to the lengths of shortest paths between other pairs. As proposed in Laumanns et al. (2014), it is possible to instead perform bundling separately for each O-D pair, and again use distribution shaping to express the marginal distribution of the shortest path length for each O-D pair. When solving the risk-neutral underlying problem, these marginal distributions will suffice, because—due to the linearity of expectation—the expected shortest path lengths for the individual O-D pairs can be aggregated into a global objective.

This approach is unfortunately no longer viable when working either in a risk-averse or in a DRO context. Expressing CVaR, or any more complex risk measure, of the global objective would require knowledge of the joint distribution of shortest path lengths, because the risk of a weighted sum is in general not equal to the weighted sum of individual risks. Similarly, the obstacle for DRO problems is that the ambiguity set around the joint distribution cannot be reduced to ambiguity sets around the marginal distributions. Therefore, while we might be able to find the worst-case distribution for each O-D pair, this does not give us the global worst-case distribution. The scope of the methods we outlined is thus limited to cases where either scenario bundling can effectively be applied in terms of the global objective function, or the overall number of scenarios is relatively small. Notably, in addition to the single O-D pair SPIPP problem we explored here, the “global bundling” approach is also reported (Haus et al., 2017) to provide good results for SNIP.