Incorporating convex risk measures into multistage stochastic programming algorithms

Oscar Dowson, David P. Morton
Department of Industrial Engineering and Management Sciences, Northwestern University, Evanston, IL 60208, USA, o.dowson@gmail.com, david.morton@northwestern.edu

Bernardo K. Pagnoncelli
SKEMA Business School Université Côte d'Azur, Lille, France. bernardo.pagnoncelli@skema.edu

Over the last two decades, coherent risk measures have been well studied as a principled, axiomatic way to characterize the risk of a random variable. Because of this axiomatic approach, coherent risk measures have a number of attractive features for computation, and they have been integrated into a variety of stochastic programming algorithms, including stochastic dual dynamic programming (SDDP), a common class of data-driven solution methods for multistage stochastic programs. Coherent risk measures and SDDP are tools used to manage risk while solving data-driven problems. Perhaps the most prominent example involves informing operations and deriving electricity prices in power systems with significant hydro-electric power, including the Brazilian interconnected power system. We focus on incorporating the more general class of *convex risk measures* into an SDDP algorithm, exemplifying our approach with the entropic risk measure. It is well-known that coherent risk measures lead to an inconsistency if agents care about their state at the end of the time horizon, but control risk in a stage-wise fashion. The entropic risk measure does not have this shortcoming. We illustrate the advantages of the entropic risk measure with two small examples from transportation and finance, and test the numerical viability of our adaptation of the SDDP decomposition scheme in a large-scale hydro-thermal scheduling problem using data from the Brazilian system.

Key words: stochastic dual dynamic programming; convex risk measures; entropic risk measure; data-driven operations

1. Introduction

When decision makers face uncertainty, they are often willing to accept modestly higher costs in expectation for a lower probability of bad outcomes. This is known as risk aversion. A convenient set of tools for modeling risk aversion are *risk measures*, which map random variables to real values. The class of *coherent* risk measures satisfy the axioms of Artzner et al. (1999), and they have been well studied as a principled and axiomatic way to incorporate risk aversion into optimization problems (e.g., Artzner et al. 1999, Riedel 2004, Ruszczyński and Shapiro 2006a, Ruszczyński 2010, Philpott and de Matos 2012, Philpott et al. 2013, Pichler and Shapiro 2021, Shapiro et al. 2013, Guigues 2016).

The monograph of Shapiro et al. (2021) provides a thorough treatment of risk-averse optimization problems, with applications to inventory and portfolio selection problems. We also point to applications in transportation (e.g., Mohammadi et al. 2021), mining (e.g., Reus et al. 2019), forestry (e.g., Bushaj et al. 2022), and finance (e.g., Guastaroba et al. 2020). Recent work has also focused on so-called *distributionally robust* risk measures, which under specific constructs form coherent risk measures via duality (e.g., Rahimian and Mehrotra 2019, Shapiro 2021).

Dynamic (multistage) problems, which are the focus of our work, are more challenging and present additional difficulties relative to the static case. Most work implementing risk aversion within a multistage stochastic programming framework is in the spirit of dynamic programming, based on the theory of dynamic risk measures (Ruszczyński and Shapiro 2006a). Most commonly, the conditional expectations at each stage are replaced by a coherent risk measure, and then the stochastic dual dynamic programming algorithm (SDDP) (Pereira and Pinto 1991) is applied recursively at each stage (e.g., Shapiro et al. 2013, Philpott et al. 2013, Kozmík and Morton 2015, Guigues 2016). At each level of the recursion, which typically corresponds to an interval in time, dynamic risk measures compute the risk of a random variable, conditioned on the history of the stochastic process up to that point in time. Because of this recursive form, dynamic risk measures are also called *nested* risk measures.

The SDDP algorithm and coherent risk measures are used in practice to manage operational risk in a range of industrial settings, exemplified by their usage for over 20 years in the Brazilian power system; see, e.g., Maceiral et al. (2018). Moreover, as we will show, a key feature of the SDDP algorithm is its ability to combine risk measures with a data-driven approach to the underlying uncertainty.

In some sequential decision-making problems, we wish to avoid bad outcomes at key points in time while otherwise ignoring the trajectory of outcomes along the way. For example, we care primarily about the value of a pension fund at the age of retirement, as opposed to its position in earlier years and the volatility of the journey. However, for reasons of tractability, current methods cannot solve large-scale instances of these *end-of-horizon* formulations with typical risk measures. The method of Baucke et al. (2018) is one exception, but it requires the expensive computation of a deterministic upper bound.

A given risk measure can be employed in an end-of-horizon formulation or in a nested formulation. As discussed in Kovacevic and Pflug (2009), if these two formulations lead to different risk preferences then the risk measure lacks information monotonicity. It is well-known that, in general, coherent risk measures lack this property, although two exceptions are the risk-neutral expectation and the worst-case risk measure.

The failure to satisfy information monotonicity has led some authors to question the validity and interpretability of using coherent risk measures in a multistage stochastic program (e.g., Homem-de-Mello and Pagnoncelli 2016, Pflug and Pichler 2016, Baucke et al. 2018). The work of Pflug and Pichler (2016) and Baucke et al. (2018) are particularly notable, because they describe a way of dynamically changing the properties of a coherent risk measure based on the history of the random variable in order to recover conditional consistency. We also point to the work of Asamov and Ruszczyński (2015), who study bounds and approximations related to the equivalence between end-of-horizon and nested formulations.

Kovacevic and Pflug (2009) indicate that the entropic risk measure—a specific risk measure which we define in Section 2.1 and which belongs to the more general class of *convex* risk measures—is information monotonic. Kupper and Schachermayer (2009) show that the entropic measure is the only translation invariant and time consistent dynamic monotonic mapping. These results motivate our paper: If we can incorporate convex risks measures into a tractable algorithm for solving multistage stochastic programs, then we can avoid the validity and interpretability concerns associated with coherent risk measures. In addition to focusing on the entropic risk measure, we also focus on SDDP because it is a tractable solution algorithm, which already supports coherent risk measures. As a result, the main contributions of this paper are:

i) To provide a computationally convenient means to represent any convex risk measure within a multistage stochastic programming problem.

ii) To provide a practical and efficient way of solving multistage stochastic programming problems using the entropic risk measure within the SDDP algorithm.

Much of the theory surrounding the entropic risk measure is well known (e.g., Detlefsen and Scandolo 2005, Frittelli and Gianin 2005, Ruszczyński and Shapiro 2006b, Ben-Tal and Teboulle 2007, Föllmer and Schied 2010), largely motivated by applications in finance and economics (e.g., Jobert and Rogers 2008, Chen and Yang 2011, Wei and Hu 2014), as well as robotics (e.g., Nass et al. 2019) and reinforcement learning (e.g., Russel et al. 2020). However, despite this fact, computational work in risk-averse stochastic programming has largely been limited to coherent risk measures. Therefore, our main goal for this paper is to derive, and demonstrate, the theory that facilitates efficient algorithms for solving multistage stochastic programming problems under the entropic risk measure. For that reason, we avoid repeating established proofs of properties of convex risk measures, and only describe what is needed to advance our algorithmic tools.

The rest of this paper is laid out as follows. In Section 2, we briefly review risk measures and the notion of *conditional consistency.* In Section 3, we discuss the dual representation of convex risk measures, and show how to compactly reformulate the entropic risk measure using conic duality. In Section 4, we show how to incorporate a convex risk measure in a stochastic dual dynamic programming algorithm as a solution method for risk-averse multistage stochastic programs. In Section 5, we illustrate our findings with two examples that are small enough to transparently highlight the drawbacks of using conditional value-at-risk in terms of information monotonicity, and show that the entropic risk measure is a viable alternative in those cases. We conclude with experiments for a large-scale hydro-thermal scheduling problem based on data from the Brazilian power system, which shows the computational tractability of our approach.

2. Risk measures

In this section we briefly review risk measures and the notion of conditional consistency.

2.1. Static risk measures

Consider a probability space $(\Omega, \mathcal{F}, P)$ with $\mathcal{Z}$ denoting a space of $\mathcal{F}$-measurable functions from Ω to $\mathbb{R}$. It is natural to restrict $\mathcal{Z}$ to an appropriately defined Orlicz heart, an analog of L^p spaces under which the requisite moment generating functions are finite; see Cheridito and Li (2009) and Ahmadi-Javid and Pichler (2017). That said, it is common in the dynamic case to make the more restrictive assumption that $\mathcal{Z}$ is a conditional L^∞ space; see, for example, Cheridito and Kupper (2013) for details.

A *risk measure* is a function $\mathbb{F}$, which maps a random variable $Z \in \mathcal{Z}$ to $\mathbb{R}$. A risk measure is said to be *coherent* if, for $Z_1, Z_2 \in \mathcal{Z}$, it satisfies the following axioms of Artzner et al. (1999):

Axiom 1. Monotonicity: $Z_1 \leq Z_2$, a.s. $\implies \mathbb{F}[Z_1] \leq \mathbb{F}[Z_2]$.

Axiom 2. Translation equivariance: $\mathbb{F}[Z + a] = \mathbb{F}[Z] + a$ for all $a \in \mathbb{R}$.

Axiom 3. Sub-additivity: $\mathbb{F}[Z_1 + Z_2] \leq \mathbb{F}[Z_1] + \mathbb{F}[Z_2]$.

Axiom 4. Positive homogeneity: $\mathbb{F}[aZ] = a\,\mathbb{F}[Z]$ for $a \geq 0$.

Commonly used coherent risk measures often involve *conditional value-at-risk* ($\mathbb{CV@R}$):

$$\mathbb{CV@R}_\gamma[Z] = \inf_{\zeta \in \mathbb{R}} \left\{ \zeta + \frac{1}{1-\gamma}\, \mathbb{E}[(Z - \zeta)_+] \right\}, \tag{1}$$

where $(x)_+ = \max\{0, x\}$ and $\gamma \in [0, 1)$. $\mathbb{CV@R}$ involves the single parameter, γ, and has two extremes: when $\gamma = 0$, $\mathbb{CV@R}_0[Z] = \mathbb{E}[Z]$, and $\lim_{\gamma \to 1} \mathbb{CV@R}_\gamma[Z] = \operatorname{ess\,sup}[Z]$. $\mathbb{CV@R}_\gamma$ can be interpreted as the expectation of the worst $1-\gamma$ fraction of outcomes; see Rockafellar and Uryasev (2002).

A slightly more general class of risk measures relaxes axioms 3 and 4 of *sub-additivity* and *positive homogeneity* for coherent risk measures, replacing them with *convexity*:

Axiom 5. Convexity: $\mathbb{F}[aZ_1 + (1-a)Z_2] \leq a\,\mathbb{F}[Z_1] + (1-a)\,\mathbb{F}[Z_2]$ for all $a \in [0, 1]$.

Definition 1. A risk measure is *convex* if it satisfies the axioms of monotonicity, translation equivariance, and convexity.

The most commonly used example of a convex risk measure that is not coherent is the entropic risk measure:

$$\mathbb{ENT}_\gamma[Z] = \frac{1}{\gamma}\log\left(\mathbb{E}[e^{\gamma Z}]\right), \tag{2}$$

which again involves a single parameter, $\gamma > 0$. Like $\mathbb{CV@R}$, in the limiting cases $\lim_{\gamma\to 0}\mathbb{ENT}_\gamma[Z] = \mathbb{E}[Z]$ and $\lim_{\gamma\to\infty}\mathbb{ENT}_\gamma[Z] = \operatorname{ess\,sup}[Z]$.

Every coherent risk measure is a convex risk measure; see, e.g., Föllmer and Schied (2002) and Rockafellar (2007). Justifications for relaxing to a convex risk measure concern both positive homogeneity and sub-additivity. The linear scaling induced by the former does not penalize concentration of risk. As discussed in Balbás et al. (2009), by increasing the holdings of a specific instrument it is reasonable to assume that quantities above some threshold introduce liquidity risk. Shortcomings associated with assuming subadditivity are the subject of work by Dhaene et al. (2008), who show that in the context of capital solvency, risk can increase by a merger. They relax subadditivity and propose a special property called a "regulator's condition" to avoid an increase in terms of shortfall risk when a merger occurs. Finally, Cominetti and Torrico (2016) and Brandtner et al. (2018) show that decisions under the (convex) entropic risk measure are not sensitive to independent background risk.

2.2. Conditional consistency in dynamic risk measures

So far we have discussed *single-period* risk measures, which measure a single random variable, Z. However, the subject of our interest is multistage stochastic programming, in which a sequence of correlated real-valued random variables, $Z = \{Z_t\}_{t=1}^T$, is induced by a policy that we select based on exogenous randomness that evolves over time. To this end we consider a probability space $(\Omega, \mathcal{F}, P)$, and we let $\mathcal{F}_1 \subset \mathcal{F}_2 \subset \cdots \subset \mathcal{F}_T$ be sub sigma-algebras of $\mathcal{F}$ that form a filtration such that $\mathcal{F}_t$ corresponds to the information available through stage t. As boundary conditions, we have $\mathcal{F}_1 = \{\emptyset, \Omega\}$ and $\mathcal{F}_T = \mathcal{F}$.

There are two common ways in the literature to measure the risk of this sequence of random variables. We refer to the first as the *end-of-horizon* approach:

$$\text{End-of-Horizon-Risk}(Z) = \mathbb{F}[Z_1 + Z_2 + \cdots + Z_T].$$

A second approach uses *nested* risk measures (Riedel 2004, Ruszczyński and Shapiro 2006a, Shapiro et al. 2021, Ruszczyński 2010, Shapiro et al. 2013), which in the Markovian setting can be written:

$$\text{Nested-Risk}(Z) = \mathbb{F}\left[Z_1 + \mathbb{F}_{Z_2|Z_1}[Z_2 + \mathbb{F}_{Z_3|Z_2}[\cdots + \mathbb{F}_{Z_T|Z_{T-1}}[Z_T]]\right].$$

We have that $\mathbb{F}_{Z_t|Z_{t-1}}$ is a mapping from the space of measurable functions with respect to $\mathcal{F}_t$ into the space of measurable functions with respect to $\mathcal{F}_{t-1}$. We link the end-of-horizon view of risk with the nested view of risk via the notion of *conditional* consistency.

Definition 2. Let (X_1, X_2) and (Y_1, Y_2) be two-dimensional vectors for which the requisite expectations are finite. A risk measure $\mathbb{F}$ is said to be *conditionally consistent* if:

$$\mathbb{F}[X_1 + X_2] \leq \mathbb{F}[Y_1 + Y_2] \iff \mathbb{F}[X_1 + \mathbb{F}_{X_2|X_1}[X_2]] \leq \mathbb{F}[Y_1 + \mathbb{F}_{Y_2|Y_1}[Y_2]].$$

In other words, conditional consistency of a risk measure says that preferring the sequence X_1, X_2 to that of Y_1, Y_2 from an end-of-horizon perspective is equivalent to preferring X_1, X_2 from a nested perspective.

Remark 1. Variations of conditional consistency have been defined and explored in a number of papers under names such as information monotonicity and time consistency (Kovacevic and Pflug 2009, 2014). We choose to present this new definition in a simplified setting of additive random variables because it most-closely matches the structure required by multistage stochastic programs. Moreover, we use the term *conditionally* consistent because time consistency is arguably an overloaded term-of-art that has been used to describe both the consistency of nested risk measures and the consistency of decision making in dynamic programs, and our definition is more limited than a proper treatment of information monotonicity requires.

It is well known in the literature that commonly used coherent risk measures, such as $\mathbb{CV@R}$, are not conditionally consistent (Homem-de-Mello and Pagnoncelli 2016, Pflug and Pichler 2016, Baucke et al. 2018). However, since

$$\mathbb{ENT}_\gamma[X + Y] = \mathbb{ENT}_\gamma[X + \mathbb{ENT}_{\gamma, Y|X}[Y]], \tag{3}$$

which is referred to as *recursivity* in Kovacevic and Pflug (2009), we can easily see that the entropic risk measure is conditionally consistent. There are two coherent risk measures that also satisfy this property, and they are expectation and worst-case. It is also straightforward to show that, in general, recursivity does not hold for any strict convex combination of the expectation and worst-case measures. This might seem surprising given that the expectation and worst-case measures are the limiting cases of $\mathbb{CV@R}$, and given that coherent risk measures have properties that might seem consistent with such a result. A related result, due to Kupper and Schachermayer (2009), shows that given a family of risk measures $(\rho_t)_{t\in\mathbb{N}}$ the *only* convex dynamic risk measure which is law invariant, time consistent and relevant is the entropic.

3. Representation and visualization of convex risk measures

In this section we discuss alternative representations for convex risk measures, which are useful for numerical computation. We start by considering the dual representation of a risk measure, which has a closed-form solution under the entropic measure. In addition, we will see that dual forms allow us to obtain a graphical illustration of why nested $\mathbb{CV@R}$ may not be suitable for multistage problems. Because of our computational focus, we use a data-driven approach and assume from now on that Ω is finite, where Ω is often composed of historical observations of data. In such a setting, the underlying probability mass function is typically $p_\omega = 1/|\Omega|$ for all $\omega \in \Omega$, but we allow for the more general mass function.

3.1. The dual form

Since we assume Ω is finite, we denote a random variable with the uppercase Z, realizations by $Z(\omega) = z_\omega$, and respective positive probabilities by p_ω, $\omega \in \Omega$. Every coherent risk measure $\mathbb{F}$ is associated with a convex *risk set* $\mathcal{M}(p) \subseteq \mathcal{P}$, where

$$\mathcal{P} = \left\{ q \in \mathbb{R}^{|\Omega|} \;\middle|\; q \geq 0, \sum_{\omega \in \Omega} q_\omega = 1 \right\},$$

such that:

$$\mathbb{F}[Z] = \sup_{q \in \mathcal{M}(p)} \mathbb{E}_q[Z]. \tag{4}$$

This representation is called the *dual form*; see Artzner et al. (1999). As an example, the risk set associated with $\mathbb{CV@R}_\gamma$, where again $\gamma \in [0,1)$, is:

$$\mathcal{M}(p) = \left\{ q \in \mathcal{P} \;\middle|\; q_\omega \leq \frac{p_\omega}{1-\gamma}, \quad \omega \in \Omega \right\}. \tag{5}$$

Convex risk measures have a similar dual formulation, differing from coherent risk measures by a penalty term $\alpha(q)$:

$$\mathbb{F}[Z] = \sup_{q \in \mathcal{M}(p)} \left\{ \mathbb{E}_q[Z] - \alpha(q) \right\}. \tag{6}$$

A convex risk measure is coherent if and only if $\alpha(\cdot) = 0$. Due to the convexity of $\mathbb{F}$ with respect to Z, it can be shown that α is a lower semi-continuous convex function of q, and thus the dual problem is concave with respect to q; see Föllmer and Schied (2002).

Föllmer and Schied also show that for the entropic risk measure we have $\mathcal{M}(p) = \mathcal{P}$ and

$$\alpha(q) = \frac{1}{\gamma} \sum_{\omega \in \Omega} q_\omega \log\left(\frac{q_\omega}{p_\omega}\right).$$

Given realizations z_ω with nominal probabilities p_ω, $\omega \in \Omega$, we can compute the probability distribution that achieves the supremum as follows:

$$
\begin{aligned}
\mathbb{ENT}_\gamma[Z] = \max_q \; & \sum_{\omega\in\Omega} z_\omega q_\omega - \frac{1}{\gamma}\sum_{\omega\in\Omega} q_\omega \log\left(\frac{q_\omega}{p_\omega}\right) \\
\text{s.t.} \; & \sum_{\omega\in\Omega} q_\omega = 1 \\
& q_\omega \geq 0, \quad \omega \in \Omega.
\end{aligned} \tag{7}
$$

It is possible to derive a closed-form expression for the optimal vector, q^*.

Theorem 1. *Let $\gamma > 0$ and $p \in \mathcal{P}$ satisfy $p > 0$. Then the unique optimal solution to problem* (7) *is given by:*

$$
q^*_\omega = \frac{p_\omega e^{\gamma z_\omega}}{\sum\limits_{\omega\in\Omega} p_\omega e^{\gamma z_\omega}}, \quad \forall \omega \in \Omega. \tag{8}
$$

Proof of Theorem 1. The unique optimal solution is obtained on the convex and compact feasible region of problem (7) because the objective function is strictly concave. Neglecting the non-negativity constraints, forming the Lagrangian, differentiating and solving for q^*_ω yields $q^*_\omega \propto p_\omega e^{\gamma z_\omega} > 0$. The fact that the probability mass function must sum to one yields the proportionality constant in the theorem's statement. Q.E.D.

Remark 2. Expression (8) for the maximizer of the dual representation of the entropic risk measure can be seen as a softmax function scaled by the probability mass, i.e., as a softmax of the terms $\gamma z_\omega + \ln(p_\omega)$, $\omega \in \Omega$.

Remark 3. The closed-form expression in Theorem 1 is an analytical result. In practice, we often compute e^x using 64-bit IEEE 754 floating-point representation (IEEE 2019). The largest value representable using 64-bit floating point is $\approx 1.8 \times 10^{308} \approx e^{709}$. Therefore, for large values of z_ω (e.g., if the random variable is a financial return in the millions of dollars), computing q^*_ω requires arbitrary precision arithmetic. The dual formulation (7) does not have this numerical issue, but it requires the solution of a nonlinear program, e.g., by an off-the-shelf solver such as Ipopt (Wächter and Biegler 2006). In practice, we have found computing q^*_ω using arbitrary precision arithmetic to be orders of magnitude faster than solving the nonlinear program.

Remark 4. Confirming the dual representation of the entropic risk measure, if we substitute q^*_ω into $\mathbb{E}_q[Z] - \frac{1}{\gamma}\sum_{\omega\in\Omega} q_\omega \log\left(\frac{q_\omega}{p_\omega}\right)$, we obtain the optimal objective value $\frac{1}{\gamma}\log\left(\sum_{\omega\in\Omega} p_\omega e^{\gamma z_\omega}\right)$, which is the primal definition of the entropic risk measure.

We now apply the dual form of the entropic risk measure in a simple example, which allows us to contrast it with the $\mathbb{CV@R}$ risk measure. Consider two independent random variables, Z_1 and Z_2, with uniform (empirical) distributions on their respective supports, $\{0, 4\}$ and $\{1, 2, 3, 4\}$.

Thus $Z_1 + Z_2$ is a random integer between 1 and 8 with uniform probability. In Figure 1, we plot the probability mass that attains the supremum over the dual set associated with: (i) expectation, (ii) the entropic risk measure with $\gamma = 0.4$, (iii) the end-of-horizon $\mathbb{CV@R}$ with $\gamma = 0.4$, and (iv) the nested $\mathbb{CV@R}$ with $\gamma = 0.4$. Note that the interpretation of γ is different for the entropic and $\mathbb{CV@R}$ risk measures; we use $\gamma = 0.4$ in both cases because it is a reasonable value of γ for the entropic risk measure in this particular example.

The risk-adjusted probabilities of the nested $\mathbb{CV@R}$ are calculated in the following way. First, we compute the adjusted probabilities for Z_2 conditioned on each outcome of Z_1. Then, using the $\mathbb{CV@R}$ of each conditional subset, we compute the adjusted probabilities for Z_1. Finally, we multiply the adjusted probabilities of Z_2 conditioned on Z_1 by the adjusted probabilities of Z_1.

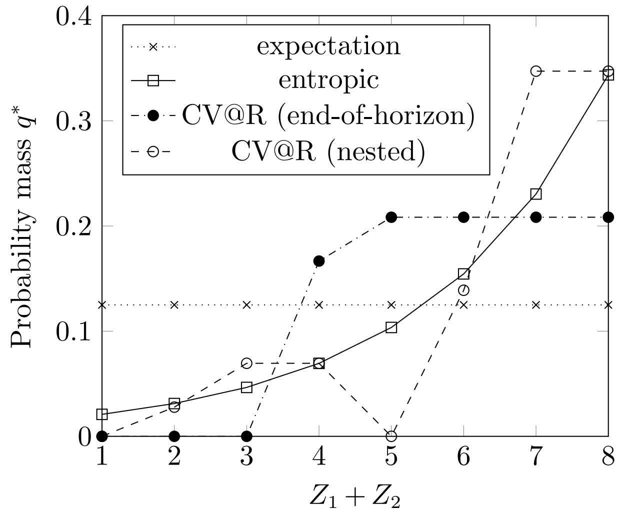

Figure 1 Probability mass q that achieves the supremum of the dual representation for four risk measures.

The results in Figure 1 clearly show the undesirability of nested $\mathbb{CV@R}$ from an end-of-horizon perspective. The main argument is that nested $\mathbb{CV@R}$ assigns probabilities that are not monotonic in the value of $Z_1 + Z_2$. According to Definition 2, conditional consistency does not hold for the nested $\mathbb{CV@R}$ because there is more probability on $Z_1 + Z_2 = 3$ (≈ 0.069) than $Z_1 + Z_2 = 5$ (0.0). In contrast, the end-of-horizon $\mathbb{CV@R}$ *does* assign non-decreasing probabilities as function of $Z_1 + Z_2$. Unlike $\mathbb{CV@R}$, the entropic risk measure assigns probabilities in a (strictly) non-decreasing and convex function of $Z_1 + Z_2$, which we argue is a more "natural" way of assigning the distribution, because undesirable outcomes are given increasing weight in the risk-adjusted expectation.

3.2. Conic dual of entropic risk measure

In Section 3.1 we dealt with the dual form of a convex risk measure. We now take the dual of the corresponding maximization problem to create a minimization problem that will show another link to data-driven distributionally robust optimization and provide us with a tighter formulation for use in Section 4. While we again focus on the entropic risk measure, the same process could be repeated for other convex risk measures.

THEOREM 2. *Let $p \in \mathcal{P}$ satisfy $p > 0$. The conic dual of problem (7) is:*

$$\begin{aligned}\mathbb{ENT}_\gamma[Z] = \min_{\mu \in \mathbb{R}^{|\Omega|+1}} & \sum_{\omega\in\Omega} p_\omega \mu_\omega + \mu_0 \\ \text{s.t. } & (-\tfrac{1}{\gamma}, \mu_0 - z_\omega, \mu_\omega) \in \mathcal{K}^*_{\exp}, \quad \omega \in \Omega,\end{aligned}$$

where the dual exponential cone is $\mathcal{K}^*_{\exp} = \left\{(u,v,w) \in \mathbb{R}^3 : -ue^{\frac{v}{u}} \le e^1 w,\ u < 0\right\}$.

Proof of Theorem 2. The exponential cone is: $\mathcal{K}_{\exp} = \left\{(x,y,z) \in \mathbb{R}^3 : ye^{\frac{x}{y}} \le z,\ y > 0\right\}$, and the dual cone of $\mathcal{K}_{\exp}$ is $\mathcal{K}^*_{\exp}$ as given in the theorem's statement. The relative entropy cone is: $\mathcal{K}_{\text{rel-ent}} = \left\{(t,p,q) \in \mathbb{R}^3 : t \ge q \log\left(\frac{q}{p}\right),\ p > 0,\ q > 0\right\}$, and the exponential and relative entropy cones are related as follows: $(t,p,q) \in \mathcal{K}_{\text{rel-ent}} \iff (-t,q,p) \in \mathcal{K}_{\exp}$; see, e.g., Chandrasekaran and Shah (2016).

Using these relations, we re-write problem (7) in conic form as:

$$\begin{aligned}\mathbb{ENT}_\gamma[Z] = \max_{q,t} & \sum_{\omega\in\Omega} z_\omega q_\omega - \frac{1}{\gamma}\sum_{\omega\in\Omega} t_\omega \\ \text{s.t. } & 1 - \sum_{\omega\in\Omega} q_\omega \in \{0\} \quad [\mu_0] \\ & (-t_\omega, q_\omega, p_\omega) \in \mathcal{K}_{\exp}, \quad [\mu_\omega] \quad \omega \in \Omega.\end{aligned} \tag{9}$$

Here, μ_0 is the scalar dual variable associated with the first constraint, and μ_ω is a three-dimensional dual variable associated with each exponential cone constraint, $\omega \in \Omega$. We denote the i^{th} component of μ_ω by $\mu_\omega^{(i)}$.

From conic duality (see, e.g., Boyd and Vandenberghe 2004) a maximization problem of form:

$$\begin{aligned}\max_{x\in\mathbb{R}^n} & \ a_0^\top x + b_0 \\ \text{s.t. } & A_i x + b_i \in \mathcal{K}_i, \quad i = 1,\dots,m,\end{aligned}$$

has a dual in minimization form:

$$\begin{aligned}\min_{y\in\mathbb{R}^m} & \sum_{i=1}^m b_i^\top y_i + b_0 \\ \text{s.t. } & a_0 + \sum_{i=1}^m A_i^\top y_i = 0 \\ & y_i \in \mathcal{K}_i^*, \quad i = 1,\dots,m,\end{aligned}$$

where $\mathcal{K}_i^*$ is the dual cone of $\mathcal{K}_i$. Therefore, taking the dual of (9), we obtain:

$$\begin{aligned}\mathbb{ENT}_\gamma[Z] = \min_{\mu} & \sum_{\omega\in\Omega} p_\omega \mu_\omega^{(3)} + \mu_0 \\ \text{s.t. } & z_\omega - \mu_0 + \mu_\omega^{(2)} = 0, \quad \omega \in \Omega \\ & -\tfrac{1}{\gamma} - \mu_\omega^{(1)} = 0, \quad \omega\in\Omega \\ & \mu_0 \in \mathbb{R} \\ & (\mu_\omega^{(1)}, \mu_\omega^{(2)}, \mu_\omega^{(3)}) \in \mathcal{K}^*_{\exp}, \quad \omega\in\Omega.\end{aligned}$$

Simplifying, we obtain:

$$\begin{aligned} \mathbb{ENT}_\gamma[Z] = \min_{\mu} \ & \sum_{\omega\in\Omega} p_\omega \mu_\omega + \mu_0 \\ \text{s.t.} \ & (-\tfrac{1}{\gamma}, \mu_0 - z_\omega, \mu_\omega) \in \mathcal{K}^*_{\exp}, \quad \omega \in \Omega. \end{aligned} \tag{10}$$

Finally, note that conic strong duality holds because (10) has a strictly feasible solution and (9) has a finite optimal solution, and so the optimal objective value of (9) is equal to that of (10). Q.E.D.

Remark 5. The dual problem can be efficiently solved as a conic program using off-the-shelf commercial solvers such as Mosek (MOSEK ApS 2020), or as a nonlinear program using solvers such as Ipopt (Wächter and Biegler 2006). There are alternative reformulations of the conic dual problem. For example:

$$\mathbb{ENT}_\gamma[Z] = \min_{\mu} \left\{ \mu_0 \;\middle|\; \sum_{\omega\in\Omega} p_\omega \mu_\omega \le 1,\ (\gamma(z_\omega - \mu_0), 1, \mu_\omega) \in \mathcal{K}_{\exp},\ \omega \in \Omega \right\},$$

which we derive directly from (2), or:

$$\mathbb{ENT}_\gamma[Z] = \min_{\mu} \left\{ \mu_0 + \sum_{\omega\in\Omega} p_\omega \mu_\omega \;\middle|\; (z_\omega - \mu_0, \frac{1}{\gamma}, e^1 \mu_\omega) \in \mathcal{K}_{\exp},\ \omega \in \Omega \right\}.$$

The relative performance of each formulation is likely problem- and solver-specific, depending, for example, on the efficiency with which the solver can handle the exponential dual cone $\mathcal{K}^*_{\exp}$. We use the view in Theorem 2 because it arises most naturally from conic duality, and in our anecdotal experience, the $\mathcal{K}^*_{\exp}$ formulation solves most quickly.

Remark 6. In a pleasing symmetry, Bayraksan and Love (2015) show that the Lagrangian dual of the entropic value-at-risk, i.e., the distributionally robust risk measure using the Kullback-Leibler divergence, is:

$$\mathbb{ENT}\,\mathrm{V@R}_\gamma[Z] = \min_{\mu_0} \left\{ \mu_0 + \frac{1}{\gamma}\varepsilon + \frac{1}{\gamma} \sum_{\omega\in\Omega} p_\omega \left(e^{\gamma(z_\omega - \mu_0)} - e^0 \right) \right\}. \tag{11}$$

The view in Theorem 2 can be re-arranged to obtain:

$$\mathbb{ENT}_\gamma[Z] = \min_{\mu_0} \left\{ \mu_0 + \frac{1}{\gamma} \sum_{\omega\in\Omega} p_\omega \frac{e^{\gamma(z_\omega - \mu_0)}}{e^1} \right\}. \tag{12}$$

These formulations differ in two terms: (i) equation (11) includes the additive constant ε/γ and (ii) equation (11) subtracts e^0 (the exponentiated additive identity), whereas in (12), we divide by e^1 (the exponentiated multiplicative identity).

Equation (12) offers insight into the interpretation of the dual, since it resembles the definition of $\mathbb{CV@R}$ in equation (1). With the change of notation, $\mu_0 = \zeta$, we obtain:

$$\mathbb{ENT}_\gamma[Z] = \inf_{\zeta\in\mathbb{R}} \left\{ \zeta + \frac{1}{\gamma e^1} \mathbb{E}\left[e^{\gamma(Z-\zeta)} \right] \right\}. \tag{13}$$

Thus, the dual variable μ_0 can be interpreted as an analog of the quantile variable in $\mathbb{CV@R}$.

Equation (13) demonstrates that the entropic risk measure belongs to the class of optimized certainty equivalent risk measures (Ben-Tal and Teboulle 2007, Vinel and Krokhmal 2017). Other convex risk measures can be obtained through the choice of an appropriate utility function $u(x)$ in the following:

$$\mathbb{F}[Z] = \inf_{\zeta\in\mathbb{R}} \{\zeta + \mathbb{E}[u(Z-\zeta)]\}.$$

See Ben-Tal and Teboulle (2007) for details.

4. Risk-averse stochastic programming

We now introduce convex risk measures into two-stage and multi-stage stochastic linear programs, solved using Benders' decomposition, known as the L-shaped method for two-stage problems (Van Slyke and Wets 1969), and solved using a multi-stage variant known as a stochastic dual dynamic programming (Pereira and Pinto 1991).

4.1. Two-stage stochastic programs

We formulate a two-stage stochastic linear program in the first-stage variables as:

$$\begin{aligned} V_1 = \min_{x_1}\ & c_1^\top x_1 + \mathbb{F}[V_2(x_1,\omega)] \\ & A_1 x_1 = b_1 \\ & x_1 \geq 0, \end{aligned} \tag{14}$$

where the second-stage problem is:

$$\begin{aligned} V_2(x_1,\omega) = \min_{\bar{x},x_2}\ & c_2^\top x_2 \\ & \bar{x} = x_1 \quad [\lambda] \\ & A_2 x_2 + B_2\bar{x} = b_2 \\ & x_2 \geq 0. \end{aligned} \tag{15}$$

Any of the vectors and matrices, c_2, A_2, B_2, and b_2, may depend on $\omega\in\Omega$, and we again assume the corresponding vector of random parameters has finite support, where again it is common in practice to follow a data-driven approach to create Ω from historical data. Note that λ is a dual variable on the *fishing* constraint $\bar{x} = x_1$, and is therefore a valid subgradient of $V_2(x_1,\omega)$ with respect to x_1. We assume the first-stage feasible region is a non-empty polytope, and that the second-stage problem is feasible, and has a finite optimal solution, given any x_1 feasible in problem (14) and given any $\omega\in\Omega$.

Clearly, $V_2(x_1,\omega)$ is convex with respect to x_1 for fixed ω. Moreover, by the convexity and monotonicity of $\mathbb{F}$, we have that $\mathbb{F}[V_2(x_1,\omega)]$ is also a convex function. Therefore, it can be approximated by the point-wise maximum of a collection of linear functions called *cuts* via a master problem:

$$\begin{aligned} \text{Single-}V_1^K = \min_{x_1,\Theta}\ & c_1^\top x_1 + \Theta \\ & A_1 x_1 = b_1 \\ & x_1 \geq 0 \\ & \Theta \geq \alpha_k + \beta_k^\top x_1, \quad k = 1,\ldots,K-1 \\ & \Theta \geq -M, \end{aligned} \tag{16}$$

where M is sufficiently large so that the final constraint provides a lower bound on $\mathbb{F}[V_2(\cdot,\omega)]$.

As in standard Benders' decomposition, the cuts are created iteratively. In the forward step, the master problem (16) is solved to obtain a feasible x_1. Then, the second-stage problems (15) are solved for each $\omega \in \Omega$. In the backward step, dual solutions from the second-stage subproblems are used to derive a cut using the process we outline next.

To compute the coefficients of the cuts, we rely on the following theorem, which is similar to a result for coherent risk measures in Philpott et al. (2013, Proposition 4) and is a special case of Danskin's Theorem; see, e.g., Shapiro et al. (2021, Theorem 6.11).

THEOREM 3. *Let $\omega \in \Omega$ index a random vector with finite support and with nominal probability mass function, $p \in \mathcal{P}$, which satisfies $p > 0$. Consider a convex risk measure, $\mathbb{F}$, with a convex risk set, $\mathcal{M}(p)$, so that $\mathbb{F}$ can be expressed as in equation (6). Let $V(x,\omega)$ be convex with respect to x for all fixed $\omega \in \Omega$. Let $\lambda(\tilde{x},\omega)$ be a subgradient of $V(x,\omega)$ with respect to x at $x = \tilde{x}$ for each $\omega \in \Omega$. Then, $\sum_{\omega\in\Omega} q^*_\omega \lambda(\tilde{x},\omega)$ is a subgradient of $\mathbb{F}[V(x,\omega)]$ at $\tilde{x}$, where $q^* \in \arg\max_{q\in\mathcal{M}(p)} \{\mathbb{E}_q[V(\tilde{x},\omega)] - \alpha(q)\}$.*

Proof of Theorem 3. By equation (6), $q^* \in \mathcal{M}(p)$, and by the subgradient inequality we have

$$\begin{aligned}
\mathbb{F}[V(x,\omega)] &= \sup_{q\in\mathcal{M}(p)} \{\mathbb{E}_q[V(x,\omega)] - \alpha(q)\} \\
&\geq \mathbb{E}_{q^*}[V(x,\omega)] - \alpha(q^*) \\
&= \sum_{\omega\in\Omega} q^*_\omega V(x,\omega) - \alpha(q^*) \\
&\geq \sum_{\omega\in\Omega} q^*_\omega \left(V(\tilde{x},\omega) + \lambda(\tilde{x},\omega)^\top (x-\tilde{x})\right) - \alpha(q^*) \\
&= \sum_{\omega\in\Omega} q^*_\omega V(\tilde{x},\omega) - \alpha(q^*) + \sum_{\omega\in\Omega} q^*_\omega \lambda(\tilde{x},\omega)^\top (x-\tilde{x}) \\
&= \mathbb{F}[V(\tilde{x},\omega)] + \left(\sum_{\omega\in\Omega} q^*_\omega \lambda(\tilde{x},\omega)\right)^\top (x-\tilde{x}).
\end{aligned}$$

Q.E.D.

Using Theorem 3, we obtain Algorithm 1 to generate Benders' cuts for the master problem (16), which approximates model (14). It is worth highlighting that for a general convex risk measure computing a cut requires solving $|\Omega|$ linear programs and one additional, potentially nonlinear, optimization problem to compute q^*, unless a closed-form expression is available as in Theorem 1.

The master problem (16) and cut-generation procedure just outlined form a *single-cut* algorithm because one cut is added at the k^{th} Benders' iteration to approximate $\mathbb{F}[V_2(x_1,\omega)]$ at $x_1 = x_1^k$. An

Algorithm 1: Risk-averse cut generator at x_1^k.

Given x_1^k at iteration k

for $\omega \in \Omega$ **do**

  solve subproblem (15) with $x_1 = x_1^k$ to obtain $V_2(x_1^k, \omega)$ and an extreme point dual solution, λ

  set $V_\omega^k = V_2(x_1^k, \omega)$

  set $\lambda_\omega^k = \lambda$

end

set $q^k \in \arg\max_{q \in \mathcal{M}(p)} \left\{ \mathbb{E}_q[V_\omega^k] - \alpha(q) \right\}$

set $\beta_k = \sum_{\omega \in \Omega} q_\omega^k \lambda_\omega^k$

set $\alpha_k = \sum_{\omega \in \Omega} q_\omega^k V_\omega^k - \alpha(q^k) - \beta_k^\top x_1^k$

return the cut $\Theta \geq \alpha_k + \beta_k^\top x_1$

alternative scheme, called a *multi-cut* algorithm, adds up to $|\Omega|$ cuts at each iteration, i.e., one cut for each realization, ω. Thus, at the K^{th} iteration, the multi-cut master problem is:

$$\begin{aligned} \text{Multi-}V_1^K = \min_{x_1, \theta, \Theta} \ & c_1^\top x_1 + \Theta \\ & A_1 x_1 = b_1 \\ & x_1 \geq 0 \\ & \theta_\omega \geq V_\omega^k + {\lambda_\omega^k}^\top (x_1 - x_1^k), \ \forall \omega \in \Omega, k = 1, \ldots, K-1 \\ & \Theta \geq \sum_{\omega \in \Omega} q_\omega^k \theta_\omega - \alpha(q^k), \quad k = 1, \ldots, K-1 \\ & \Theta \geq -M, \end{aligned} \tag{17}$$

where again M is sufficiently large so that the final constraint provides a lower bound on $\mathbb{F}[V_2(\cdot, \omega)]$. The superscript k indicates optimal primal solutions—both x_1^k and q^k—from the k^{th} iteration. As in the single-cut master problem (16) and cut-generation Algorithm 1, in the multi-cut procedure we again have:

$$q^k \in \arg\max_{q \in \mathcal{M}(p)} \left\{ \mathbb{E}_q[V_\omega^k] - \alpha(q) \right\}. \tag{18}$$

The following lemma establishes the relationship between the optimal values of the single-cut and multi-cut master problems, setting up a follow-on result regarding a third algorithmic variant using a conic formulation.

LEMMA 1. *Let $p \in \mathcal{P}$ satisfy $p > 0$. Assume that at the K^{th} iteration of the single-cut and multi-cut algorithms, the respective master problems* (16) *and* (17) *had cuts computed at the same sequence of first-stage decisions x_1^k; i.e., V_ω^k and λ_ω^k, $\omega \in \Omega$, are identical in each algorithm, $k = 1, \ldots, K-1$. Further assume that in each algorithm the corresponding q^k terms, $k = 1, \ldots, K-1$, are identical; i.e., multiple optima in equation* (18) *are resolved consistently. Then,*

$$\textit{Single-}V_1^K \leq \textit{Multi-}V_1^K.$$

Proof of Lemma 1. Let $x_1 \in \mathcal{X}$ denote the constraints of (14). We can then re-write each master problem using an explicit maximization in terms of the second-stage coefficients V_ω^k and λ_ω^k, giving:

$$\text{Single-}V_1^K = \min_{x_1 \in \mathcal{X}} \left[c_1^\top x_1 + \max_{k=1,\ldots,K-1} \left\{ \sum_{\omega \in \Omega} q_\omega^k [V_\omega^k + {\lambda_\omega^k}^\top (x_1 - x_1^k)] - \alpha(q^k) \right\} \right],$$

and

$$\text{Multi-}V_1^K = \min_{x_1 \in \mathcal{X}} \left[c_1^\top x_1 + \max_{k=1,\ldots,K-1} \left\{ \sum_{\omega \in \Omega} q_\omega^k \max_{j=1\ldots,K-1} \left\{ V_\omega^j + {\lambda_\omega^j}^\top (x_1 - x_1^j) \right\} - \alpha(q^k) \right\} \right].$$

Note that the q^k terms are identical in both master problems by hypothesis. The lemma follows immediately since $V_\omega^k + {\lambda_\omega^k}^\top (x_1 - x_1^k) \le \max_{j=1\ldots,K-1} \left\{ V_\omega^j + {\lambda_\omega^j}^\top (x_1 - x_1^j) \right\}$ for all $k = 1, \ldots, K-1$. Q.E.D.

Remark 7. Lemma 1 suggests that the multi-cut formulation will result in a tighter lower bound on model (14)'s optimal value than the single-cut formulation. (This holds when the iterates are identical as in the lemma's hypothesis, but it is possible to construct counterexamples when the iterates differ; see Birge and Louveaux 1988.) However, compared to the single-cut formulation, the multi-cut formulation requires $|\Omega|(K-1)$ more linear constraints and $|\Omega|$ more variables. Therefore, the relative performance of each algorithm is problem- and solver-specific.

It is also possible to combine the single-cut and multi-cut formulations into a hybrid master program:

$$\begin{aligned} V_1^K = \min_{x_1, \theta, \Theta} \; & c_1^\top x_1 + \Theta \\ & A_1 x_1 = b_1 \\ & x_1 \ge 0 \\ & \theta_\omega \ge V_\omega^k + {\lambda_\omega^k}^\top (x_1 - x_1^k),\ \forall \omega \in \Omega, k \in \mathcal{K}_m \\ & \Theta \ge \sum_{\omega \in \Omega} q_\omega^k \theta_\omega - \alpha(q^k), \quad k \in \mathcal{K}_m \\ & \Theta \ge \alpha_k + \beta_k^\top x_1, \quad k \in \mathcal{K}_s, \\ & \Theta \ge -M, \end{aligned} \tag{19}$$

where $\mathcal{K}_s \subseteq \{1, \ldots, K-1\}$ are the iterations at which a single-cut is added, and $\mathcal{K}_m \subseteq \{1, \ldots, K-1\}$ are the iterations at which a multi-cut is added, $\mathcal{K}_s \cap \mathcal{K}_m = \emptyset$. Here, we refer to cuts containing only θ_ω as *multi-cuts*; cuts containing only Θ as *single-cuts*; and the set of inequalities linking Θ and θ as *risk-set cuts*. We might, for example, add single-cuts in early iterations of the algorithm but use multi-cuts as the algorithm nears convergence.

Using the result in Theorem 2 we can, for the entropic risk measure, develop a variant of the multi-cut master problem (17) as follows:

$$\begin{aligned} \text{Conic-}V_1^K = \min_{x_1, \theta, \mu} \; & c_1^\top x_1 + \sum_{\omega \in \Omega} p_\omega \mu_\omega + \mu_0 \\ & A_1 x_1 = b_1 \\ & x_1 \ge 0 \\ & \theta_\omega \ge V_\omega^k + {\lambda^k}^\top (x_1 - x_1^k),\ \forall \omega \in \Omega, k = 1, \ldots, K-1 \\ & \theta_\omega \ge -M \\ & (\tfrac{-1}{\gamma}, \mu_0 - \theta_\omega, \mu_\omega) \in \mathcal{K}_{\exp}^*, \quad \forall \omega \in \Omega. \end{aligned} \tag{20}$$

The following theorem builds on Lemma 1 by relating the optimal value of the conic master problem (20) to its single- and multi-cut counterparts under the entropic risk measure.

THEOREM 4. *Let $\gamma > 0$, let $p \in \mathcal{P}$ satisfy $p > 0$, and let $\mathbb{F}$ denote the entropic risk measure. Assume that at the K^{th} iteration the respective master problems* (16), (17), *and* (20) *had cuts computed at the same sequence of first-stage decisions x_1^k; i.e., V_ω^k and λ_ω^k, $\omega \in \Omega$, are identical in each algorithm, $k = 1, \ldots, K-1$. Further assume that in the single- and multi-cut algorithms the corresponding q^k terms, $k = 1, \ldots, K-1$, are identical. Then, Single-$V_1^K \le$ Multi-$V_1^K \le$ Conic-V_1^K.*

Proof of Theorem 4. By Lemma 1, Single-$V_1^K \le$ Multi-V_1^K. Therefore, it remains to show that Multi-$V_1^K \le$ Conic-V_1^K. This follows using Theorem 2, by noticing that for any fixed values of θ_ω, $\omega \in \Omega$:

$$
\begin{aligned}
&\max_{k=1,\ldots,K-1} \left\{ \sum_{\omega\in\Omega} q_\omega^k \theta_\omega - \frac{1}{\gamma} \sum_{\omega\in\Omega} q_\omega^k \log\left(\frac{q_\omega^k}{p_\omega}\right) \right\} \\
&\le \max_{q\in\mathcal{P}} \left\{ \sum_{\omega\in\Omega} q_\omega \theta_\omega - \frac{1}{\gamma} \sum_{\omega\in\Omega} q_\omega \log\left(\frac{q_\omega}{p_\omega}\right) \right\} \\
&= \min_{\mu\in\mathbb{R}^{|\Omega|+1}} \left\{ \sum_{\omega\in\Omega} p_\omega \mu_\omega + \mu_0 : \, (-\tfrac{1}{\gamma}, \mu_0 - \theta_\omega, \mu_\omega) \in \mathcal{K}_{\text{exp}}^*, \ \forall \omega \in \Omega \right\}.
\end{aligned}
$$

Q.E.D.

REMARK 8. Theorem 4 implies that the conic formulation results in a tighter lower bound on model (14)'s optimal value than either the single- or multi-cut formulations, at least when the iterates are identical. However, compared to the multi-cut formulation, the conic formulation needs $|\Omega|$ exponential dual cone constraints instead of K linear risk-set constraints. The relative performance of each formulation is therefore likely problem- and solver-specific.

REMARK 9. Duque and Morton (2020) derive a similar result to Theorem 4, showing that for a data-driven distributionally robust problem with the Wasserstein distance, the dual of the inner maximization problem (i.e., our Conic$-V_1^K$ problem) forms a tighter lower bound than the multi-cut formulation. It is easy to see how, in the presence of conic strong duality, this result should hold for any convex risk measure.

4.2. Multi-stage stochastic programs

Our results naturally extend to the multistage case, in which Benders' decomposition extends to nested Benders' decomposition (Birge 1985) and stochastic dual dynamic programming (Pereira and Pinto 1991). Because our results closely relate to those for coherent risk measures—the key difference being the penalty term in the intercept of the cuts—and because of the extensive literature on the subject (e.g., Philpott and de Matos 2012, Philpott et al. 2013, Shapiro et al. 2013, Guigues 2016), we will be brief.

We consider a T-stage multistage stochastic program with the problem in the first-stage variables given by:

$$\begin{aligned} V_1 = \min_{x_1} \; & c_1^\top x_1 + \mathbb{F}_{\omega_2 \in \Omega_2}[V_2(x_1, \omega_2)] \\ & A_1 x_1 = b_1 \\ & x_1 \geq 0, \end{aligned}$$

where the second-stage problem is now replaced by a generic t-stage problem for $t = 2, \dots, T$:

$$\begin{aligned} V_t(x_{t-1}, \omega_t) = \min_{\bar{x}_t, x_t} \; & c_t^\top x_t + \mathbb{F}_{\omega_{t+1} \in \Omega_{t+1}}[V_{t+1}(x_t, \omega_{t+1})] \\ & \bar{x}_t = x_{t-1} \quad [\lambda] \\ & A_t x_t + B_t \bar{x}_t = b_t \\ & x_t \geq 0, \end{aligned}$$

and where we assume that $V_{T+1}(\cdot, \cdot) = 0$. Any of the vectors and matrices, A_t, B_t, b_t, and c_t, may depend on $\omega_t \in \Omega_t$, and we make similar assumptions as in the two-stage case with respect to the finite support of Ω, and the existence of feasible and finite optimal solutions for all t-stage problems. While forms of dependency can be introduced (e.g., De Queiroz and Morton 2013, Rebennack 2016, Löhndorf and Shapiro 2019), for simplicity we assume inter-stage independence of the stochastic process. Returning to our discussion in Sections 2.2 and 3 we note how the recursive nature of V_{t+1} in V_t results in a *nested* formulation of the risk measure.

Like the two-stage case, by the monotonicity and convexity of the convex risk measure $\mathbb{F}$, $\mathbb{F}_{\omega_{t+1} \in \Omega_{t+1}}[V_{t+1}(x_t, \omega_{t+1})]$ is convex with respect to x_t. Therefore, we can form single-cut master problems that approximate the stage-wise problems as follows:

$$\begin{aligned} V_1^K = \min_{x_1, \theta_2} \; & c_1^\top x_1 + \theta_2 \\ & A_1 x_1 = b_1 \\ & x_1 \geq 0 \\ & \theta_2 \geq \alpha_{2,k} + \beta_{2,k}^\top x_1, \quad k = 1, \dots, K-1 \\ & \theta_2 \geq -M_2, \end{aligned} \tag{21}$$

and:

$$\begin{aligned} V_t^K(x_{t-1}, \omega_t) = \min_{\bar{x}_t, x_t, \theta_{t+1}} \; & c_t^\top x_t + \theta_{t+1} \\ & \bar{x}_t = x_{t-1} \quad [\lambda] \\ & A_t x_t + B_t \bar{x}_t = b_t \\ & x_t \geq 0 \\ & \theta_{t+1} \geq \alpha_{t+1,k} + \beta_{t+1,k}^\top x_t, \quad k = 1, \dots, K-1 \\ & \theta_{t+1} \geq -M_{t+1}, \end{aligned} \tag{22}$$

where M_{t+1} is sufficiently large so that the final constraint provides a lower bound on $\mathbb{F}[V_{t+1}(\cdot, \omega_{t+1})]$.

Similar to the two-stage case, the cuts are created in an iterative process comprised of two phases. A *forward pass*, which simulates a sequence of state variables, x_t, for $t = 1, \dots, T-1$, and a *backward pass*, which adds a cut to the master problem at each stage t at the point, x_t, in the state-space visited on the forward pass. Simplified pseudo-code is given in Algorithm 2.

Algorithm 2: Stochastic dual dynamic programming algorithm with a convex risk measure.

```
Set K = 1
while not converged do
    // Forward pass
    solve master problem (21) and obtain solution x_1^K
    for t = 2, ..., T − 1 do
        sample ω_t from Ω_t
        solve master problem (22) given (x_{t−1}^K, ω_t) and obtain solution x_t^K
    end
    // Backward pass
    for t = T, ..., 2 do
        for ω_t ∈ Ω_t do
            solve (22) given (x_{t−1}^K, ω_t) to obtain V_t^K(x_{t−1}^K, ω_t) and an extreme point dual
              solution, λ
            set V_{ω_t}^K = V_t^K(x_{t−1}^K, ω_t)
            set λ_{ω_t}^K = λ
        end
        set q^K ∈ arg max_{q∈M(p)} {E_q[V_{ω_t}^K] − α(q)}
        set β_{t,K} = Σ_{ω_t∈Ω_t} q_{ω_t}^K λ_{ω_t}^*
        set α_{t,K} = Σ_{ω_t∈Ω_t} q_{ω_t}^K V_{ω_t}^K − α(q^K) − β_{t,K}^⊤ x_{t−1}^K
        Add the cut θ_t ≥ α_{t,K} + β_{t,K}^⊤ x_t to (22) for t − 1, i.e., updating the model with value
          V_{t−1}^K to V_{t−1}^{K+1}
    end
    K ← K + 1
end
```

Remark 10. If the risk measure $\mathbb{F}$ is *coherent*, then under appropriate technical assumptions, the stochastic dual dynamic programming algorithm has been shown to converge to an ε-optimal solution almost surely in a finite number of iterations (e.g., Guigues 2016). Somewhat surprisingly, to the best of our knowledge, none of the related convergence results in the stochastic dual dynamic programming literature that include coherent risk measures rely on sub-additivity and positive homogeneity directly; instead, they rely on these axioms only to derive convexity, which convex risk measures assume directly. This means that convergence results such as in Guigues (2016) can be directly applied to convex risk measures.

The implication of Remark 10 is that we do not need a new proof of convergence of the SDDP-style Algorithm 2 because the corresponding proofs for coherent risk measures (see, e.g., Guigues 2016) do not use the more restrictive axioms of coherent risk measures, relative to those of convex risk measures.

REMARK 11. Under certain technical assumptions, stochastic dual dynamic programming can be extended to problems involving discrete variables (Zou et al. 2019). The key insight is that instead of computing λ in the backward pass using linear programming duality, we compute λ using the Lagrangian dual formed by relaxing the $\bar{x}_t = x_{t-1}$ constraint. The inclusion of convex risk measures does not change this; i.e., the SDDiP algorithm of Zou et al. can include directly convex risk measures.

4.3. Implementation

We provide a public implementation of the entropic risk measure in `SDDP.jl` (Dowson and Kapelevich 2021), a free and open-source Julia (Bezanson et al. 2017) package for solving multistage stochastic programs. Source-code and documentation for `SDDP.jl` is available at `https://github.com/odow/SDDP.jl`. `SDDP.jl` makes it easy for practitioners to implement their own models and compare the resulting policies under a range of different risk measures. Notably, `SDDP.jl` implements both the single-cut and multi-cut formulations, and allows the user to add cuts of both types to the same model; i.e., we implement a variant of (19). To the best of our knowledge, this hybrid formulation is also a novel, if modest, contribution, although we point to work regarding dynamic aggregation of cuts by Trukhanov et al. (2010).

5. Numerical experiments

We start by describing two simple examples that demonstrate the attractiveness of using the entropic risk measure to control operational risk, before turning to a larger example. The code and data for all of our experiments are available at `https://github.com/odow/SDDP.jl`.

5.1. Example I: Road networks

Consider a road network with three arcs as shown in Figure 2. The travel time on each arc is an independent random variable, denoted X, Y, and Z, and we want to travel from left to right via roads with travel times X and Z, or Y and Z, in order to minimize total travel time, i.e., we will choose between $\mathbb{F}[X+Z]$ and $\mathbb{F}[Y+Z]$, according to the travel time distributions in Table 1. Cominetti and Torrico (2016) use an example with this topology to show that all risk measures except the entropic risk measure violate a definition they call *additive consistency*, which states that $\mathbb{F}[X+Z] \leq \mathbb{F}[Y+Z] \iff \mathbb{F}[X] \leq \mathbb{F}[Y]$.

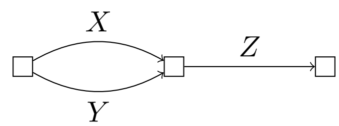

Figure 2 Road network with three roads with random travel times, X, Y, and Z, which are independent.

$\mathbb{P}(\omega)$	Ω_X	Ω_Y	Ω_Z
0.9	1.8	1.7	1.0
0.1	2.0	2.2	2.0

Table 1 Support and corresponding probabilities of the three random variables, X, Y, and Z.

This simple example, with probability mass function given in Table 1, is structured to expose the conditional inconsistency of $\mathbb{CV@R}$. That said, the example's simplicity should not belie the fact that this inconsistency can also arise in large-scale real-world models. In practice, a common approach to generating Table 1 would be to measure the travel times on a large set of representative days, and then use the days as a data-driven empirical distribution with uniform probability.

Our problem can be formulated as a multistage stochastic integer program and solved using `SDDP.jl` or by enumeration. To aid readability, we slightly abuse notation so that $\Omega_{X,Y,Z} = \Omega_X \times \Omega_Y \times \Omega_Z$. We first consider the *end-of-horizon* formulation:

$$V_1 = \min_{x\in\{0,1\}} \underset{\omega\in\Omega_{X,Y,Z}}{\mathbb{F}} [V_2(x,\omega_x,\omega_y,\omega_z)],$$

where $V_2(x,\omega_x,\omega_y,\omega_z) = X(\omega_x)x + Y(\omega_y)(1-x) + Z(\omega_z)$. This formulation is equivalent to choosing between $\mathbb{F}[X+Z]$ and $\mathbb{F}[Y+Z]$.

We also consider the *nested* formulation:

$$V_1 = \min_{x\in\{0,1\}} \underset{\omega\in\Omega_{X,Y}}{\mathbb{F}} [V_2(x,\omega_x,\omega_y)],$$

where:

$$V_2(x,\omega_x,\omega_y) = X(\omega_x)x + Y(\omega_y)(1-x) + \underset{\omega\in\Omega_Z}{\mathbb{F}} [V_3(x,\omega_z)],$$

and $V_3(x,\omega_z) = Z(\omega_z)$. This formulation is equivalent to choosing between $\mathbb{F}[X + \mathbb{F}_{Z|X}[Z]]$ and $\mathbb{F}[Y + \mathbb{F}_{Z|Y}[Z]]$.

The results are shown in Figure 3, in which we plot the first-stage decision x against the risk-aversion parameter γ for the end-of-horizon formulation and the nested formulation. Decision $x = 1$ corresponds to choosing the road with travel time denoted by X, i.e., "road X."

In Figure 3a, we use the $\mathbb{CV@R}$ risk measure. As $\gamma \to 1$, the measure becomes more risk-averse. When the risk-aversion parameter is low, both the end-of-horizon and nested formulations choose $x = 0$ as the first-stage decision, which corresponds to taking road Y. When the risk-aversion parameter takes value $\gamma = 0.7$, both the end-of-horizon model and the nested model switch and choose road X. However, for $\gamma \in (0.865, 0.97)$, the optimal decision to the end-of-horizon model is to switch back and take road Y. This difference in decision making is due to the fact that $\mathbb{CV@R}$ is not conditionally consistent.

In Figure 3b, we use the entropic risk measure. As $\gamma \to \infty$, the measure becomes more risk-averse. Since the end-of-horizon and nested formulations are equivalent, there is only one line visible on the graph (both lines lie on top of each other). When γ is small (i.e., $\gamma < 4.2$), the optimal decision is to take road Y. For larger values of γ the optimal decision is to take road X.

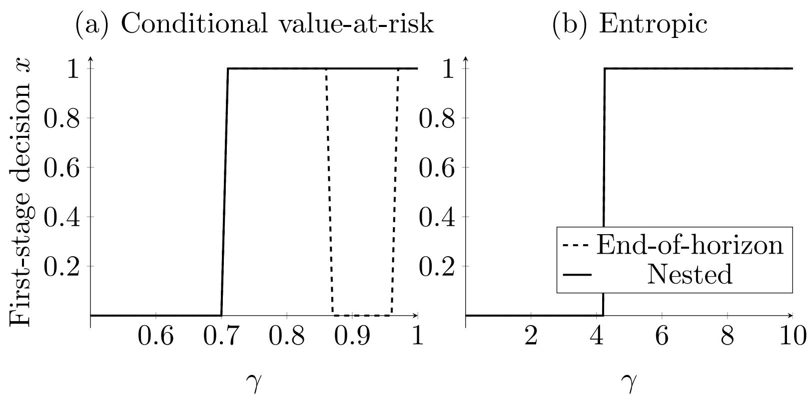

Figure 3 Plot of optimal first-stage decision x against risk aversion parameter γ using the $\mathbb{CV@R}$ (subplot a) and entropic (subplot b) risk measures. Decision $x = 0$ indicates choosing road Y, and $x = 1$ indicates choosing road X.

5.2. Example II: Portfolio management

As a second example, we consider the problem of optimizing a financial portfolio consisting of stocks and bonds. We formulate our problem as a sequential decision model with five stages, $t = 1, 2, \ldots, 5$. There are two state variables in the model: x_t^s, the quantity of stocks held at the end of stage t; and x_t^b, the quantity of bonds held at the end of stage t. For our control variables, we assume that the assets can be rebalanced without transaction costs, and we introduce a consumption variable, u_t, which is the quantity of cash consumed by the agent in stage t.

In each stage t, we represent the market returns by ω_t^s and ω_t^b for the stocks and bonds respectively, and we denote the sample space from which the returns are drawn from by Ω_t. In the first stage, we assume that Ω_t is the singleton $(\omega_1^s, \omega_1^b) = (1, 1)$, and in all other stages that $(\omega_t^s, \omega_t^b) = (1.11, 1.02)$ with probability 0.2, and $(\omega_t^s, \omega_t^b) = (1.04, 1.06)$ with probability 0.8. We assume the returns are independent across the stages. Finally, we assume that the agent initially holds $(x_0^s, x_0^b) =$

$(0,1)$. The goal of the agent is to choose a policy of investment and consumption that maximizes cumulative consumption over the time horizon, accounting for the risk measure.

In the recursive definition of a multistage stochastic program, our model has value functions:

$$\begin{aligned} V_t(x_{t-1},\omega_t) = \min_{u_t,x_t} \ & -u_t + \mathbb{F}_{\omega_{t+1}\in\Omega_{t+1}}[V_{t+1}(x_t,\omega_{t+1})] \\ & x_t^s + x_t^b + u_t = \omega_t^s x_{t-1}^s + \omega_t^b x_{t-1}^b \\ & x_t \geq 0 \\ & u_t \geq 0, \end{aligned}$$

for $t = 1,\ldots,5$, where we assume $V_6(\cdot,\cdot) = 0$. The goal of the agent is to minimize $V_1((0,1),(1,1))$, which amounts to maximizing expected consumption when $\gamma = 0$ and is increasingly averse to low values of consumption as γ grows.

Using `SDDP.jl`, we solved this problem using the entropic and nested $\mathbb{CV@R}$ risk measures over a range of risk aversion parameters, γ (for the entropic, between 0 and 50, and for $\mathbb{CV@R}$, between 0 and 1). Then, we compute the distribution of consumption under each policy. Because returns are positive, and we do not discount, all consumption occurs in the final time period. The results are shown in Figure 4, along with the fraction of stocks chosen in the first stage, i.e., x_1^s.

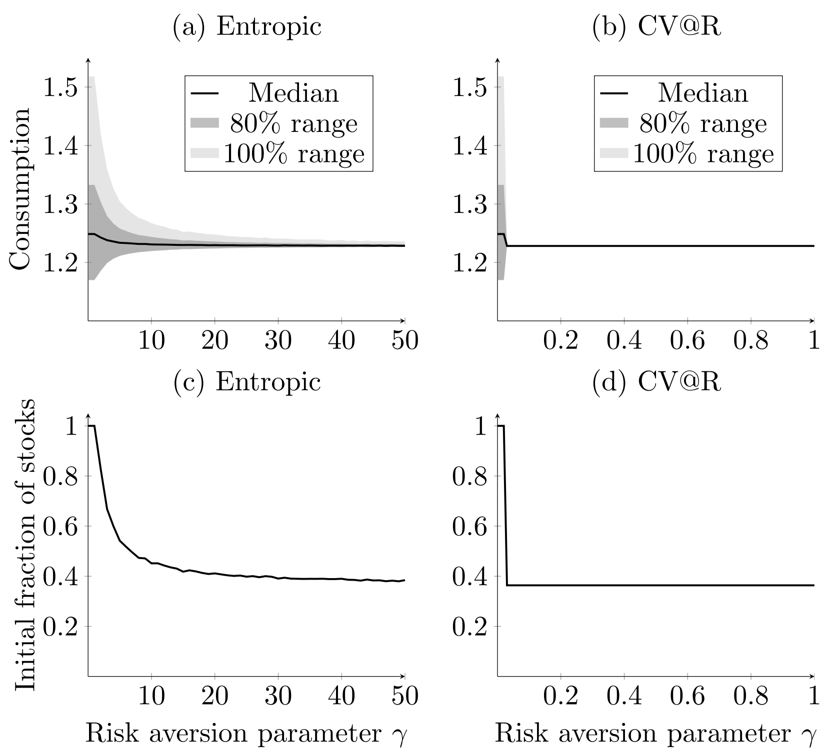

Figure 4 Distribution of consumption and initial fraction of wealth placed in stocks, x_1^s, against the risk aversion parameter γ for the entropic risk measure (a) and (c) and nested $\mathbb{CV@R}$ risk measure (b) and (d).

When the risk aversion parameter is zero, both risk measures seek to maximize expected consumption and hence maximize expected return. Therefore, they both invest all of their wealth in stocks for the highest expected return (23%). In contrast, when the risk aversion parameter is large both risk measures approach the worst-case risk measure. In this situation, both policies invest 4/11 of their wealth in stocks and 7/11 of their wealth in bonds, since this combination yields a return of 5.27% in each stage regardless of the realization of ω. In between these extremes, the two risk measures exhibit large differences. First, the entropic risk measure smoothly grows more risk-averse as γ increases. This is reflected in a reduced mean level of consumption, and also in the spread of the distribution of consumption. Driving this change is the initial investment in stocks, which gradually decreases from 1 towards the worst-case level of 4/11. In contrast, the nested $\mathbb{CV@R}$ risk measure exhibits a classic "bang-bang" solution common in linear programming. When $\gamma \leq 0.022$, the policy is risk-neutral and invests everything in stocks. However, once γ exceeds the threshold, the policy switches to its worst case, and invests 4/11 in stocks. Therefore, small changes in the risk aversion parameter can have large changes in the optimal policy, and nested $\mathbb{CV@R}$ does not have a range of policy options for the agent to choose from.

5.3. Large-scale application: Hydro-thermal scheduling

We consider a hydro-thermal scheduling problem in the Brazilian interconnected power system (Shapiro et al. 2013), using the open-source model and historical data provided by Ding et al. (2019). This problem is based on NEWAVE, an important data-driven model that is deployed and solved on a daily operational basis using nested $\mathbb{CV@R}$ as a risk measure (Maceiral et al. 2018). We now sketch the main details of the model; see Ding et al. (2019), Shapiro et al. (2013) for a full description.

Our hydro-thermal scheduling problem seeks to manage the medium-term electricity generation of the Brazilian national grid. The country is partitioned into four regions, each with an aggregate reservoir, hydro-power station, and thermal power stations. Energy can be transferred between regions. The goal of the agent is to construct a policy for hydro and thermal electricity generation with minimum expected cost under uncertain inflows into the four reservoirs. The cost is composed of two main components: (i) deficit cost, which is incurred if there is insufficient electricity generated to meet demand; and (ii) thermal cost, which is incurred when thermal generation is used. Thermal cost is a convex increasing function of the quantity of thermal generation. As originally formulated, the objective in each stage is composed by the addition of these terms, so the problem can be stated as:

$$
\begin{aligned}
V_t(x_{t-1}, \omega_t) = \min_{x_t} \;& c^d(x_t) + c^t(x_t) + \mathbb{E}[V_{t+1}(x_t, \omega_{t+1})] \\
\text{s.t. } & x_t \in \mathcal{X}(x_{t-1}, \omega_t),
\end{aligned}
$$

where x are the decision variables, c^d is the deficit cost, c^t is the thermal cost, and $\mathcal{X}$ are the various constraints of the problem. In the full problem there are 60 stages—representing five years of 12 monthly periods—four state variables, and $|\Omega_t| = 40$, composed of historical data.

Using `SDDP.jl` we solved this problem using the entropic risk measure for a range of risk aversion parameters γ. We ran the algorithm on each instance for 10,000 iterations, which required 24 million solves of a linear program and appeared to suffice to achieve convergence. The resulting solutions are visualized in Figures 5 and 6.

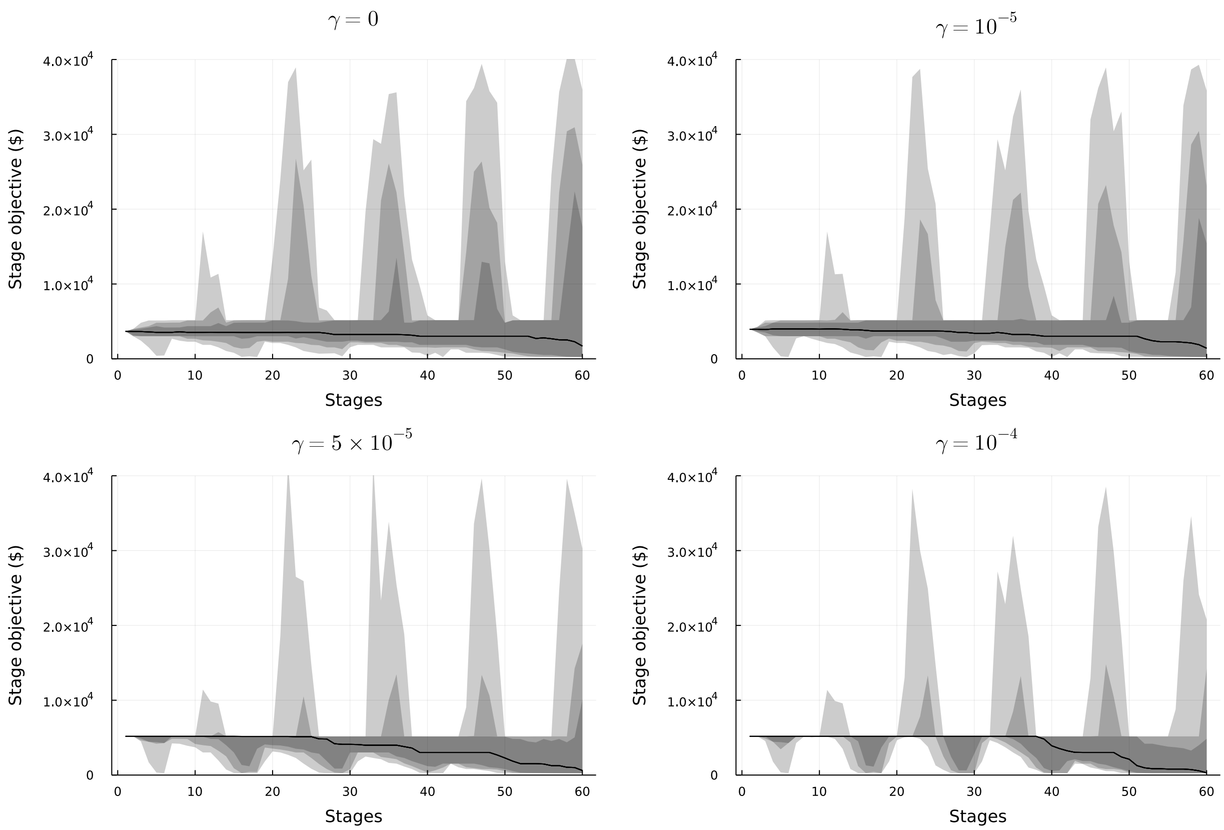

Figure 5 Distribution of stage costs for varying γ. In order of increasing darkness are the 0-100, 1-99, and 5-95 percentiles. The solid line is the median. For each instance, the percentiles of costs are estimated using the same 1,000 sample paths through the scenario tree.

When the risk measure is the expectation ($\gamma = 0$), there are large spike in the stage objective costs on an annual basis. This corresponds to water shortages requiring expensive load-shedding. As the risk-aversion parameter γ is increased, the magnitude of the spikes decreases, and the median cost increases. This corresponds to the system using more thermal generation to save water for future periods of shortage. In Figure 6 we plot the cumulative distribution function of total cost over the five-year period. As the risk-aversion parameter is increased, the bulk of the distribution shifts to the right (more expensive), and the high-cost tail shifts to the left (lower risk). For example, the 99th percentile of total cost is reduced from \$368,000 to \$247,000 (-5.7%) as we move from $\gamma = 0$ to $\gamma = 10^{-4}$, but the mean increases from \$213,000 to \$247,000 (+16%). The limited ability of the

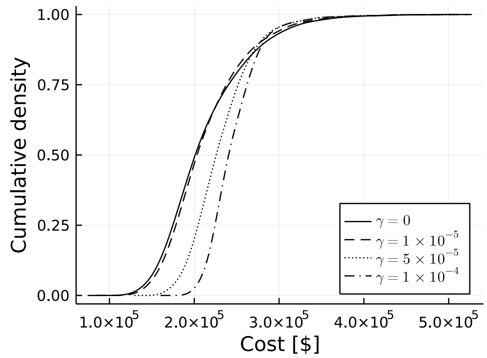

Figure 6 Cumulative distribution function of the end-of-horizon costs with varying γ.

system to reduce tail risk without incurring large growth in the expected cost, suggests that the system is highly constrained. One reason is that our model does not capture all actions available to decision makers, such as the ability to import energy from neighboring countries.

6. Conclusion

This paper has demonstrated the desirability and tractability of using the entropic risk measure to manage risk in multistage stochastic programming, a paradigm well-designed for many operational problems in which similar decisions are made repeatedly. The conditional consistency of the entropic risk measure alleviates the criticism in the literature that coherent risk measures can lead to sub-optimal and counter-intuitive results when viewed from an end-of-horizon perspective. Using conic duality, in the dual form of a convex risk measure, we developed a variant of SDDP, which provides provably tighter approximations of the value function at each stage when compared to classical single- and multi-cut procedures.

However, we note that the entropic risk measure is not a panacea. What we have gained in conditional consistency has come at a cost. Whereas many coherent risk measures have an interpretable meaning (e.g., $\mathbb{CV@R}$ is the expectation of the worst $1-\gamma$ fraction of outcomes), the entropic risk measure does not. There is no clear, general rule for choosing γ, although we know that as γ increases, the measure becomes more risk-averse, and for values close to zero the measure can be approximated by a combination of the expected value plus $\gamma/2$ times the variance (Asienkiewicz and Jaśkiewicz 2017). Future work includes additional computational studies in multistage problems in finance, transportation, and natural resource management to gain a better understanding on the effect that convex risk measures have in practical operations.

Acknowledgments

This research was supported, in part, by Northwestern University's Center for Optimization & Statistical Learning (OSL). We thank Benoît Legat for discussions on conic duality and Roberto Cominetti for valuable comments on an earlier version of this manuscript.

References

Ahmadi-Javid A, Pichler A (2017) An analytical study of norms and Banach spaces induced by the entropic value-at-risk. *Mathematics and Financial Economics* 11(4):527–550.

Artzner P, Delbaen F, Eber JM, Heath D (1999) Coherent measures of risk. *Mathematical Finance* 9(3):203–228.

Asamov T, Ruszczyński A (2015) Time-consistent approximations of risk-averse multistage stochastic optimization problems. *Mathematical Programming* 153(2):459–493.

Asienkiewicz H, Jaśkiewicz A (2017) A note on a new class of recursive utilities in Markov decision processes. *Applicationes Mathematicae* 44:149–161.

Balbás A, Garrido J, Mayoral S (2009) Properties of distortion risk measures. *Methodology and Computing in Applied Probability* 11(3):385.

Baucke R, Downward A, Zakeri G (2018) A deterministic algorithm for solving multistage stochastic minimax dynamic programmes. *Optimization Online* URL `https://www.optimization-online.org/DB_FILE/2018/02/6449.pdf`.

Bayraksan G, Love DK (2015) Data-driven stochastic programming using phi-divergences. Aleman D, Thiele A, Smith J, eds., *The Operations Research Revolution*, 1–19, TutORials in Operations Research (INFORMS).

Ben-Tal A, Teboulle M (2007) An old-new concept of convex risk measures: The optimized certainty equivalent. *Mathematical Finance* 17(3):449–476.

Bezanson J, Edelman A, Karpinski S, Shah VB (2017) Julia: A fresh approach to numerical computing. *SIAM Review* 59(1):65–98.

Birge JR (1985) Decomposition and partitioning methods for multistage stochastic linear programs. *Operations Research* 33(5):989–1007.

Birge JR, Louveaux FV (1988) A multicut algorithm for two-stage stochastic linear programs. *European Journal of Operational Research* 34:384–392.

Boyd S, Vandenberghe L (2004) *Convex Optimization* (Cambridge University Press).

Brandtner M, Kürsten W, Rischau R (2018) Entropic risk measures and their comparative statics in portfolio selection: Coherence vs. convexity. *European Journal of Operational Research* 264(2):707–716.

Bushaj S, Büyüktahtakın İE, Haight RG (2022) Risk-averse multi-stage stochastic optimization for surveillance and operations planning of a forest insect infestation. *European Journal of Operational Research* 299(3):1094–1110.

Chandrasekaran V, Shah P (2016) Relative entropy optimization and its applications. *Mathematical Programming* 161:1–32.

Chen Z, Yang L (2011) Nonlinearly weighted convex risk measure and its application. *Journal of Banking & Finance* 35(7):1777–1793.

Cheridito P, Kupper M (2013) Composition of time-consistent dynamic monetary risk measures in discrete time. *Finance at Fields*, 181–206 (World Scientific).

Cheridito P, Li T (2009) Risk measures on Orlicz hearts. *Mathematical Finance: An International Journal of Mathematics, Statistics and Financial Economics* 19(2):189–214.

Cominetti R, Torrico A (2016) Additive consistency of risk measures and its application to risk-averse routing in networks. *Mathematics of Operations Research* 41(4):1510–1521.

De Queiroz AR, Morton DP (2013) Sharing cuts under aggregated forecasts when decomposing multi-stage stochastic programs. *Operations Research Letters* 41(3):311–316.

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

Dhaene J, Laeven RJ, Vanduffel S, Darkiewicz G, Goovaerts MJ (2008) Can a coherent risk measure be too subadditive? *Journal of Risk and Insurance* 75(2):365–386.

Ding L, Ahmed S, Shapiro A (2019) A Python package for multi-stage stochastic programming. *Optimization Online* http://www.optimization-online.org/DB_HTML/2019/05/7199.html.

Dowson O, Kapelevich L (2021) SDDP.jl: a Julia package for stochastic dual dynamic programming. *INFORMS Journal on Computing* 33(1):27–33.

Duque D, Morton DP (2020) Distributionally robust stochastic dual dynamic programming. *SIAM Journal on Optimization* 30(4):2841–2865.

Föllmer H, Schied A (2002) Convex measures of risk and trading constraints. *Finance and Stochastics* 6(4):429–447.

Föllmer H, Schied A (2010) Convex risk measures. *Encyclopedia of Quantitative Finance* .

Frittelli M, Gianin ER (2005) Law invariant convex risk measures. *Advances in Mathematical Economics*, 33–46 (Springer).

Guastaroba G, Mansini R, Ogryczak W, Speranza MG (2020) Enhanced index tracking with CVaR-based ratio measures. *Annals of Operations Research* 292(2):883–931.

Guigues V (2016) Convergence analysis of sampling-based decomposition methods for risk-averse multistage stochastic convex programs. *SIAM Journal on Optimization* 26(4):2468–2494.

Homem-de-Mello T, Pagnoncelli BK (2016) Risk aversion in multistage stochastic programming: A modeling and algorithmic perspective. *European Journal of Operational Research* 249(1):188–199.

IEEE (2019) 754-2019 - IEEE standard for floating-point arithmetic. Technical report, IEEE.

Jobert A, Rogers LCG (2008) Valuations and dynamic convex risk measures. *Mathematical Finance: An International Journal of Mathematics, Statistics and Financial Economics* 18(1):1–22.

Kovacevic R, Pflug GC (2009) Time consistency and information monotonicity of multiperiod acceptability functionals. *Advanced Financial Modelling* 8:347–370.

Kovacevic RM, Pflug GC (2014) Are time consistent valuations information monotone? *International Journal of Theoretical and Applied Finance* 17(01):1450003.

Kozmík V, Morton DP (2015) Evaluating policies in risk-averse multi-stage stochastic programming. *Mathematical Programming* 152(1-2):275–300.

Kupper M, Schachermayer W (2009) Representation results for law invariant time consistent functions. *Mathematics and Financial Economics* 2(3):189–210.

Löhndorf N, Shapiro A (2019) Modeling time-dependent randomness in stochastic dual dynamic programming. *European Journal of Operational Research* 273(2):650–661.

Maceiral M, Penna D, Diniz A, Pinto R, Melo A, Vasconcellos C, Cruz C (2018) Twenty years of application of stochastic dual dynamic programming in official and agent studies in Brazil: Main features and improvements on the NEWAVE model. *2018 Power Systems Computation Conference (PSCC)*, 1–7.

Mohammadi M, Shahparvari S, Soleimani H (2021) Multi-modal cargo logistics distribution problem: Decomposition of the stochastic risk-averse models. *Computers & Operations Research* 131:105280.

MOSEK ApS (2020) *MOSEK Optimization Suite 9.1.13*. URL `https://docs.mosek.com/9.1/intro/index.html`, date accessed: August 21, 2021.

Nass D, Belousov B, Peters J (2019) Entropic risk measure in policy search. *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 1101–1106 (IEEE).

Pereira M, Pinto L (1991) Multi-stage stochastic optimization applied to energy planning. *Mathematical Programming* 52:359–375.

Pflug GC, Pichler A (2016) Time-inconsistent multistage stochastic programs: Martingale bounds. *European Journal of Operational Research* 249(1):155–163.

Philpott A, de Matos V (2012) Dynamic sampling algorithms for multi-stage stochastic programs with risk aversion. *European Journal of Operational Research* 218(2):470–483.

Philpott A, de Matos V, Finardi E (2013) On solving multistage stochastic programs with coherent risk measures. *Operations Research* 61(4):957–970.

Pichler A, Shapiro A (2021) Mathematical foundations of distributionally robust multistage optimization. *SIAM Journal on Optimization* 31(4):3044–3067.

Rahimian H, Mehrotra S (2019) Distributionally robust optimization: A review. *arXiv preprint arXiv:1908.05659* .

Rebennack S (2016) Combining sampling-based and scenario-based nested Benders decomposition methods: Application to stochastic dual dynamic programming. *Mathematical Programming* 156(1-2):343–389.

Reus L, Pagnoncelli B, Armstrong M (2019) Better management of production incidents in mining using multistage stochastic optimization. *Resources Policy* 63:101404.

Riedel F (2004) Dynamic coherent risk measures. *Stochastic Processes and their Applications* 112(2):185–200.

Rockafellar RT (2007) Coherent approaches to risk in optimization under uncertainty. *OR Tools and Applications: Glimpses of Future Technologies*, 38–61 (INFORMS).

Rockafellar RT, Uryasev SP (2002) Conditional value-at-risk for general loss distributions. *Journal of Banking and Finance* 26:1443–1471.

Russel RH, Behzadian B, Petrik M (2020) Entropic risk constrained soft-robust policy optimization. *arXiv preprint arXiv:2006.11679* .

Ruszczyński A (2010) Risk-averse dynamic programming for Markov decision processes. *Mathematical Programming* 125(2):235–261.

Ruszczyński A, Shapiro A (2006a) Conditional risk mappings. *Mathematics of Operations Research* 31(3):544–561.

Ruszczyński A, Shapiro A (2006b) Optimization of convex risk functions. *Mathematics of Operations Research* 31(3):433–452.

Shapiro A (2021) Tutorial on risk neutral, distributionally robust and risk averse multistage stochastic programming. *European Journal of Operational Research* 288(1):1–13.

Shapiro A, Dentcheva D, Ruszczynski A (2021) *Lectures on Stochastic Programming: Modelling and Theory* (Philadelphia, PA: Society for Industrial and Applied Mathematics).

Shapiro A, Tekaya W, da Costa JP, Soares MP (2013) Risk neutral and risk averse stochastic dual dynamic programming method. *European Journal of Operational Research* 224(2):375–391.

Trukhanov S, Ntaimo L, Schaefer A (2010) Adaptive multicut aggregation for two-stage stochastic linear programs with recourse. *European Journal of Operational Research* 206(2):395–406.

Van Slyke RM, Wets R (1969) L-shaped linear programs with applications to optimal control and stochastic programming. *SIAM Journal on Applied Mathematics* 17(4):638–663.

Vinel A, Krokhmal PA (2017) Certainty equivalent measures of risk. *Annals of Operations Research* 249(1-2):75–95.

Wächter A, Biegler LT (2006) On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. *Mathematical Programming* 106(1):25–57.

Wei L, Hu Y (2014) Coherent and convex risk measures for portfolios with applications. *Statistics & Probability Letters* 90:114–120.

Zou J, Ahmed S, Sun XA (2019) Stochastic dual dynamic integer programming. *Mathematical Programming* 175(1-2):461–502.