Multistage stochastic programs with the entropic risk measure

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

Bernardo K. Pagnoncelli
School of Business, Universidad Adolfo Ibáñez, Santiago, Chile, bernardo.pagnoncelli@uai.cl

Over the last two decades, coherent risk measures have been well studied as a principled, axiomatic way to measure 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 algorithms, a common class of solution methods for multistage stochastic programs. However, even though they facilitate computational tractability, they can suffer from a type of time inconsistency, which we call *conditional inconsistency*. This inconsistency can lead to sub-optimal policies if agents care about their state at the end of the time horizon, but control risk in a stage-wise fashion. The more general class of *convex risk measures* includes the entropic risk measure, which is conditionally consistent and computationally tractable. We explore the relationship between convex risk measures and distributionally robust optimization, and we discuss how to incorporate general convex risk measures into a stochastic dual dynamic programming algorithm.

Key words: risk; entropic; multistage stochastic program; stochastic dual dynamic programming

---

1. Introduction

In order to make decisions under uncertainty, it is convenient to summarize a random variable into a single statistic that can be optimized. Functions which map random variables to real values are called *risk measures*, and over the last two decades *coherent* risk measures have been well studied as a principled and axiomatic way to measure the risk of a random variable (e.g., Artzner et al. 1999, Riedel 2004, Ruszczyński and Shapiro 2006, Shapiro et al. 2009, Ruszczyński 2010, Philpott and de Matos 2012, Philpott et al. 2013, Shapiro et al. 2013, Guigues 2016).

Coherent risk measures have many desirable properties for computation, which have been leveraged so that we can incorporate such measures within stochastic programming algorithms. Miller and Ruszczyński (2011) discuss modeling and algorithmic aspects of static (two-stage) problems for a variety of risk measures, and the monograph of Shapiro et al. (2009) provides a complete treatment of the topic with applications to inventory and portfolio selection problems. In many cases we employ risk aversion to avoid bad outcomes or extreme events, at the sacrifice of minimizing expected costs, or maximizing expected profits.

Dynamic (multistage) problems, which are the focus of our work, are more challenging and present additional difficulties that are absent in 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 2006). The conditional expectations at every stage—which characterize the risk-neutral formulation—are replaced by some coherent risk measure, often involving conditional value-at-risk, and the stochastic dual dynamic programming algorithm (Pereira and Pinto 1991) is applied recursively at each stage. 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 (e.g. Shapiro et al. 2013, Philpott et al. 2013, Kozmík and Morton 2015, Guigues 2016).

In some sequential decision making problems, we wish to avoid bad outcomes at key points in time while ignoring the trajectory of outcomes along the way. For example, we care primarily about the value of a retirement 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 an 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. If these two formulations lead to different risk preferences then we say that the risk measure is *conditionally inconsistent*, a notion made precise in Section 3. Translating into the optimization context, we say that a risk measure is conditionally consistent if and only if any policy P that is preferred over a policy Q under a nested formulation respects that preference order in an end-of-horizon formulation.

It is well-known that, in general, coherent risk measures are not conditionally consistent, although two exceptions are the risk-neutral expectation and the worst-case risk measure. This lack of conditional consistency 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) is 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 highlight the work of Asamov and Ruszczyński (2015), who study bounds and approximations related to a property we establish in Theorem 1.

Unlike Pflug and Pichler (2016), we show that a conditionally consistent risk measure can be obtained by relaxing two key axioms of coherent risk measures to obtain the class of *convex* risk measures. In particular, we recast existing results to show that a popular convex risk measure, the entropic risk measure, is conditionally consistent.

Conditional consistency is particularly desirable for multistage stochastic programming because it allows us to find optimal policies to previously intractable end-of-horizon formulations by using a conditionally consistent risk measure in a nested formulation. In practice, measures that are conditionally consistent exhibit a cohesion in terms of preferences between their nested and end-of-horizon versions. Therefore, motivated by the goal of incorporating the entropic risk measure in stochastic dual dynamic programming algorithms, we extend this algorithmic framework to include any convex risk measure.

To summarize, the main contributions of this paper are:

i) to define *conditional consistency*, and to interpret existing results in a new light showing that the entropic risk measure is conditionally consistent; and

ii) to provide a practical way of incorporating convex risk measures in stochastic dual dynamic programming, allowing us to solve multistage stochastic programming problems with the entropic risk measure.

The rest of this paper is laid out as follows. In Section 2, we briefly review risk measures. In Section 3, we introduce the notion of *conditional consistency* and explain why this is a desirable property when making decisions under uncertainty. In Section 4, we discuss the dual representation of convex risk measures and their relationship to distributionally robust optimization. In Section 5, we explain how to incorporate convex risk measures in stochastic dual dynamic programming as a solution method for risk-averse multistage stochastic programs. Section 6 provides two examples.

2. Risk measures

Consider a probability space $(\Omega, \mathcal{F}, P)$, and 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 $\mathcal{F}_T = \mathcal{F}$. Let $\mathcal{Z}_t$ denote a space of $\mathcal{F}_t$-measurable functions from Ω to $\mathbb{R}$, and $\mathcal{Z} := \mathcal{Z}_1 \times \cdots \times \mathcal{Z}_T$.

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. (2007):

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 $a \geq 0$.

A commonly used coherent risk measure is the *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 $\gamma \in [0, 1)$. $\mathbb{CV@R}$ 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:

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

where $\gamma > 0$. 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. Problems with subadditivity are the subject of the 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.

3. Conditional consistency

The previous section discusses *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 random variables is induced by a policy that we select based on exogenous randomness that evolves over time. We represent this process as $Z = \{Z_t\}_{t=1}^T$, where $Z_t \in \mathcal{Z}_t$.

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 2006, Shapiro et al. 2009, Ruszczyński 2010, Shapiro et al. 2013):

$$\text{Nested-Risk}(Z) = \mathbb{F}\left[Z_1 + \mathbb{F}[Z_2 + \mathbb{F}[\cdots + \mathbb{F}[Z_T]\cdots] \mid Z_2] \mid Z_1]\right].$$

Here the inner evaluations of risk measure are conditioned on the realizations of random variables in the outer layers of nesting. 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 requisite expectations are finite. A risk measure $\mathbb{F}$ is said to be *conditionally consistent* if:

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

In other words, conditional consistency says that if we prefer the sequence X_1, X_2 to that of Y_1, Y_2 when viewed from an end-of-horizon perspective, then we should also prefer X_1, X_2 from a nested perspective.

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). To demonstrate, consider the scenario trees in Figure 1, which we have borrowed and adapted slightly from Pflug and Pichler (2016). The numbers inside each node of the tree represent the cost incurred by the agent at that node, and the numbers on the arcs are the corresponding transition probabilities.

First, consider the scenario tree in Figure 1(a). In the upper branch from the second stage, there are two possible future outcomes: 5 with probability 0.1 and 4 with probability 0.9. Therefore, the $\mathbb{CV@R}$ at the worst 10% quantile ($\mathbb{CV@R}_{1-0.1}$) is 5. So, $2 + \mathbb{F}[Z_2|Z_1 = 2] = 2 + 5$. In the analogous lower branch, the $\mathbb{CV@R}$ is 3, so we have $0 + 3$. Stepping back to the first stage, the nested risk measure sees a cost-to-go of $2 + 5 = 7$ with probability 0.1 and $0 + 3 = 3$ with probability 0.9. Therefore, the nested $\mathbb{CV@R}$ of the tree is 7. However, if we calculate the $\mathbb{CV@R}$ over the end-of-horizon outcomes, i.e., based on the cumulative costs at the leaf nodes in Figure 1(a), we find that the end-of-horizon $\mathbb{CV@R}$ is 6.1.

Now, consider the scenario tree in Figure 1(b). From the upper branch in the second stage, there are two possible future outcomes: 6 with probability 0.1 and 3 with probability 0.9, giving

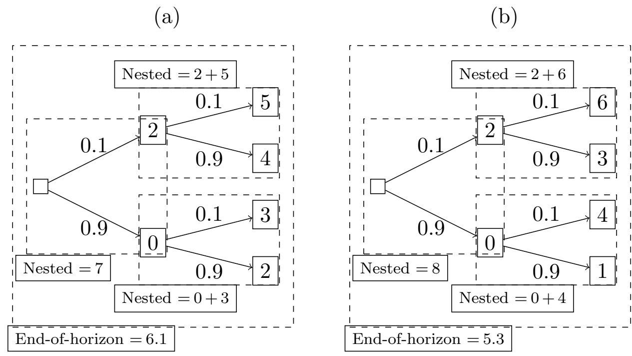

Figure 1 Different views of risk over two scenario trees (a) and (b) using the $\mathbb{CV@R}_{0.9}$ risk measure. Values in each square node indicate the cost incurred by the agent in that node. Values on the arcs are transition probabilities.

$2+\mathbb{F}[Z_2|Z_1=2]=2+6=8$. In the lower branch, the $\mathbb{CV@R}$ is 4. Stepping back to the first stage, the nested risk measure sees a cost-to-go of 8 with probability 0.1 and 4 with probability 0.9. Therefore, the nested $\mathbb{CV@R}$ of the tree is 8. However, if we calculate the $\mathbb{CV@R}$ over the end-of-horizon outcomes, we find that the end-of-horizon $\mathbb{CV@R}$ is 5.3.

Suppose that we are taking some actions in stage one, which determine the costs as given by scenario tree (a) versus (b). From a nested viewpoint, we prefer scenario tree (a) over scenario tree (b). However, from an end-of-horizon viewpoint, we prefer scenario tree (b) over scenario tree (a). Thus, if we use the nested formulation of risk to guide our decision-making in stage one, we may make sub-optimal decisions from an end-of-horizon point-of-view.

This example demonstrates that $\mathbb{CV@R}$ is not conditionally consistent. However, there are two coherent risk measures that do satisfy this property, and they are expectation and worst-case. More specifically, it is straightforward to show that for the expectation and worst-case risk measures:

$$\mathbb{F}[Z_1+Z_2+\cdots+Z_T]=\mathbb{F}[Z_1+\mathbb{F}[Z_2+\mathbb{F}[\cdots+\mathbb{F}[Z_T]\cdots]\mid Z_2]\mid Z_1]].$$

Moreover, it is also straightforward to show that, in general, this equality 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. However, the following theorem suggests that it may be useful to focus on the entropic risk measure as a family of risk measures that have expectation and worst-case measures as the limiting cases and provides conditional consistency.

Theorem 1. *Let* $\mathbb{ENT}_\gamma[Z]=\frac{1}{\gamma}\log\left(\mathbb{E}\left[e^{\gamma Z}\right]\right)$*, and let* (X,Y) *be such that* $\mathbb{E}\left[e^{\gamma(X+Y)}\right]<\infty$*, where* $\gamma>0$*. Then,*

$$\mathbb{ENT}_\gamma[X+Y]=\mathbb{ENT}_\gamma[X+\mathbb{ENT}_\gamma[Y\mid X]].$$

Proof of Theorem 1.

$$\begin{aligned}
\mathbb{ENT}_\gamma[X+Y] =& \frac{1}{\gamma}\log\left(\mathbb{E}\left[e^{\gamma(X+Y)}\right]\right)\\
=& \frac{1}{\gamma}\log\left(\mathbb{E}_X\left[\mathbb{E}_{Y|X}\left[e^{\gamma(X+Y)}\right]\right]\right)\\
=& \frac{1}{\gamma}\log\left(\mathbb{E}\left[e^{\gamma X}\,\mathbb{E}\left[e^{\gamma Y} \mid X\right]\right]\right)\\
=& \frac{1}{\gamma}\log\left(\mathbb{E}\left[e^{\gamma X}e^{\gamma\frac{1}{\gamma}\log\left(\mathbb{E}\left[e^{\gamma Y} \mid X\right]\right)}\right]\right)\\
=& \frac{1}{\gamma}\log\left(\mathbb{E}\left[e^{\gamma\left(X+\frac{1}{\gamma}\log\left(\mathbb{E}\left[e^{\gamma Y} \mid X\right]\right)\right)}\right]\right)\\
=& \mathbb{ENT}_\gamma[X+\mathbb{ENT}_\gamma[Y \mid X]]. \qquad \text{Q.E.D.}
\end{aligned}$$

COROLLARY 1. *If the requisite moment generating functions exist then the entropic risk measure is conditionally consistent.*

Proof of Corollary 1. Both directions of the implication in Definition 2 are immediately satisfied when $\mathbb{ENT}_\gamma[X_1+\mathbb{ENT}_\gamma[X_2 \mid X_1]] = \mathbb{ENT}_\gamma[X_1+X_2]$ and $\mathbb{ENT}_\gamma[Y_1+\mathbb{ENT}_\gamma[Y_2 \mid Y_1]] = \mathbb{ENT}_\gamma[Y_1+Y_2]$. Q.E.D.

We state and prove Theorem 1 to give a simple setting under which it is possible to show that the end-of-horizon risk is equal to the nested risk. A more general result, due to Kupper and Schachermayer (2009), shows that the *only* class of risk measures for which the end-of-horizon view is equal to the nested view is the entropic risk measure. We restate their theorem for completeness without proof (our emphasis added).

THEOREM 2. *(Kupper and Schachermayer 2009, Theorem 1.10) Let $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t\in\mathbb{N}_0}, \mathbb{P})$ be a standard filtered probability space. The family $(\mathbb{F}_t)_{t\in\mathbb{N}_0}$ is a law invariant, **time consistent**, relevant dynamic risk measure **if and only if** there is $\gamma \in (-\infty, \infty]$ such that:*

$$\mathbb{F}_t[Z] = \frac{1}{\gamma}\log\mathbb{E}[e^{\gamma X} \mid \mathcal{F}_t] \textit{ for all } t \in \mathbb{N}_0.$$

The limiting cases are defined as $\mathbb{F}_t[Z] = \mathbb{E}[Z \mid \mathcal{F}_t]$ when $\gamma = 0$, and $\mathbb{F}_t[Z] = \operatorname{ess\,sup}[Z \mid \mathcal{F}_t]$ when $\gamma = \infty$.

Note that there are multiple notions of time consistency in the literature. In this case, Kupper and Schachermayer (2009) use *time consistency* in the same sense as Artzner et al. (2007), in that $\mathbb{F}[X+Y] = \mathbb{F}[X+\mathbb{F}[Y \mid X]]$. Definition 2 instead only requires consistent risk preferences.

To conclude this section, we repeat the analysis of the scenario trees in Figure 1, this time using the entropic risk measure with $\gamma = 1$. The adjusted scenario trees are given in Figure 2. Using the entropic risk measure, we now prefer scenario tree (b) over scenario tree (a) regardless of whether

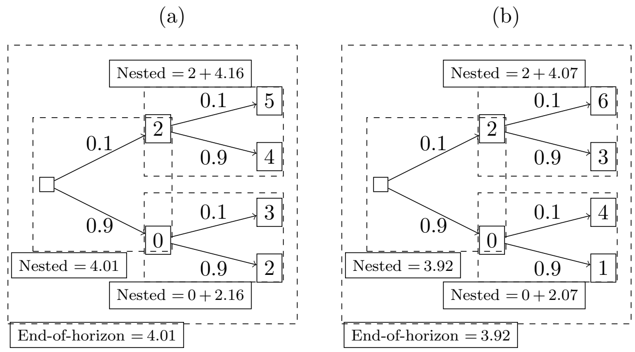

Figure 2 Different views of risk over two scenario trees (a) and (b) using the entropic risk measure with $\gamma = 1$. Values in each square node indicate the cost incurred by the agent in that node. Values on the arcs are transition probabilities.

we calculate the end-of-horizon risk or the nested risk. Moreover, as shown in Theorem 1, the end-of-horizon risk is equal to the nested risk of the root node.

In the remainder of this paper, we consider the dual representation of a risk measure with the view to constructing computationally tractable algorithms. In the context of a stochastic dual dynamic programming algorithm—as well as most other approaches rooted in sample average approximation or other discretization schemes—computational tractability hinges on employing random variables at each stage with finite support, and so to simplify our analysis, we restrict to random variables with a finite sample space, Ω, and respective positive probabilities, p_ω, $\omega \in \Omega$. We denote a random variable with the uppercase Z and realizations, typically costs, by $Z(\omega) = z_\omega$.

4. Representation and visualization of convex risk measures

In this section we discuss alternative representations for convex risk measures. Some representations are useful for our numerical computations, while others establish connections with distinct paradigms of decision making under uncertainty. We start by considering the dual representation of a risk measure, which in the case of the entropic risk measure has a closed-form solution. 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.

4.1. The dual form

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:

$$\mathbb{F}[Z] = \sup_{q \in \mathcal{M}(p)} \mathbb{E}_q[Z]. \tag{3}$$

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:

$$\mathcal{M}(p) = \left\{ q \in \mathcal{P} \;\middle|\; q_\omega \leq \frac{p_\omega}{1-\gamma}, \quad \omega \in \Omega \right\}. \tag{4}$$

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

$$\mathbb{F}[Z] = \sup_{q \in \mathcal{M}(p)} \left\{ \mathbb{E}_q[Z] - \alpha(q) \right\}. \tag{5}$$

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 must be 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, $\mathcal{M}(p) = \mathcal{P}$, 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:

$$\begin{aligned} \mathbb{ENT}_\gamma[Z] = \max_q \; & \left[\sum_{\omega \in \Omega} z_\omega q_\omega - \frac{1}{\gamma} \sum_{\omega \in \Omega} q_\omega \log\left(\frac{q_\omega}{p_\omega}\right) \right] \\ \text{s.t.} \; & \sum_{\omega \in \Omega} q_\omega = 1 \\ & q_\omega \geq 0, \quad \omega \in \Omega. \end{aligned} \tag{6}$$

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

Theorem 3. *Let $\gamma > 0$ and $p \in \mathcal{P}$ satisfy $p > 0$. Then the unique optimal solution to problem* (6) *is given by:*

$$q^*_\omega = \frac{p_\omega e^{\gamma z_\omega}}{\sum_{\omega \in \Omega} p_\omega e^{\gamma z_\omega}}.$$

Proof of Theorem 3. The unique optimal solution is obtained on the convex and compact feasible region of problem (6) because the objective function is strictly concave. Neglecting the non-negativity constraints, forming the Lagrangian, differentiating and solving for q^*_ω yields $q^*_\omega \propto p_\omega e^{\gamma z_\omega}$. The fact that the probability mass function must sum to one yields the proportionality constant in the theorem's statement and satisfies non-negativity. Q.E.D.

Remark 1. The closed-form expression in Theorem 3 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 (6) 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 2. 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.

As we show in Theorem 3, the dual form of a risk measure offers an alternative representation, which sometimes is more tractable than its original primal form. In addition, the dual form is useful because the lack of conditional consistency in coherent risk measures can be visualized by plotting the probability mass function q that attains the supremum in the dual representation (3) with its risk set. Next we illustrate a simple example involving $\mathbb{CV@R}$ and risk set (4).

Consider two independent random variables, Z_1 and Z_2, that are uniformly distributed 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 3, 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.

The results in Figure 3 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, e.g., 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. It is this difference in monotonicity that gives rise to the conditional inconsistency in Figure 1. 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.

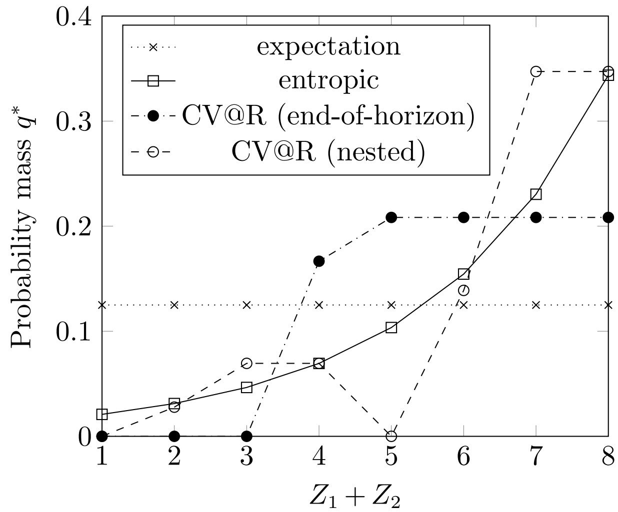

Figure 3 Probability mass q that achieves the supremum of the dual representation for four risk measures.

4.2. Distributionally robust interpretation of a convex risk measure

Our main motivation for using the entropic risk measure is its conditional consistency property in multistage problems. As we now discuss, the entropic risk measure can be interpreted as a Lagrangian relaxation of a distributionally robust risk measure using the Kullback-Leibler phi-divergence. In general, a distributionally robust risk measure with phi-divergence can be formulated as

$$\mathbb{F}[Z] = \max_{q \in \mathcal{M}(p)} \mathbb{E}_q[Z], \tag{7}$$

where

$$\mathcal{M}(p) = \left\{ q \in \mathcal{P} \mid \partial^\phi(q,p) \leq \varepsilon \right\},$$

and where ∂^ϕ is a function specific to the chosen phi-divergence. The Kullback-Leibler divergence function is given by:

$$\partial^\phi(q,p) = \sum_{\omega \in \Omega} q_\omega \log\left(\frac{q_\omega}{p_\omega}\right). \tag{8}$$

For an extensive discussion on distributionally robust risk measures via phi-divergence, see Bayraksan and Love (2015).

Formulation (7), in which the Kullback-Leibler divergence defined in (8) is bounded by ε is a coherent risk measure called the *entropic value-at-risk* (Ahmadi-Javid 2012). However, since the entropic risk measure is not coherent, it does not correspond to this formulation directly, but rather it can be recovered from a reformulation using a Lagrangian relaxation of the bounding constraint:

$$\mathcal{L}(\mu) = \max_{q \in \mathcal{P}} \left[\mathbb{E}_q[Z] - \mu \left(\partial^\phi(q,p) - \varepsilon \right) \right], \tag{9}$$

where μ is the non-negative Lagrangian multiplier. When $\partial^\phi(q,p)$ is given by equation (8) this is equivalent to the entropic risk measure—see formulation (6)—yielding the same optimal q^* when $\mu = \frac{1}{\gamma}$ for $\gamma > 0$ and the same value when $\varepsilon = 0$. Note that the phi-divergence function of equation (8) is zero if and only if $q = p$, which explains why the entropic risk measure is equivalent to expectation when $\gamma \to 0$, i.e., $\mu \to \infty$. When $\varepsilon = 0$, the original distributionally robust risk measure is equivalent to the expectation operator. More generally, because $\partial^\phi(\cdot,\cdot) \geq 0$, a relaxed risk measure defined via equation (9) can be interpreted as a risk measure in which we penalize deviations from expectation, where deviations are measured according to the corresponding phi-divergence function.

If $\partial^\phi(q,p)$ is given by equation (8), $\varepsilon \geq 0$, and $\mu = \frac{1}{\gamma}$ then equation (9) is equivalent to the entropic risk measure plus an additive constant ε/γ. By the translation equivariance property, conditional consistency is maintained.

THEOREM 4. *Let* $\mathbb{ENT}_\gamma[Z]$ *be defined by* (6). *The risk measure* $\mathbb{F}_{\varepsilon,\gamma}[Z] = \mathbb{ENT}_\gamma[Z] + \frac{\varepsilon}{\gamma}$, *where* $\varepsilon \geq 0$ *and* $\gamma > 0$, *is conditionally consistent.*

Proof of Theorem 4. Consider (X_1, X_2) and (Y_1, Y_2). From the statement of the theorem, we have:

$$\begin{aligned}\mathbb{F}_{\varepsilon,\gamma}[X_1 + \mathbb{F}_{\varepsilon,\gamma}[X_2|X_1]] =& \mathbb{ENT}_\gamma\left[X_1 + \mathbb{ENT}_\gamma[X_2|X_1] + \frac{\varepsilon}{\gamma}\right] + \frac{\varepsilon}{\gamma}\\ =& \mathbb{ENT}_\gamma[X_1 + \mathbb{ENT}_\gamma[X_2|X_1]] + 2\frac{\varepsilon}{\gamma}\\ =& \mathbb{ENT}_\gamma[X_1 + X_2] + 2\frac{\varepsilon}{\gamma}\\ =& \mathbb{F}_{\varepsilon,\gamma}[X_1 + X_2] + \frac{\varepsilon}{\gamma}.\end{aligned}$$

Therefore:

$$\begin{aligned}\mathbb{F}_{\varepsilon,\gamma}[X_1 + \mathbb{F}_{\varepsilon,\gamma}[X_2|X_1]] \leq \mathbb{F}_{\varepsilon,\gamma}[Y_1 + \mathbb{F}_{\varepsilon,\gamma}[Y_2|Y_1]] &\iff \mathbb{F}_{\varepsilon,\gamma}[X_1 + X_2] + \frac{\varepsilon}{\gamma} \leq \mathbb{F}_{\varepsilon,\gamma}[Y_1 + Y_2] + \frac{\varepsilon}{\gamma}\\ &\iff \mathbb{F}_{\varepsilon,\gamma}[X_1 + X_2] \leq \mathbb{F}_{\varepsilon,\gamma}[Y_1 + Y_2].\end{aligned}$$

Q.E.D.

4.3. Conic dual of entropic risk measure

In Section 4.1 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 distributionally robust optimization and provide us with a tighter formulation for use in Section 5. While we again focus on the entropic risk measure, the same process could be repeated for other convex risk measures.

THEOREM 5. *Let $p \in \mathcal{P}$ satisfy $p > 0$. The conic dual of problem* (6) *(or, equivalently, problem* (9) *with $\varepsilon = 0$) 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 5. 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 (6) in conic form as:

$$\begin{aligned} \mathbb{ENT}_\gamma[Z] = \max_{q,t} & \left[\sum_{\omega \in \Omega} z_\omega q_\omega - \frac{1}{\gamma} \sum_{\omega \in \Omega} t_\omega\right] \\ \text{s.t. } & 1 - \sum_{\omega \in \Omega} q_\omega \in \{0\} \quad [\mu_0] \\ & (-t_\omega, q_\omega, p_\omega) \in \mathcal{K}_{\exp}, \quad [\mu_\omega] \quad \omega \in \Omega. \end{aligned} \tag{10}$$

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:

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

has a dual in minimization form:

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

where $\mathcal{K}_i^*$ is the dual cone of $\mathcal{K}_i$. Therefore, taking the dual of (10), 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:

$$\begin{aligned} \mathbb{ENT}_\gamma[Z] = \min_{\mu} & \sum_{\omega \in \Omega} p_\omega \mu_\omega + \mu_0 \\ \text{s.t. } & (-\tfrac{1}{\gamma}, \mu_0 - z_\omega, \mu_\omega) \in \mathcal{K}^*_{\exp}, \quad \omega \in \Omega. \end{aligned} \tag{11}$$

Finally, note that conic strong duality holds because (11) has a strictly feasible solution and (10) has a finite optimal solution, and so the optimal objective value of (10) is equal to that of (11). Q.E.D.

Remark 3. The dual problem can be efficiently solved as a conic program using off-the-shelf commercial solvers such as Mosek (MOSEK ApS 2019), 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 5 because it arises most naturally from conic duality, and in our anecdotal experience, the $\mathcal{K}^*_{\exp}$ formulation solves most quickly.

Remark 4. 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:

$$\mathbb{ENT}\,\mathrm{V@R}_\gamma[Z] = \min_{\mu_0} \left[\mu_0 + \frac{1}{\gamma}\varepsilon + \frac{1}{\gamma} \sum_{\omega\in\Omega} p_\omega \left(e^{\gamma(z_\omega - \mu_0)} - e^0 \right) \right]. \tag{12}$$

The view in Theorem 5 can be re-arranged to obtain:

$$\mathbb{ENT}_\gamma[Z] = \min_{\mu_0} \left[\mu_0 + \frac{1}{\gamma} \sum_{\omega\in\Omega} p_\omega \frac{e^{\gamma(z_\omega - \mu_0)}}{e^1} \right]. \tag{13}$$

These formulations differ in two terms: (i) equation (12) includes the additive constant ε/γ discussed in Section 4.2; and (ii) equation (12) subtracts e^0 (the exponentiated additive identity), whereas in (13), we divide by e^1 (the exponentiated multiplicative identity).

Equation (13) 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:

$$\mathbb{ENT}_\gamma[Z] = \inf_{\zeta\in\mathbb{R}} \left\{ \zeta + \frac{1}{\gamma e^1} \mathbb{E}\left[e^{\gamma(Z-\zeta)} \right] \right\}. \tag{14}$$

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

5. 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 (Benders 1962, Van Slyke and Wets 1969), and solved using a multi-stage variant known as a stochastic dual dynamic programming (Pereira and Pinto 1991).

5.1. Two-stage stochastic programs

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

$$\begin{aligned} V_1 = \min_{x_1} \; & c_1^\top x_1 + \mathbb{F}[V_2(x_1,\omega)] \\ & A_1 x_1 = b_1 \\ & x_1 \geq 0, \end{aligned} \tag{15}$$

where the second-stage problem is:

$$\begin{aligned} V_2(x_1,\omega) = \min_{\bar{x},x_2} \; & c_2^\top x_2 \\ & \bar{x} = x_1 \quad [\lambda] \\ & A_2 x_2 + B_2 \bar{x} = b_2 \\ & x_2 \geq 0. \end{aligned} \tag{16}$$

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 vector of random parameters has finite support. 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 (15) 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:

$$\begin{aligned} \text{Single-}V_1^K = \min_{x_1,\Theta} \; & c_1^\top x_1 + \Theta \\ & A_1 x_1 = b_1 \\ & x_1 \geq 0 \\ & \Theta \geq \alpha_k + \beta_k^\top x_1, \quad k = 1,\ldots,K-1 \\ & \Theta \geq -M, \end{aligned} \tag{17}$$

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 (17) is solved to obtain a feasible x_1. Then, the second-stage problems (16) are solved for each $\omega \in \Omega$. In the backward step, 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. (2009, Theorem 6.11).

THEOREM 6. *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 (5). 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 6. By equation (5), $q^* \in \mathcal{M}(p)$, and 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 6, we obtain Algorithm 1 to generate Benders' cuts for the master problem (17), which approximates model (15). 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 3.

The master problem (17) and cut-generation procedure just outlined form a *single-cut* algorithm because one cut is added at each Benders' iteration to approximate $\mathbb{F}[V_2(x_1,\omega)]$ at $x_1 = x_1^k$. An 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:

$$\begin{aligned}
\text{Multi-}V_1^K = \min_{x_1,\theta,\Theta}\ & c_1^\top x_1 + \Theta \\
& A_1 x_1 = b_1 \\
& x_1 \geq 0 \\
& \theta_\omega \geq V_\omega^k + {\lambda_\omega^k}^\top (x_1 - x_1^k),\ \forall \omega \in \Omega, k = 1,\dots,K-1 \\
& \Theta \geq \sum_{\omega\in\Omega} q_\omega^k \theta_\omega - \alpha(q^k), \quad k = 1,\dots,K-1 \\
& \Theta \geq -M,
\end{aligned} \tag{18}$$

where 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

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

Given x_1^k at iteration k

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

  solve subproblem (16) 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 $\Theta \geq \alpha_k + \beta_k^\top x_1$

the single-cut master problem (17) and cut-generation Algorithm 1, in the multi-cut procedure we again have:

$$q^k \in \underset{q \in \mathcal{M}(p)}{\arg\max} \left\{ \mathbb{E}_q[V_\omega^k] - \alpha(q) \right\}. \tag{19}$$

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* (17) *and* (18) *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* (19) *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 (15). 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 $V_\omega^k + {\lambda_\omega^k}^\top (x_1 - x_1^k) \leq \max_{j=1\dots,K-1} \left\{ V_\omega^j + {\lambda_\omega^j}^\top (x_1 - x_1^j) \right\}$ for all $k = 1, \dots, K-1$. Q.E.D.

Remark 5. Lemma 1 suggests that the multi-cut formulation will result in a tighter lower bound on model (15)'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:

$$
\begin{aligned}
V_1^K = \min_{x_1,\theta,\Theta} \ & c_1^\top x_1 + \Theta \\
& A_1 x_1 = b_1 \\
& x_1 \geq 0 \\
& \theta_\omega \geq V_\omega^k + {\lambda_\omega^k}^\top (x_1 - x_1^k), \ \forall \omega \in \Omega, k \in \mathcal{K}_m \\
& \Theta \geq \sum_{\omega \in \Omega} q_\omega^k \theta_\omega - \alpha(q^k), \quad k \in \mathcal{K}_m \\
& \Theta \geq \alpha_k + \beta_k^\top x_1, \quad k \in \mathcal{K}_s, \\
& \Theta \geq -M,
\end{aligned}
\tag{20}
$$

where $\mathcal{K}_s \subseteq \{1, \dots, K-1\}$ are the iterations at which a single-cut is added, and $\mathcal{K}_m \subseteq \{1, \dots, K-1\}$ 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 5 we can, for the entropic risk measure, develop a variant of the multi-cut master problem (18) as follows:

$$
\begin{aligned}
\text{Conic-}V_1^K = \min_{x_1,\theta,\mu} \ & c_1^\top x_1 + \sum_{\omega \in \Omega} p_\omega \mu_\omega + \mu_0 \\
& A_1 x_1 = b_1 \\
& x_1 \geq 0 \\
& \theta_\omega \geq V_\omega^k + {\lambda^k}^\top (x_1 - x_1^k), \ \forall \omega \in \Omega, k = 1, \dots, K-1 \\
& \theta_\omega \geq -M \\
& (\tfrac{-1}{\gamma}, \mu_0 - \theta_\omega, \mu_\omega) \in \mathcal{K}_{\text{exp}}^*, \quad \forall \omega \in \Omega.
\end{aligned}
\tag{21}
$$

The following theorem builds on Lemma 1 by relating the optimal value of the conic master problem (21) to its single- and multi-cut counterparts under the entropic risk measure.

Theorem 7. *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* (17), (18), *and* (21) *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 7. 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 5, by noticing that for any fixed values of θ_ω, $\omega \in \Omega$:

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

Q.E.D.

Remark 6. Theorem 7 implies that the conic formulation results in a tighter lower bound on model (15)'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 7. Duque and Morton (2019) derive a similar result to Theorem 7, showing that for a 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.

5.2. 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 are so closely related to coherent risk measures—the only 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:

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

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

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

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., Infanger and Morton 1996, 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 3 and 4 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:

$$\begin{aligned} V_1^K = \min_{x_1, \theta_2} \ & c_1^\top x_1 + \theta_2 \\ & A_1 x_1 = b_1 \\ & x_1 \geq 0 \\ & \theta_2 \geq \alpha_{2,k} + \beta_{2,k}^\top x_1, \quad k = 1, \ldots, K-1 \\ & \theta_2 \geq -M_2, \end{aligned} \tag{22}$$

and:

$$\begin{aligned} V_t^K(x_{t-1}, \omega_t) = \min_{\bar{x}_t, x_t, \theta_{t+1}} \ & c_t^\top x_t + \theta_{t+1} \\ & \bar{x}_t = x_{t-1} \quad [\lambda] \\ & A_t x_t + B_t \bar{x}_t = b_t \\ & x_t \geq 0 \\ & \theta_{t+1} \geq \alpha_{t+1,k} + \beta_{t+1,k}^\top x_t, \quad k = 1, \ldots, K-1 \\ & \theta_{t+1} \geq -M_{t+1}, \end{aligned} \tag{23}$$

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 $t = 1, \ldots, T-1$, 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.

Remark 8. 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

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

Set $K = 1$

while *not converged* **do**

  `// Forward pass`

  solve master problem (22) and obtain solution x_1^K

  for $t = 2, \ldots, T-1$ **do**

    sample ω_t from Ω_t

    solve master problem (23) given (x_{t-1}^K, ω_t) and obtain solution x_t^K

  end

  `// Backward pass`

  for $t = T, \ldots, 2$ **do**

    for $\omega_t \in \Omega_t$ **do**

      solve (23) given (x_{t-1}^K, ω_t) to obtain $V_t^K(x_{t-1}^K, \omega_t)$ and an extreme point dual solution, λ

      set $V_{\omega_t}^K = V_t^K(x_{t-1}^K, \omega_t)$

      set $\lambda_{\omega_t}^K = \lambda$

    end

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

    set $\beta_{t,K} = \sum_{\omega_t \in \Omega_t} q_{\omega_t}^K \lambda_{\omega_t}^*$

    set $\alpha_{t,K} = \sum_{\omega_t \in \Omega_t} q_{\omega_t}^K V_{\omega_t}^K - \alpha(q^K) - \beta_{t,K}^\top x_{t-1}^K$

    Add the cut $\theta_t \geq \alpha_{t,K} + \beta_{t,K}^\top x_t$ to (23) for $t-1$, i.e., updating the model with value V_{t-1}^K to V_{t-1}^{K+1}

  end

  $K \leftarrow K + 1$

end

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 Guigues (2016) can be directly applied to convex risk measures.

Remark 9. Under certain technical assumptions, stochastic dual dynamic programming can be extended to problems involving binary variables, resulting in the SDDiP method of 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. We use SDDiP in the example in Section 6.

5.3. Implementation

We provide an implementation of the entropic risk measure in `SDDP.jl` (Dowson and Kapelevich 2020), a Julia (Bezanson et al. 2017) package for solving multistage stochastic programs via stochastic dual dynamic programming. `SDDP.jl` exposes modular support for researchers to easily implement and test other convex 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 (20). To the best of our knowledge, this formulation is also a novel, if modest, contribution, although we point to work regarding dynamic aggregation of cuts by Trukhanov et al. (2010).

6. Examples

We describe two simple examples that demonstrate the desirability of the entropic risk measure.

6.1. Road networks

Consider a road network with three arcs as shown in Figure 4. 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 $\mathbb{F}[X+Z] \leq \mathbb{F}[Y+Z] \iff \mathbb{F}[X] \leq \mathbb{F}[Y]$.

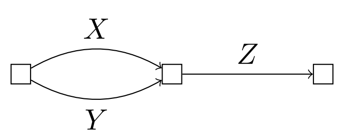

Figure 4 Road network with three roads with random travel times, X, Y, and Z, which are independent.

$\mathbb{P}(\omega)$	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.

Our problem can be formulated as a multistage stochastic integer program and solved using `SDDP.jl` with the SDDiP method of Zou et al. (2019) or by enumeration. We first consider the *end-of-horizon* formulation:

$$V_1 = \min_{x \in \{0,1\}} \underset{(\omega_x, \omega_y, \omega_z) \in \Omega_X \times \Omega_Y \times \Omega_Z}{\mathbb{F}} [V_2(x, \omega_x, \omega_y, \omega_z)],$$

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:

$$V_1 = \min_{x \in \{0,1\}} \underset{(\omega_x, \omega_y) \in \Omega_X \times \Omega_Y}{\mathbb{F}} [V_2(x, \omega_x, \omega_y)],$$

where:

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

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

The results are shown in Figure 5, 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 5a, 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 5b, we use the entropic risk measure. As $\gamma \to \infty$, the measure becomes more risk-averse. By Theorem 1, the end-of-horizon and nested formulations are equivalent. Therefore, there is only one line visible on the graph (both lines lie on top of each other). When γ is small (i.e., < 4.2), the optimal decision is to take road Y. For larger valued of γ the optimal decision is to take road X.

This example shows that the end-of-horizon $\mathbb{CV@R}$ policy is not monotonic in the choice of action with respect to the risk-aversion parameter. However, it is possible to derive examples with non-monotonic policies for the nested $\mathbb{CV@R}$ and entropic risk measure, as we will show in Section 6.2.

This example demonstrates the conditional inconsistency of the $\mathbb{CV@R}$ risk measure. That is, from an end-of-horizon view of risk (e.g., risk with respect to the total travel time), nested $\mathbb{CV@R}$ can lead to a sub-optimal policy. In contrast, due to Theorem 1, when using the entropic risk measure, the end-of-horizon view of risk is equivalent to the nested view of risk.

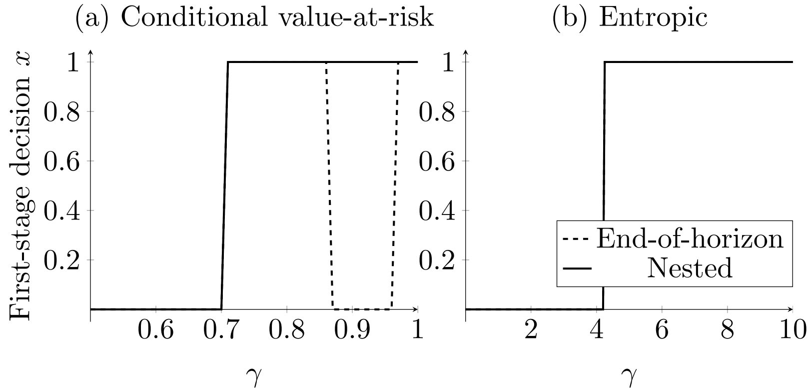

Figure 5 Plot of optimal first-stage decision x against risk aversion parameter γ using the $\mathbb{CV@R}$ (subplot a) and entropic (subplot b) risk measures. Decision $x = 0$ implies choosing road Y, and $x = 1$ implies choosing road X.

6.2. 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. We assume the returns are realized at the start of stage t, *before* the rebalancing decision is made. 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.

Like the road network example, our data are structured in such a way as to expose the conditional inconsistency of $\mathbb{CV@R}$, but again we emphasize this inconsistency is not limited to simple examples and can also arise in large-scale real-world models. We choose to demonstrate the effect on simple models for clarity of presentation and to avoid the complications that inevitably come with larger models.

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 6, along with the fraction of stocks chosen in the first stage, i.e., x_1^s.

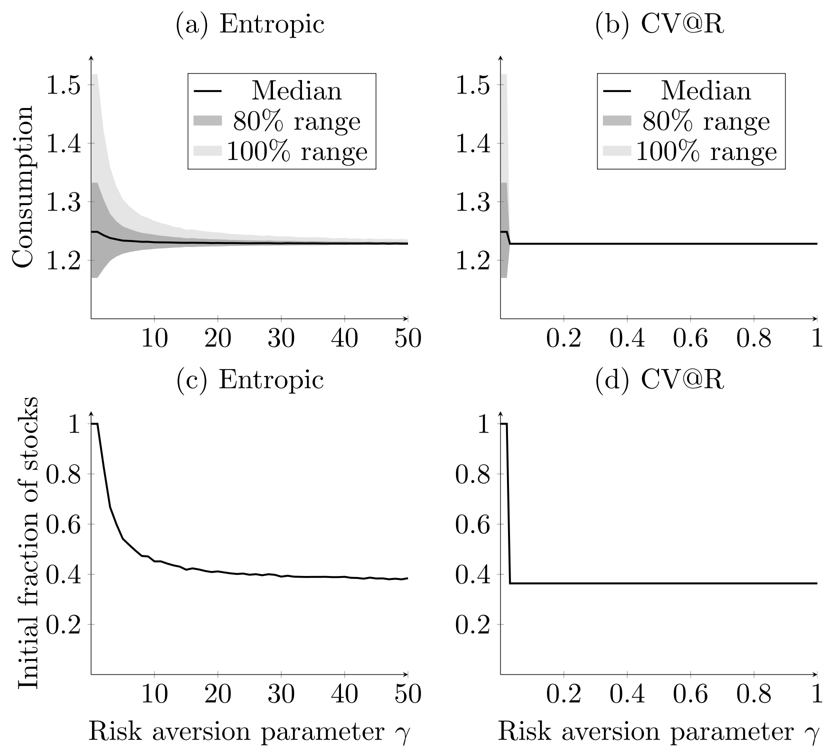

Figure 6 Distribution of consumption and initial fraction of wealth placed in stocks, x_1^s, against the risk aversion parameter γ for the entropic risk measure (a) and (c) and nested $\mathbb{CV@R}$ risk measure (b) and (d).

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.

Finally, we analyze the situation in which the agent can pick stocks or bonds only, must stay with that choice to the horizon, and again consumes only in the final stage. In Figure 7, we plot the first-stage objective function value, i.e., $-V_1((0,1),(1,1))$, against the risk aversion parameter, γ, for the end-of-horizon $\mathbb{CV@R}$, nested $\mathbb{CV@R}$, and entropic risk measures for the "*Stocks Only*" and "*Bonds Only*" decisions. Note that since all of the states and controls are fixed, we can compute this value exactly without needing to solve an optimization problem.

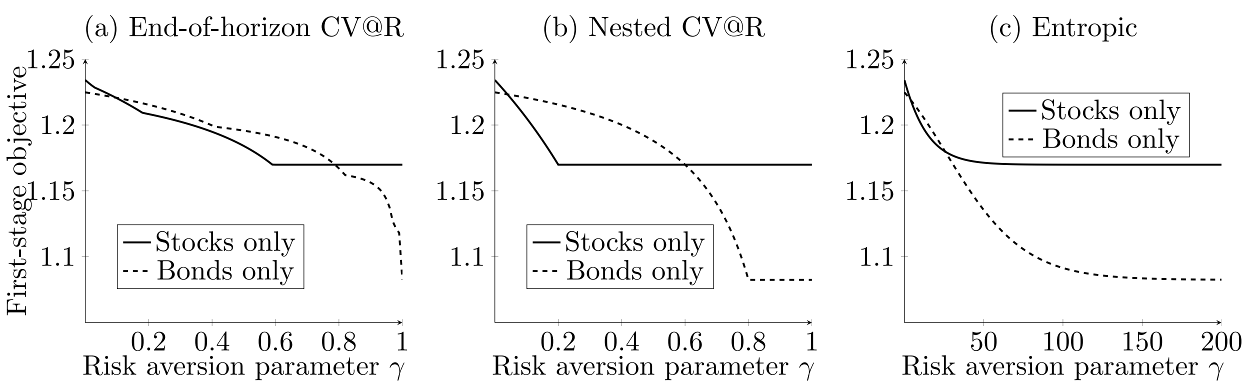

Figure 7 First-stage objective value against the risk aversion parameter γ for end-of-horizon $\mathbb{CV@R}$, nested $\mathbb{CV@R}$, and entropic risk measures.

As we saw in Figure 6, in Figure 7 when the risk aversion parameter is zero, all measures are equivalent to expectation. Therefore, the *Stocks Only* option has a larger first-stage objective than *Bonds Only* because the stocks have a higher expected return. Moreover, when the risk aversion parameter is large, the risk measures are equivalent to the worst-case risk measure. Note that for the entropic risk measure, this occurs when $\gamma \to \infty$, and so the *Bonds Only* line is still decreasing to the right of the plot. For the worst-case risk measure, the *Stocks Only* option has a larger first-stage objective because in the worst-case it returns 4% per stage, compared to only 2% for bonds. However, for all risk measures, there is an intermediate range in which *Bonds Only* has a larger objective. Thus, in the same way that Figure 5 shows the end-of-horizon $\mathbb{CV@R}$ measure can yield non-monotonicity of a first-stage decision as γ grows, we see analogous behavior for all three risk measures in Figure 7.

7. Conclusion

This paper has demonstrated the desirability and viability of using the entropic risk measure in multistage stochastic programming. In particular, we proved that the entropic risk measure is conditionally consistent, and showed how to incorporate any convex risk measure into the stochastic dual dynamic programming algorithm. 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.

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. As a consequence, there is no clear rule for choosing γ; we only know that as γ increases, the measure becomes more risk-averse. Future work includes large-scale computational studies on specific problems on energy, finance, and transportation to gain a better understanding on the effect that convex risk measures have in practice. Finally, to our knowledge it is an open question whether the entropic risk measure is, in some appropriate sense, the only conditionally consistent risk measure.

Acknowledgements

This research was supported, in part, by Northwestern University's Center for Optimization & Statistical Learning (OSL). We thank Benoît Legat for his discussions on conic duality.

References

Ahmadi-Javid A (2012) Entropic value-at-risk: A new coherent risk measure. *Journal of Optimization Theory and Applications* 155(3):1105–1123.

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

Artzner P, Delbaen F, Eber JM, Heath D, Ku H (2007) Coherent multiperiod risk adjusted values and Bellman's principle. *Annals of Operations Research* 152(1):5–22.

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

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).

Benders J (1962) Partitioning procedures for solving mixed-variables programming problems. *Numerische Mathematik* 4:238–252.

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.

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

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.

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.

Dowson O, Kapelevich L (2020) SDDP.jl: A Julia package for stochastic dual dynamic programming. *INFORMS Journal on Computing* To appear.

Duque D, Morton DP (2019) Distributionally robust stochastic dual dynamic programming. *Optimization Online* URL `https://www.optimization-online.org/DB_HTML/2019/12/7539.html`.

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

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.

Infanger G, Morton DP (1996) Cut sharing for multistage stochastic linear programs with interstage dependency. *Mathematical Programming* 75(2):241–256.

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.

Miller N, Ruszczyński A (2011) Risk-averse two-stage stochastic linear programming: Modeling and decomposition. *Operations research* 59(1):125–132.

MOSEK ApS (2019) *MOSEK Optimization Suite 9.1.9*. URL `https://docs.mosek.com/9.1/intro/index.html`.

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.

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.

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

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

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

Shapiro A, Dentcheva D, Ruszczynski A (2009) *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.

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.

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