Risk aversion in multistage stochastic programming: a modeling and algorithmic perspective

Tito Homem-de-Mello
School of Business
Universidad Adolfo Ibañez
Santiago, Chile
tito.hmello@uai.cl

Bernardo K. Pagnoncelli
Department of Systems and Industrial Engineering
School of Business
Universidad Adolfo Ibañez
Santiago, Chile
bernardo.pagnoncelli@uai.cl

February 4, 2015

Abstract

We discuss the incorporation of risk measures into multistage stochastic programs. While much attention has been recently devoted in the literature to this type of model, it appears that there is no consensus on the best way to accomplish that goal. In this paper, we discuss pros and cons of some of the existing approaches. A key notion that must be considered in the analysis is that of consistency, which roughly speaking means that decisions made today should agree with the planning made yesterday for the scenario that actually occurred. Several definitions of consistency have been proposed in the literature, with various levels of rigor; we provide our own definition and give conditions for a multi-period risk measure to be consistent. A popular way to ensure consistency is to nest the one-step risk measures calculated in each stage, but such an approach has drawbacks from the algorithmic viewpoint. We discuss a class of risk measures—which we call expected conditional risk measures—that address those shortcomings. We illustrate the ideas set forth in the paper with numerical results for a pension fund problem in which a company acts as the sponsor of the fund and the participants' plan is defined-benefit.

Keywords: Stochastic programming; risk aversion; multistage; consistency; pension funds

1 Introduction

The evolution and widespread use of stochastic programming is closely related to the increasing computing power made available since the foundation of the field. The important class of two-stage stochastic programs found immediate use in applications since its general framework of first- and second-stage decisions is suitable for a number of real-world problems, see for instance Wallace and Ziemba (2005). Later, the attention turned to multistage stochastic programs (MSSPs), which are a natural extension of two-stage models. In those problems the sequence of events starts with a decision, followed by a realization of a random vector, and then a decision is made knowing the outcome of the random vector, a new realization occurs, and so on. Randomness is often described by a continuous stochastic process or by a discrete process with a very high number of possible outcomes. A common approach is to build a *scenario tree*, which is a discrete representation that in some sense is close to the original process according to some distance. The generation of scenario trees has received a great deal of attention in the literature: see, for instance, Pflug (2001); Høyland and Wallace (2001); Dupačová et al. (2003); Heitsch and Römisch (2009); Pflug and Pichler (2011, 2012); Mehrotra and Papp (2013). Scenario trees are crucial for numerical solution of the problem, as algorithms used in practice to solve the problem are typically rooted in some decomposition principles such as the Nested Decomposition (Donohue and Birge, 2006) or the Stochastic Dual Dynamic Programming scheme (Pereira and Pinto, 1991). MSSPs have been used in a number of areas, including finance, revenue management, energy planning, and natural resources management, among others.

The classical formulation of stochastic programs (in two or more stages) optimizes the expected value of an objective function that depends on the decision variables as well as on the random variables that represent the uncertainty in the problem. Such a formulation assumes that the decision maker is risk-neutral, i.e., he or she will not mind large losses in some scenarios as long as those are offset by large gains in other scenarios. While such an approach is useful in a number of applications, it does not reflect the situation where the decision is very concerned about large losses—in other words, such a decision maker is *risk-averse*. It is natural then to consider risk-averse formulation of stochastic programs.

In the case of two-stage models, the structure of a first-stage deterministic cost plus a random recourse cost in the second stage makes the extension to the risk-averse case immediate from a modeling perspective, in the sense that the natural choice is to replace the expectation of the second stage cost with some other risk measure; see, for instance, Schultz and Tiedemann (2006), Fábián (2008), Shapiro et al. (2009) and Miller and Ruszczyński (2011). The difficulty associated with the risk-averse model depends on the choice of the risk measure. Ahmed (2006) shows that the if the risk measure is the variance, then the resulting problem is NP-hard. Furthermore, monotonicity in the second stage would be lost, and the cost units of first and second stage would be different unless the standard deviation was used, but the problem would likely become intractable in this case. Rockafellar and Uryasev (2000) show that if the Conditional Value-at-Risk is chosen, the sampled version of the continuous problem can be approximated by a linear programming problem. Noyan (2012) proposes two decomposition algorithms to efficiently solve a disaster management problem with the Conditional Value-at-Risk as the risk measure.

For multistage stochastic programming the picture is quite different and several questions arise. When sequential decision is involved, there is no natural or obvious way of measuring risk. Should risk be measured at every stage separately? Should it be applied to the several scenario paths in the tree? Or should risk be measured in a nested way, in the spirit of dynamic programming? What if only the risk at the end of the time horizon is relevant and we do not want to measure risk at the other stages? The difficulty in extending risk measures to the multistage setting has been discussed in several papers and it can be argued that the differences among the approaches are far more significant than the two-stage case.

A number of recent papers have considered the importance of measuring risk in MSSPs (see, for instance, Eichhorn and Römisch 2005; Pflug and Römisch 2007; Collado et al. 2012; Shapiro 2012a; Philpott and de Matos 2012; Philpott et al. 2013; Shapiro et al. 2013; Kozmık and Morton 2013; Pagnoncelli and Piazza 2012; Guigues and Sagastizábal 2013; Pflug and Pichler 2014b). Several of these papers focus on how to adapt existing algorithms from the risk-neutral case to the risk-averse case, often with the Conditional Value-at-Risk as the risk measure. One of the goals of our paper is to address some of the popular ways to measure risk and discuss their advantages and drawbacks. In addition, we revisit and extend a class of multi-period risk measures proposed by Pflug and Ruszczyński (2005) (see also Pflug 2006 for a more extensive discussion), which we call *expected conditional risk measures (ECRMs)*, and discuss how the resulting problem can be efficiently solved. ECRMs combine two attractive features: on the one hand, ECRMs can be represented in a nested form, a feature that is desirable and the focus of much of the recent literature, as we shall see later; on the other hand, we show that when ECRMs are applied with the Conditional Value-at-Risk (CVaR) as the underlying risk measure, the resulting MSSP can be represented by a simpler risk-neutral MSSP with additional variables, much in the spirit of the polyhedral risk measures introduced by Eichhorn and Römisch (2005).

As it has been observed in the literature, one very important issue that arises when modeling risk-averse MSSPs is that of *time consistency*. Time consistency in MSSPs has been highlighted by several authors in recent years as a desirable property a problem should have. Informally, time consistency means that if you solve an MSSP today and find solutions for each node of a tree, you should find the same solutions if you re-solve the problem tomorrow given what was observed and decided today. The definitions in the literature differ mainly by their focus: the works of Ruszczyński (2010) and Kovacevic and Pflug (2014) deal with sequences of random variables, while Detlefsen and Scandolo (2005), Cheridito et al. (2006), and Bion-Nadal (2008), define time consistency for continuous-time dynamic models. The definitions in Shapiro (2009), Carpentier et al. (2012), Rudloff et al. (2014) and De Lara and Leclère (2014) are centered on optimization and on the stability of decision variables at every stage. Xin et al. (2013) propose definitions of time-consistency of policies in the context of distributionally robust MSSPs, whereas Pflug and Pichler (2014a) propose a related notion of time-consistent decisions. We propose a new definition of consistency, closer to the optimization-oriented papers. Our definition is suitable for MSSPs that can be represented via scenario trees. Using a simple three-stage inventory problem we show that several natural ways of measuring risk lead to inconsistent formulations, according to our definition. We also show the class of ECRMs we study in this paper is time-consistent.

We illustrate the applicability of ECRMs by using it in a pension fund problem proposed by Haneveld et al. (2010). This numerical example illustrates two important aspects of ECRMs: first, the simplicity of implementation when the CVaR is used as an ingredient for the ECRM—indeed, we use standard software for risk-neutral multistage programs available in the literature to solve the corresponding risk-averse problem. The second important aspect is the flexibility allowed by the model to represent the change in the degree of risk aversion over time; for example, the decision maker may be more risk-averse about the earlier stages and less risk-averse about the stages farther in the future. We also use the numerical example to propose a (to the best of our knowledge) novel way to compare optimal solutions of MSSPs. The majority of applications only analyzes the first-stage solution since in most cases a rolling-horizon procedure will be implemented in practice and the solutions of other stages will not be implemented. We show in our pension fund example that the solutions of subsequent stages carry important information concerning the quality and robustness of the first-stage solution. By using first- and second-order dominance we show that despite having an attractive first-stage allocation, some solutions exhibit a poor behavior in subsequent stages, such as having a very high probability of needing extra money injection in the fund.

The rest of the paper is organized as follows. Section 2 defines our notion of consistency. In Section 3 we present an inventory problem that illustrates our notion of consistency and discuss several modeling paradigms for risk-averse MSSP. We prove some results that characterize consistency according to our definition in Section 4. In Section 5 we introduce the notion of ECRMs and study in detail their properties, including consistency and the equivalent risk-neutral formulation of the case with CVaR. The pension fund example that illustrates our approach is presented in Section 6, while Section 7 presents some concluding remarks.

2 Consistency

We start by defining precisely the notation and the class of problems we want to study. Consider a probability space $(\Omega, \mathcal{F}, P)$, and let $\mathcal{F}_1 \subset \mathcal{F}_2 \subset \ldots \mathcal{F}_T$ be sigma sub-algebras of $\mathcal{F}$ such that each $\mathcal{F}_t$ corresponds to the information available up to (and including) stage t, with $\mathcal{F}_1 = \{\emptyset, \Omega\}$ 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 let $\mathcal{Z} := \mathcal{Z}_1 \times \ldots \times \mathcal{Z}_T$. We define a *multi-period risk function* $\mathbb{F}$ as a mapping from $\mathcal{Z}$ to $\mathbb{R}$. For example, we may have, for $Z \in \mathcal{Z}$,

$$\mathbb{F}(Z) \;=\; \mathbb{F}_1(Z_1) + \ldots + \mathbb{F}_T(Z_T),$$

where each $\mathbb{F}_t$ is a *one-period risk function*, i.e., a mapping from $\mathcal{Z}_t$ to $\mathbb{R}$.

Consider now the space $\mathcal{D}_T$ of distributions of T-dimensional random vectors in $\mathcal{Z}$. That is, each element $G \in \mathcal{D}_T$ — which is a mapping from $\mathcal{B}_T$ to $[0,1]$, where $\mathcal{B}_T$ is the Borel sigma-algebra in $\mathbb{R}^T$ — can be written as the distribution function G_Z of some $Z = (Z_1, \ldots, Z_T) \in \mathcal{Z}$, defined as

$$G_Z(B_T) \;:=\; P(Z \in B_T).$$

Note that Z is not uniquely defined, i.e., we may have $G = G_Z = G_Y$ for two different random vectors Z and Y. For the purpose of the developments in this paper, we shall restrict ourselves to multi-period risk functions $\mathbb{F}$ that are *law-invariant*, in the sense that if Z and Y are elements of $\mathcal{Z}$ such that Z and Y have the same distribution (i.e., $G_Z(B_T) = G_Y(B_T)$ for all $B_T \in \mathcal{B}_T$), then $\mathbb{F}(Z) = \mathbb{F}(Y)$. In that case, given a multi-period risk function $\mathbb{F}$ we can associate $\mathbb{F}$ with a unique mapping $\tilde{\mathbb{F}} : \mathcal{D}_T \mapsto \mathbb{R}$ such that

$$\tilde{\mathbb{F}}(G_Z) \ := \ \mathbb{F}(Z).$$

Following this line of reasoning, we can define the notion of a conditional risk function corresponding to a multi-period risk function $\mathbb{F}$ as follows. Given $Z \in \mathcal{Z}$ and $(z_1, \ldots, z_t) \in \mathbb{R}^t$, let $G_{Z \,|\, Z_1=z_1,\ldots,Z_t=z_t}$ denote the conditional distribution of Z given $Z_1 = z_1, \ldots, Z_t = z_t$, which lies in $\mathcal{D}_T$. We then define the *conditional risk function* $\mathbb{F}$ *of* Z *given* $Z_1, \ldots, Z_t$ (written as $\mathbb{F}^{Z_1,\ldots,Z_t}(Z)$) as a mapping from $\mathcal{Z}$ to $\mathcal{Z}_t$ such that on the event $\{\omega \,:\, Z_1(\omega) = z_1, \ldots, Z_t(\omega) = z_t\}$ we have

$$\mathbb{F}^{Z_1,\ldots,Z_t}(Z)(\omega) := \ \tilde{\mathbb{F}}(G_{Z \,|\, Z_1=z_1,\ldots,Z_t=z_t}). \tag{2.1}$$

Note that $\mathbb{F}^{Z_1,\ldots,Z_t}(Z)$ is an element of $\mathcal{Z}_t$.

Remark 1: It is useful to compare the above notion of a conditional risk function with that of a *conditional risk mapping* defined in Ruszczyński and Shapiro (2006). In that paper, the authors define a conditional risk mapping as a mapping between two linear spaces (for example, from $\mathcal{Z}_t$ to $\mathcal{Z}_{t-1}$) that satisfies some axioms of convexity, monotonicity and translation invariance. Our notion is different in that it presupposes the existence of a multi-period risk function $\mathbb{F}$ on the entire product space $\mathcal{Z}$, and then the conditional risk functions are defined in terms of conditional distributions. Although our notion is in a sense less general than that of Ruszczyński and Shapiro (2006) (as it requires the risk functions to be law-invariant), we believe it is convenient from a practical viewpoint. A risk-averse decision maker might want to simply replace the single external expected value in MSSP by another risk measure, without having to define the appropriate conditional risk mappings in each stage as in Ruszczyński and Shapiro (2006). Note, however, that we can create conditional risk mappings from a special class of the conditional risk functions defined in (2.1) using an additive risk function $\mathbb{F}$, as discussed in Proposition 2.1 below.

Proposition 2.1. *Let $\mathbb{F}$ be a multi-period risk function (as defined earlier) such that*

$$\mathbb{F}(Z_1, \ldots, Z_T) \ = \ \mathbb{F}_0(Z_1 + \ldots + Z_T), \tag{2.2}$$

for some one-period risk function $\mathbb{F}_0 : \mathcal{Z}_T \mapsto \mathbb{R}$. Given $Z_t \in \mathcal{Z}_t$, consider the conditional risk function $\mathbb{F}^{Z_t}$ constructed from $\mathbb{F}$ as in (2.1) and define, accordingly, $\mathbb{F}_0^{Z_t}$ (so $\mathbb{F}_0^{Z_t}$ is a mapping from $\mathcal{Z}_T$ to $\mathcal{Z}_t$). Suppose that the one-period risk function $\mathbb{F}_0$ satisfies the following properties:

(i) Convexity *: If $\alpha \in [0, 1]$ and $Y, W \in \mathcal{Z}_T$, then*

$$\mathbb{F}_0(\alpha W + (1 - \alpha)Y) \ \leq \ \alpha \mathbb{F}_0(W) + (1 - \alpha)\mathbb{F}_0(Y).$$

(ii) Monotonicity *: If $W, Y \in \mathcal{Z}_T$ are such that $W \geq Y$ w.p.1, then $\mathbb{F}_0(W) \geq \mathbb{F}_0(Y)$.*

(iii) Translation invariance *: If $W \in \mathcal{Z}_T$ and c is a constant, then*

$$\mathbb{F}_0(W + c) = \mathbb{F}_0(W) + c.$$

Then, $\mathbb{F}_0^{Z_t}$ satisfies the conditions to be a conditional risk mapping as defined in Ruszczyński and Shapiro (2006).

Proof: The proof follows directly from the fact that, for each realization $Z_t = z_t$, the conditional risk function $\mathbb{F}_0^{Z_t}$ defined from $\mathbb{F}^{Z_t}$ in (2.1) corresponds to $\mathbb{F}_0$ applied to some element of $\mathcal{Z}_T$ (namely, the random variable whose distribution is $G(\cdot \mid Z_t = z_t)$), so properties (A1) and (A2) in Ruszczyński and Shapiro (2006) follow immediately from (i) and (ii). Property (A3) in Ruszczyński and Shapiro (2006) follows from (iii) and the observation that on the event $\{Z_t = z_t\}$ the random variable Z_t is a constant. □

In the course of our discussion we will often refer to one-period conditional risk measures; those should be understood as the $\mathbb{F}_0^{Z_t}$ in Proposition 2.1. Note that several one-period risk measures can be expressed as functions of expectations; this is the case, for example, of CVaR and mean semi-deviations, among others. In such cases, the conditional risk function $\mathbb{F}^{Z_t}$ corresponds to replacing the expectations with conditional expectations $\mathbb{E}[\cdot \mid Z_t]$. More generally, one can apply a similar procedure to the dual representation of a coherent risk measures to obtain its conditional version, as done in Pflug and Pichler (2014a).

Remark 2: It is important to notice that the notion of law invariance defined above is different from that in Shapiro (2012b). In that paper, a risk function of the form (2.2) is called law invariant if $\mathbb{F}_0(Z_1 + \ldots + Z_T) = \mathbb{F}_0(W_1 + \ldots + W_T)$ whenever $Z_1 + \ldots + Z_T$ is equal in distribution to $W_1 + \ldots + W_T$. Our notion, in contrast, requires that the values of the risk functions be equal whenever the *joint* distribution of $Z_1, \ldots, Z_T$ is the same as the joint distribution of $W_1, \ldots, W_T$. In other words, our requirement for law invariance is weaker in the sense that we only require the risk function values to be the same when the whole processes are equivalent. This is an important distinction, since in our view the decision maker may assign different risks to the total cost $Z_1 + \ldots + Z_T$ depending on how that cost was achieved — for example, a portfolio manager may consider a portfolio with many ups-and-downs a lot riskier than one with constant returns, even if the final wealths are equal in distribution. Figure 1 illustrates the issue. An examination of the two trees shows that the distribution of total costs (i.e., second plus third period) is the same for both trees. However, the tree on the left-side of the figure has more variability in period 2 than the one on the right side, which may be important for the decision maker. Under the notion of law invariance in Shapiro (2012b), the two trees will be assigned the same risk, whereas our definition allows for different values. □

With the above notation at hand, we formulate the class of problems we are interested in studying. Below and henceforth, the notation $a_{[t]}$ indicates the collection $a_1, \ldots, a_t$.

$$\begin{aligned} \min_{x_1,\ldots,x_T} \; & \mathbb{F}\left(f_1(x_1,\xi_1), \ldots, f_T(x_T,\xi_T)\right) \\ \text{s.t.} \quad & x_t \in \mathcal{X}_t\left(x_{[t-1]}, \xi_{[t]}\right), \quad t = 1, \ldots, T. \end{aligned} \tag{$\mathcal{P}$}$$

In the above model, $x_t \in \mathbb{R}^{n_t}$ denotes the decision made in stage t; ξ_t is an m_t-dimensional random vector representing the uncertainty observed in stage t, i.e., ξ_t is an $\mathcal{F}_t$-measurable mapping from Ω to $\mathbb{R}^{m_t}$; f_t is a function from $\mathbb{R}^{n_t} \times \mathbb{R}^{m_t}$ that corresponds to the cost of decision x_t given the observed uncertainty $\xi_t(\omega)$ in that stage; $\mathcal{X}_t\left(x_{[t-1]}, \xi_{[t]}\right)$ denotes

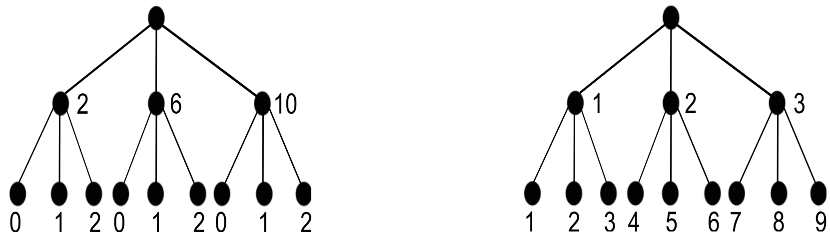

Figure 1: Two trees with different joint distributions but identical distributions of total costs.

the feasibility set in stage t, which may depend on previous decisions as well as on the observed uncertainty. We write $\mathbb{F}(f_1(x_1,\xi_1),\ldots,f_T(x_T,\xi_T))$ as a short for $\mathbb{F}(Z_1,\ldots,Z_T)$, with $Z_t \in \mathcal{Z}_t$ defined as $Z_t(\omega) := f_t(x_t,\xi_t(\omega))$. It is clear from the above definition that any feasible solution $x := [x_1,\ldots,x_T]$ to $(\mathcal{P})$ is such that each x_t is actually a function of $\xi_2,\ldots,\xi_t$, though we make that dependence explicit only when necessary to avoid cluttering the notation.

We propose now our notion of *consistency*, which is inspired by the different definitions available in the literature. Consider the problem of solving $(\mathcal{P})$ at a given stage t, when all the information from previous stages (given by $\hat{x}_{[t-1]}$ and $\hat{\xi}_{[t]}$) is known. That is, we have the following optimization problem to solve:

$$\begin{aligned} \min_{x_t,\ldots,x_T} \quad & \mathbb{F}^{\hat{\xi}_{[t]}}\left(f_1(x_1,\xi_1),\ldots,f_T(x_T,\xi_T)\right) \\ \text{s.t.} \quad & x_\tau \in \mathcal{X}_\tau\left(\hat{x}_{[t-1]},x_t,\ldots,x_{\tau-1},\hat{\xi}_{[t]},\xi_{t+1},\ldots,\xi_\tau\right), \quad \tau=t,\ldots,T. \end{aligned} \tag{$\mathcal{P}_t$}$$

In the above, the notation $\mathbb{F}^{\hat{\xi}_{[t]}}$ indicates a *conditional risk function* as defined earlier, i.e., the multi-period risk function $\mathbb{F}$ in $(\mathcal{P})$ applied to the random vector $f_1(x_1,\xi_1),\ldots,f_T(x_T,\xi_T)$, conditional on a given realization $\hat{\xi}_1,\ldots,\hat{\xi}_t$ (and implicitly on $\hat{x}_{[t-1]}$). Note that under such conditions, $f_1(\hat{x}_1,\hat{\xi}_1),\ldots,f_{t-1}(\hat{x}_{t-1},\hat{\xi}_{t-1})$ are constants, and so is $\hat{\xi}_t$.

Let $\left[\bar{x}_\tau^{t,\hat{x}_{[t-1]},\hat{\xi}_{[t]}} : \tau=t,\ldots,T\right]$ denote an optimal solution of $(\mathcal{P}_t)$. We include t, $\hat{x}_{[t-1]}$ and $\hat{\xi}_{[t]}$ as superscripts to emphasize that such a solution is calculated at time t, given the previous stages decisions $\hat{x}_1,\ldots,\hat{x}_{t-1}$ and conditional on a given realization $\hat{\xi}_1,\ldots,\hat{\xi}_t$. Moreover, as mentioned earlier, each $\bar{x}_\tau^{t,\hat{x}_{[t-1]},\hat{\xi}_{[t]}}$, $\tau=t,\ldots,T$, is a function of $\xi_{t+1},\ldots,\xi_\tau$.

Definition 2.2. *We say that the* inherited optimality property (henceforth called IOP) *holds for an instance*[1] *of problem $(\mathcal{P})$ if, given any time period t such that $1<t\le T$ and any realization $\hat{\xi}_1,\ldots,\hat{\xi}_t$, there exists an optimal solution x^* of $(\mathcal{P})$ such that the solution "inherited" from x^* at $\hat{\xi}_2,\ldots,\hat{\xi}_t$ (denoted as $\left[x_\tau^*(\hat{\xi}_2,\ldots,\hat{\xi}_t,\cdot):\tau=t,\ldots,T\right]$, where "$(\cdot)$" indicates this is a function of $\xi_{t+1},\ldots,\xi_\tau$) coincides with an optimal solution of $(\mathcal{P}_t)$ for those t, $\hat{\xi}$, and $\hat{x}:=x^*$.*

[1]An instance is an specification of functions $f_t(x_t,\xi_t)$, the feasible sets $\mathcal{X}_t$ as well as the distribution of random vector $(\xi_1,\ldots,\xi_T)$.

It is important to observe in the above definition that, in general, problem $(\mathcal{P}_t)$ may have many optimal solutions of the form $\left[\bar{x}_\tau^{t,\hat{x}_{[t-1]},\hat{\xi}_{[t]}}(\cdot) : \tau = t, \ldots, T\right]$. Likewise, problem $(\mathcal{P})$ may have multiple optimal solutions as well. When solving problem $(\mathcal{P}_t)$ one might hit upon an optimal solution which is different from the one inherited from period 1. The IOP only requires that *one* of the optimal solutions of $(\mathcal{P}_t)$ coincide with *some* inherited solution from period 1. Of course, when $(\mathcal{P})$ and $(\mathcal{P}_t)$ have *unique* optimal solutions, then the optimal solution of $(\mathcal{P}_t)$ must be the solution "inherited" from the optimal solution in period 1.

We are now in a position to define consistency:

Definition 2.3. *We say that the multi-period risk measure* $\mathbb{F}$ *is* consistent *for problems of the form* $(\mathcal{P})$ *if the IOP holds for any particular instance of that problem.*

Note that, by definition, our notion of consistency is independent of any particular instance of the problem; for example, when working with scenario trees, a multi-period risk measure $\mathbb{F}$ is consistent no matter which realization of the tree we are considering. We believe consistency is a desired property a risk-averse multistage stochastic program should possess. It can be understood as some sort of stability property: the decision you make today should agree with *some* optimal plan made yesterday given what was observed today. While a somewhat natural property, consistency does not hold automatically. In the next section we present an example of a simple problem that illustrates several different modeling frameworks for multistage risk-averse stochastic programming and discuss their advantages and disadvantages, including lack of consistency.

3 An inventory problem

In this section we present a small inventory problem that illustrates our notion of consistency. To keep the calculations simple, the problem has only three stages. We write the formulations and calculations explicitly to illustrate the characteristics of the resulting models, a feature that in our view helps to illustrate the consequences of choosing among the various ways of measuring risk.

3.1 Conditional Value-at-Risk

We consider two different ways of incorporating risk into the problem, and for those examples we will use the Conditional Value-at-Risk (CVaR) as our risk measure. The choice is justified by the extensive use of this risk measure in a vast array of applications, as well as by its desired properties. A risk measure is said to be coherent according to Artzner et al. (1999) if it satisfies the following properties (below, $\mathcal{W}$ is a linear space of random variables):

Translation invariance: If $a \in \mathbb{R}$ and $W \in \mathcal{W}$, then $\rho(W + a) = a + \rho(W)$.

Positive homogeneity: If $\mathrm{c} > 0$ e $Z \in \mathcal{W}$ then $\rho(\mathrm{c}W) = \mathrm{c}\rho(W)$.

Monotonicity: If $W_1, W_2 \in \mathcal{W}$ and $W_1 \leq W_2$, then $\rho(W_1) \leq \rho(W_2)$.

Convexity: If $W_1, W_2 \in \mathcal{W}$ and $\lambda \in (0,1)$, then $\rho(\lambda W_1 + (1-\lambda)W_2) \le \lambda\rho(W_1) + (1-\lambda)\rho(W_2)$.

It can be shown that the CVaR is a coherent risk measure Pflug (2000).

A key result in Rockafellar and Uryasev (2002) is the proof that CVaR can be expressed as the optimal value of the following optimization problem:

$$\mathrm{CVaR}_\alpha[X] = \min_{\eta\in\mathbb{R}} \left\{ \eta + \frac{1}{1-\alpha}\mathbb{E}\left[(X-\eta)_+\right] \right\}, \tag{3.1}$$

where $(a)_+ := \max(a, 0)$. It is easy to see from representation (3.1) that CVaR is law-invariant. and that the conditional version of CVaR (as in the context of Proposition 2.1) is obtained by conditioning the expectation in (3.1).

3.2 Problem formulation

Consider a decision maker trying to sell a certain product. Each unit produced has a cost c, and there are two opportunities to sell the product, at time $t = 2$ and time $t = 3$, at prices s_2 and s_3 respectively, with $s_3 > s_2$. The demand D_2 at time 2 is revealed after the decision x_1 of how many products should be manufactured is made, whereas the demand D_3 at time 3 is revealed after the decision x_2 of how many products to sell at price s_2 is made. The decision of how many products to sell at the end of the horizon at price s_3 — when all the uncertainty has been realized — is denoted by x_3. Unsold products at the end of the horizon have no value.

Let us write a minimization problem, so costs are positive and gains are negative. In that case, the total cost is $cx_1 - s_2x_2 - s_3x_3$. A perhaps intuitive way to formulate the problem with risk-aversion is to measure the risk of the total cost as a single quantity, i.e., to take $\mathbb{F}(Z_1, \ldots, Z_T) = \rho(Z_1 + \ldots + Z_T)$, where ρ is a one-period law-invariant risk function. Clearly, this implies that $\mathbb{F}$ is a law-invariant multi-period risk function. Suppose also that ρ satisfies the translation-invariance property (iii) of Proposition 2.1. In this particular example we have

$$\mathbb{F}\big(f_1(x_1,\xi_1), f_2(x_2,\xi_2), f_3(x_3,\xi_3)\big) = \mathbb{F}(cx_1, -s_2x_2, -s_3x_3) = \rho(cx_1 - s_2x_2 - s_3x_3).$$

For an arbitrary one-period risk function ρ satisfying the above conditions the problem can be formulated as follows:

$$\begin{aligned} \min \;\; & cx_1 + \rho\left(-s_2x_2 - s_3x_3\right) \\ \text{s.t.} \;\; & x_2 \le \min\{D_2,\ x_1\} \\ & x_3 \le \min\{D_3,\ x_2\} \\ & x_1, x_2, x_3 \ge 0, \end{aligned} \tag{3.2}$$

where x_2 is a function of D_2 and x_3 is a function of D_2 and D_3. Note that the term cx_1 moves out of the risk measure due to the translation-invariance property, since in this case cx_1 is not random. Suppose now that we use the risk measure $\rho = \mathrm{CVaR}_\alpha$. Using the optimization formulation (3.1) for CVaR, we can write problem (3.2) in the following equivalent form:

$$\begin{aligned} \min \; & cx_1 + \eta + \frac{1}{1-\alpha}\mathbb{E}\left[(-s_2 x_2(D_2) - s_3 x_3(D_2, D_3) - \eta)_+\right] \\ \text{s.t.} \quad x_2(D_2) & \le \min\{D_2,\ x_1\} \\ x_3(D_2, D_3) & \le \min\{D_3|D_2,\ x_1 - x_2(D_2)\} \\ \eta & \in \mathbb{R}. \end{aligned} \tag{3.3}$$

In the above formulation, we abuse the notation by writing $D_3|D_2$ to indicate a collection of random variables indexed by the values taken by D_2, such that the random variable indexed by $D_2 = d_2$ is defined on $\Omega_{d_2} := \Omega \cap \{\omega \in \Omega : D_2(\omega) = d_2\}$, and is given by the restriction of D_3 to Ω_{d_2}. It is clear that when D_2 and D_3 have finite support the above problem can be written as a linear program. Note also that the auxiliary variable η is not a function of D_2 or D_3.

Suppose that demand is distributed as in Figure 2, with equal probabilities on each branch. The purchase price is $c = 2$, and the sale prices are $s_2 = 3$ in the second stage and $s_3 = 10$ in the third stage.

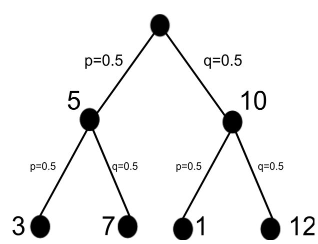

Figure 2: The demand tree.

The case of α close to 1. Let us first take $\alpha = 1 - \epsilon$ for some arbitrarily small ϵ. Such a choice corresponds to using the one-period worst-case risk measure given by $\rho(X) = \text{ess sup}\,(X)$. By solving problem (3.3) we obtain the following optimal policy:

$$\begin{aligned} \hat{x}_1 &= 11, \\ \hat{x}_2(5) = 5; &\quad \hat{x}_2(10) = 10, \\ \hat{x}_3(5,3) = 3; &\quad \hat{x}_3(5,7) = 3; \quad \hat{x}_3(10,1) = 1; \quad \hat{x}_3(10,12) = 1. \end{aligned} \tag{3.4}$$

Now suppose we advance one period and consider the subproblem consisting of the left-hand side of the tree in Figure 2, that is, demand at $t = 2$ is $\hat{D}_2 = 5$ and we want to solve the problem at this node. Since $\hat{x}_1 = 11$, our decisions are what to do at times 2 and 3 with the 11 units we have at hand. To test whether the IOP holds or not, let us solve problem $(\mathcal{P}_t)$ for $t = 2$. In line with the notation of problem $(\mathcal{P}_t)$, the variables in that problem can be identified as $x_t^{2,\hat{x}_1,\hat{D}_2}$, $t = 2, 3$. Since we are considering the realization $\hat{D}_2 = 5$, we shall write $x_t^{2,5}$ for short. The problem can then be written as

$$\begin{aligned}
\min \ \ & -s_2 x_2^{2,5} + \eta^{2,5} + \frac{1}{1-\alpha}\mathbb{E}\left[(-s_3 x_3^{2,5}(D_3) - \eta^{2,5})_+ \,|\, D_2 = 5\right] \\
\text{s.t.} \quad & x_2^{2,5} \ \le \ \min\{5,\ \hat{x}_1\} \\
& x_3^{2,5}(D_3) \ \le \ \min\{D_3|(D_2=5),\ \hat{x}_1 - x_2^{2,5}\} \\
& x_2^{2,5}, x_3^{2,5} \ \ge \ 0. \\
& \eta^{2,5} \ \in \ \mathbb{R}.
\end{aligned} \tag{3.5}$$

Note the presence of the auxiliary variable $\eta^{2,5}$ that is introduced to model the conditional risk measure $\mathrm{CVaR}_\alpha(\cdot|D_2 = 5)$. By solving this problem we obtain

$$\begin{aligned}
\bar{x}_2^{2,5} &= 5, \\
\bar{x}^{2,5}(3) = 3; &\quad \bar{x}^{2,5}(7) = 3,
\end{aligned} \tag{3.6}$$

with optimal value -45. We see that this solution actually coincides with the solution obtained earlier calculated at $\hat{D}_2 = 5$, i.e., $\bar{x}_2^{2,5} = \hat{x}_2(5) = 5$, $\bar{x}_3^{2,5}(3) = \hat{x}_3(5,3) = 3$, and $\bar{x}_3^{2,5}(7) = \hat{x}_3(5,7) = 3$. It is easy to see that the same phenomenon occurs when $\hat{D}_2 = 10$.

To finish the test for the IOP, we need to solve problem $(\mathcal{P}_t)$ for $t = 3$. Since the third-stage variable x_3 is essentially the minimum between the observed demand in that stage and the slack between the purchased amount x_1 and the amount x_2 sold in the second stage, it is easy to verify that there exists some optimal solution $\check{x}$ of (3.3) — in this case, $\check{x} = \hat{x}$ except that $\check{x}_3(5,7) = 6$ — such that the solution inherited from $\check{x}$ is optimal for $(\mathcal{P}_t)$. Thus, the IOP holds for this problem instance when the risk measure used is $\mathbb{F}(Z_1, \ldots, Z_T) = \text{ess sup } (Z_1 + \ldots + Z_T)$.

It is worthwhile observing that the solution

$$\begin{aligned}
\tilde{x}_1 &= 11, \\
\tilde{x}_2(5) = 4; &\quad \tilde{x}_2(10) = 10, \\
\tilde{x}_3(5,3) = 3; \quad \tilde{x}_3(5,7) = 3; &\quad \tilde{x}_3(10,1) = 1; \quad \tilde{x}_3(10,12) = 1
\end{aligned} \tag{3.7}$$

is also an optimal solution of (3.3) that coincides with $\hat{x}$ in period 1, but the solution of problem (3.5) obtained by fixing $x_2^{2,5} = \tilde{x}_2(5) = 4$, $x_3^{2,5}(3) = \tilde{x}_3(5,3) = 3$, and $x_3^{2,5}(7) = \tilde{x}_3(5,7) = 3$ yields an objective value equal to -42. Thus, the solution "inherited" from this $\tilde{x}$ is not optimal for the subproblem. The goal of this remark is to emphasize that the proposed notion of consistency requires only that *some* inherited solution be optimal for the subproblem, we do not require that *all* inherited solutions be optimal — it is tempting to call the latter property "strong consistency" but, as the above example demonstrates, such a property is unlikely to hold in practice and it would be difficult to verify it unless the optimal solutions are unique, in which case it would be the same as consistency.

The case of intermediate α. Suppose we take again $\mathbb{F}(cx_1, -s_2x_2, -s_3x_3) = \mathrm{CVaR}_\alpha(cx_1 - s_2x_2 - s_3x_3)$, but with $\alpha = 0.3$. By solving problem (3.3) we obtain the following optimal policy:

$$
\begin{aligned}
&\hat{x}_1 = 12, \\
&\hat{x}_2(5) = 5; \quad \hat{x}_2(10) = 5, \\
&\hat{x}_3(5,3) = 3; \quad \hat{x}_3(5,7) = 7; \quad \hat{x}_3(10,1) = 1; \quad \hat{x}_3(10,12) = 7.
\end{aligned}
\tag{3.8}
$$

An analysis of the reduced costs shows that such this optimal solution is *unique*. Now, the solution of the subproblem (3.5) yields

$$
\begin{aligned}
&\bar{x}_2^{2,5} = 5, \\
&\bar{x}^{2,5}(3) = 3; \quad \bar{x}^{2,5}(7) = 7,
\end{aligned}
$$

so we see that this solution actually coincides with the solution $\hat{x}$ in (3.8) calculated at $\hat{D}_2 = 5$. However, the solution of the subproblem analogous to (3.5) for $\hat{D}_2 = 10$ yields

$$
\begin{aligned}
&\bar{x}_2^{2,10} = 10, \\
&\bar{x}^{2,10}(1) = 1; \quad \bar{x}^{2,10}(12) = 2,
\end{aligned}
$$

with an optimal value of -42.85, which is better than the solution inherited from $\hat{x}$, that is, $\hat{x}_2(10) = 5$, $\hat{x}_3(10,1) = 1$, $\hat{x}_3(10,12) = 7$, which yields an objective value equal to -42.14. Thus, we see that the risk measure $\mathbb{F}(Z_1, \ldots, Z_T) = \text{CVaR}_{0.3}(Z_1 + \ldots + Z_T)$ is *inconsistent* for the problem.

Remark: As discussed in Carpentier et al. (2012), it is possible that by enlarging the state space consistency may be recovered. However, the resulting dynamic programming equations would be intractable since one might end up with an infinite dimensional problem. Since we have algorithmic concerns, our definition of consistency does not allow for infinite dimensional state spaces.

3.3 Measuring risk separately

As another example, suppose we measure risk in each stage separately, i.e., we use the risk measure $\mathbb{F}(Z_1, \ldots, Z_T) = Z_1 + \rho_2(Z_2) + \ldots + \rho_T(Z_T)$ where each ρ_t is a one-period risk measure applied to period t. In this particular example we have $\mathbb{F}(cx_1, -s_2x_2, -s_3x_3) = cx_1 + \rho_2(-s_2x_2) + \rho_3(-s_3x_3)$. In particular, suppose that we use the risk measures $\rho_2 = \rho_3 = \text{CVaR}_\alpha$ for $\alpha = 1 - \epsilon$ with arbitrarily small ϵ, which as seen before is the same as using $\rho_2 = \rho_3 = \text{ess sup}$. Then, we obtain the problem

$$
\begin{aligned}
\min \;\; & cx_1 + \eta_2 + \frac{1}{1-\alpha}\mathbb{E}\left[(-s_2x_2(D_2) - \eta_2)_+\right] + \eta_3 + \frac{1}{1-\alpha}\mathbb{E}\left[(-s_3x_3(D_2, D_3) - \eta_3)_+\right] \\
\text{s.t.} \;\; & \text{the same constraints as problem (3.2),}
\end{aligned}
\tag{3.9}
$$

which yields the following optimal policy:

$$
\begin{aligned}
&\hat{x}_1 = 6, \\
&\hat{x}_2(5) = 5; \quad \hat{x}_2(10) = 5, \\
&\hat{x}_3(5,3) = 1; \quad \hat{x}_3(5,7) = 1; \quad \hat{x}_3(10,1) = 1; \quad \hat{x}_3(10,12) = 1.
\end{aligned}
\tag{3.10}
$$

Note that the subproblem consisting of the left-hand side of the tree in Figure 2 coincides with (3.5), except that $\hat{x}_1 = 6$. By solving this problem we obtain

$$\bar{x}_2^{2,5} = 3, \qquad (3.11)$$
$$\bar{x}^{2,5}(3) = 3; \quad \bar{x}^{2,5}(7) = 3.$$

The optimal value of the subproblem is -39. However, the solution inherited from $\hat{x}$, which is $\hat{x}_2(5) = 5$, $\hat{x}_3(5,3) = 1$, and $\hat{x}_3(5,7) = 1$ yields an objective value equal to -25. Thus, the solution "inherited" from solving the problem at time 1 is not optimal for the subproblem. Moreover, by conducting a sensitivity analysis for problem (3.9) it can be verified that any optimal solution of the problem at $t = 1$ must have $\hat{x}_1 = 6$ and $\hat{x}_2(5) = 5$, which leads to an inherited sub-optimal solution for the subproblem. We conclude that the risk measure $\mathbb{F}(Z_1, \ldots, Z_T) = Z_1 + \rho_2(Z_2) + \ldots + \rho_T(Z_T)$ with $T = 3$ is *inconsistent* for this problem.

The situations shown by the above examples illustrate well the concepts of consistency. In the first example with $\alpha = 1 - \epsilon$, the optimal policy (3.4) protects against the worst-case path. When we move to the subproblem, the same principle applies. When $\alpha = 0.3$ the decision is not guided by the worst case criterion anymore, and we showed that inconsistency occurs. In the second example, the optimal policy (3.10) is protecting against the worst-case outcome in each stage — i.e., $D_2 = 5$, $D_3 = 1$, which leads to $\hat{x}_1 = 6$ — but it is clear such a scenario cannot happen since those two realizations are not in the same path of the tree in Figure 2. That is, *no* optimal solution generated by the original problem will be optimal for the subproblem. These simple examples show that a risk measure for multistage stochastic programs that measures risk either based on complete paths of the scenario tree or separately per stage is inconsistent in general. So if consistency is a desired property, then such risk measures should be avoided.

3.4 Discussion on inconsistency of CVaR

We close this section by noting that the analysis of some of the above examples where $\text{CVaR}_\alpha = \operatorname{ess\,sup}$ could have been made much shorter since in reality there is no need to solve linear programs to find the best solutions. Nevertheless, we chose to present the material as done above for two reasons: first, to illustrate how consistency can be checked in a more systematic way for a general problem (in fact, even the same problem but for a smaller α, as done for the case $\alpha = 0.3$); and second, to illustrate the role of the auxiliary variables η_t. The latter is particular important because Shapiro (2009) justifies inconsistency of stage-wise risk formulations such as (3.9) based on the fact that the auxiliary variables η_t can be viewed as stage-1 variables as they are not functions of ξ_t. While that statement is certainly true for the problem at $t = 1$, we believe it does not quite explain the issues of inconsistency since in reality these variables are not "inherited" by the subproblems. For example, we cannot fix the values of the auxiliary variable η_2 when solving the subproblem at time $t = 2$, otherwise the subproblem would not calculate the CVaR correctly. Indeed, we can see that in the subproblem formulation (3.5) it is essential to define a new variable $\eta^{2,5}$ to model the conditional risk measure $\rho(\cdot|D_2 = 5)$, but even then inconsistency is observed. That is, inconsistency is a characteristic of the risk measure itself, and not of one of its particular representations.

4 Consistent risk measures

We discuss now some consistent risk measures.

4.1 Nested risk measures

We consider now *nested* risk measures of the form

$$\mathbb{F}(Z_1, \ldots, Z_T) \; = \; \rho_2 \circ \ldots \circ \rho_T(Z_1 + \ldots + Z_T), \tag{4.1}$$

where each ρ_t is a one-period conditional risk measure, defined in terms of a realization $\hat{\xi}_2, \ldots, \hat{\xi}_{t-1}$. Such nested risk measures have been object of extensive study, see for instance Shapiro (2009), Ruszczyński (2010), Philpott and de Matos (2012) and Philpott et al. (2013). Nested risk measures are somewhat equated with their respective definitions of consistency. We show next that nested risk measures satisfy our notion of consistency.

Theorem 4.1. *Consider the risk measure $\mathbb{F}$ given by $\mathbb{F}(Z_1, \ldots, Z_T) = \rho_2 \circ \rho_3^{\xi_{[2]}} \circ \ldots \circ \rho_T^{\xi_{[T-1]}}(Z_1 + \ldots + Z_T)$, where the notation indicates that each $\rho_t^{\xi_{[t-1]}}$ is a one-period conditional risk measure, defined in terms of the history $\xi_2, \ldots, \xi_{t-1}$. Suppose also that each $\rho_t^{\xi_{[t-1]}}$ is translation-invariant and monotone, i.e.,*

(i) $\rho_t^{\xi_{[t-1]}}(Z + X) = Z + \rho_t^{\xi_{[t-1]}}(X)$ *whenever Z is measurable with respect to the sigma-algebra generated by $\xi_1, \ldots, \xi_{t-1}$;*

(ii) $\rho_t^{\xi_{[t-1]}}(X) \leq \rho_t^{\xi_{[t-1]}}(Y)$ *whenever $X \leq Y$ w.p.1.*

Then, $\mathbb{F}$ is consistent for problem $(\mathcal{P})$.

Proof: Observe initially that, given x, each term $f_t(x_t, \xi_t)$ is constant given the history $\xi_1, \ldots, \xi_t$, so it "moves out" of the argument of the conditional risk measure $\rho_{t+1}^{\xi_{[t]}}$ due to the translation-invariant property. It follows that

$$\min_{x : x_t \in \mathcal{X}_t\left(x_{[t-1]}, \xi_{[t]}\right)} \rho_2 \circ \ldots \circ \rho_T^{\xi_{[T-1]}}\Big(f_1(x_1, \xi_1) + \ldots + f_T(x_T, \xi_T)\Big) \tag{4.2}$$

$$= \min_{x_1 \in \mathcal{X}_1} \min_{x_2 \in \mathcal{X}_2\left(x_1, \xi_{[2]}\right)} \ldots \min_{x_T \in \mathcal{X}_T\left(x_{[T-1]}, \xi_{[T]}\right)} f_1(x_1, \xi_1) + \rho_2\Bigg(f_2(x_2, \xi_2) + \rho_3^{\xi_{[2]}}\Bigg(f_3(x_3, \xi_3) + + \rho_4^{\xi_{[3]}}\Big(\ldots + \rho_T^{\xi_{[T-1]}}\big(f_T(x_T, \xi_T)\big)\Big)\ldots\Bigg)\Bigg)$$

$$= \min_{x_1 \in \mathcal{X}_1} \Bigg\{ f_1(x_1, \xi_1) + \rho_2\Bigg(\min_{x_2 \in \mathcal{X}_2\left(x_1, \xi_{[2]}\right)} \Big\{ f_2(x_2, \xi_2) + \rho_3^{\xi_{[2]}}\Bigg(\min_{x_3 \in \mathcal{X}_3\left(x_{[2]}, \xi_{[3]}\right)} \Big\{ f_3(x_3, \xi_3) + \ldots \ldots + \rho_T^{\xi_{[T-1]}}\Big(\min_{x_T \in \mathcal{X}_T\left(x_{[T-1]}, \xi_{[T]}\right)} f_T(x_T, \xi_T)\Big) \ldots \Big\}\Bigg)\Bigg\}. \tag{4.3}$$

Note that the second equality in the above development (i.e., the interchange between $\rho_t^{\xi_{[t-1]}}$ and min) follows from the assumed monotonicity of each $\rho_t^{\xi_{[t-1]}}$. Thus, we see that solving the original problem ($\mathcal{P}$) is equivalent to solving (4.3).

Let t be an arbitrary time period such that $1 < t \leq T$. Note that we can rewrite (4.3) as

$$\min_{x: x_\tau \in \mathcal{X}_\tau\left(x_{[\tau-1]}, \xi_{[\tau]}\right),\ \tau=1,\ldots,t-1} \Bigg\{ f_1(x_1,\xi_1) + \rho_2\Bigg(f_2(x_2,\xi_2) + \rho_3^{\xi_{[2]}}\Bigg(f_3(x_3,\xi_3) + \rho_4^{\xi_{[3]}}\Big(\ldots$$
$$\ldots f_{t-1}(x_{t-1},\xi_{t-1}) + \rho_t^{\xi_{[t-1]}}\Big(\min_{x: x_\tau \in \mathcal{X}_\tau\left(x_{[\tau-1]}, \xi_{[\tau]}\right), \tau=t,\ldots,T} \Big\{ f_t(x_t,\xi_t) + \rho_{t+1}^{\xi_{[t]}}\Big(\ldots$$
$$\ldots f_{T-1}(x_{T-1},\xi_{T-1}) + \rho_T^{\xi_{[T-1]}}\big(f_T(x_T,\xi_T)\big)\ldots\Big)\Big\}\ldots\Big)\Bigg)\Bigg)\Bigg\} \tag{4.4}$$

Let $\hat{x} := [\hat{x}_\tau : \tau = 1,\ldots,T]$ be an optimal solution of problem (4.4) (and hence of (4.2)). Note that, by construction, the solution inherited from $\hat{x}$ is an optimal solution of the inner minimization problem in (4.4).

We now show that the solution inherited from $\hat{x}$ is optimal for the subproblem ($\mathcal{P}_t$). Since $\mathbb{F}$ is assumed to be a nested risk measure, we have that

$$\begin{aligned}
\mathbb{F}^{\xi_{[t]}}(Z_1,\ldots,Z_T) &= \rho_{t+1}^{\xi_{[t]}} \circ \rho_{t+2}^{\xi_{[t]},\xi_{t+1}} \ldots \circ \rho_T^{\xi_{[t]},\xi_{t+1},\ldots,\xi_{T-1}}(Z_1+\ldots+Z_T) \\
&= Z_1+\ldots+Z_t + \rho_{t+1}^{\xi_{[t]}}\Big(Z_{t+1} + \rho_{t+2}^{\xi_{[t]},\xi_{t+1}}\ \big(\ldots Z_{T-1} + \rho_T^{\xi_{[t]},\xi_{t+1},\ldots,\xi_{T-1}}(Z_T)\big)\Big)
\end{aligned}$$

so problem ($\mathcal{P}_t$) becomes

$$\min_{x: x_\tau \in \mathcal{X}_\tau\left(x_{[\tau-1]}, \xi_{[\tau]}\right), \tau=t,\ldots,T} \mathbb{F}^{\xi_{[t]}}\Big(f_1(x_1,\xi_1),\ldots,f_T(x_T,\xi_T)\Big)$$
$$= f_1(x_1,\xi_1)+\ldots+f_{t-1}(x_{t-1},\xi_{t-1}) + \min_{x: x_\tau \in \mathcal{X}_\tau\left(x_{[\tau-1]}, \xi_{[\tau]}\right), \tau=t,\ldots,T} \Big\{ f_t(x_t,\xi_t) + \rho_{t+1}^{\xi_{[t]}}\Big(\ldots$$
$$\ldots f_{T-1}(x_{T-1},\xi_{T-1}) + \rho_T^{\xi_{[T-1]}}\big(f_T(x_T,\xi_T)\big)\ldots\Big)\Big\}. \tag{4.5}$$

It becomes clear that the inner minimization problem in (4.4) coincides with the minimization problem in (4.5). It follows that the solution inherited from $\hat{x}$ is optimal for the subproblem ($\mathcal{P}_t$), so the inherited optimality property holds. The above argument is valid regardless of the problem instance, so we conclude that $\mathbb{F}$ is consistent. □

It is worthwhile noticing that problem (4.3) essentially corresponds to a Bellman formulation of the original problem, which allows for the use of recursive algorithms commonly used in dynamic programming. Thus, we see that nested risk measures are consistent and naturally lead to Bellman formulations. Of course, many authors have discussed the equivalence between consistency and recursive formulations; the value of Theorem 4.1 lies in showing that our proposed notion of consistency — which, as mentioned before, attempts to formalize some of the notions found in the literature — is also implied by the Bellman formulation of nested risk measures.

A particular case of interest is the risk neutral measure $\mathbb{F}(Z_1, \ldots, Z_T) = \mathbb{E}[Z_1 + \ldots + Z_T]$. By the well-known "tower property", we have that $\mathbb{E} = \mathbb{E} \circ \mathbb{E}^{\xi_{[2]}} \circ \ldots \circ \mathbb{E}^{\xi_{[T-1]}}$, where $\mathbb{E}^{\xi_{[t]}}$ (written this way to agree with our notation) is the conditional expectation $\mathbb{E}[\cdot \mid \xi_{[t]}]$. Another case of interest is the worst-case risk measure $\mathbb{F}(Z_1, \ldots, Z_T) = \text{ess sup } (Z_1 + \ldots + Z_T)$. Since we can also write that risk measure in similar composite form, we see that Theorem 4.1 provides a general framework that includes the first example of Section 3. The conclusion of the corollary below is immediate from Theorem 4.1.

Corollary 4.2. *The risk neutral measure* $\mathbb{F}(Z_1, \ldots, Z_T) = \mathbb{E}[Z_1 + \ldots + Z_T]$ *and the worst-case risk measure* $\mathbb{F}(Z_1, \ldots, Z_T) = \text{ess sup } (Z_1 + \ldots + Z_T)$ *are consistent for problems of the form* $(\mathcal{P})$.

It is worthwhile mentioning here that Shapiro (2012b) shows that the *only* consistent law invariant risk measures are the ones in Corollary (4.2). While such a result is interesting in its own, we believe it may lead one to think that it is dangerous to use other consistent risk measures as they are not law invariant. It should be stressed however that this limitation is a direct consequence of the notion of law invariance used in Shapiro (2012b), which is somewhat strict as discussed in Remark 2 earlier in this paper. Once that notion is relaxed (as in our definition), nested risk measures of the form (4.1) become law invariant as well, provided each ρ_i is law invariant.

5 A period-wise composite measure

5.1 Expected conditional risk measures (ECRMs)

Consider now the class of multi-period risk measures $\mathbb{F}$ defined as follows:

$$\mathbb{F}(Z_1, \ldots, Z_T) = Z_1 + \rho_2(Z_2) + \mathbb{E}_{\xi_{[2]}}\left[\rho_3^{\xi_{[2]}}(Z_3)\right] + \mathbb{E}_{\xi_{[3]}}\left[\rho_4^{\xi_{[3]}}(Z_4)\right] + \ldots + \mathbb{E}_{\xi_{[T-1]}}\left[\rho_T^{\xi_{[T-1]}}(Z_T)\right], \tag{5.1}$$

where the subscript in $\mathbb{E}$ indicates that the expectation is with respect to the corresponding variables. We shall call multi-period risk measures defined this way *expected conditional risk measures* (ECRMs). This class of risk measures includes the mCVaR risk measure that was introduced in Pflug and Ruszczyński (2005), which is defined as (5.1) for the case where $\rho_t = \text{CVaR}_{\alpha_t}$. Pflug (2006) provides a more extensive study of mCVaR and proves some fundamental properties of that risk measure. We show now that any $\mathbb{F}$ defined as in (5.1) is consistent, provided that each $\rho_t^{\xi_{[t-1]}}$ satisfies some basic properties that automatically hold, for example, for coherent risk measures.

Theorem 5.1. *Consider the expected conditional risk measure* $\mathbb{F}$ *given by* (5.1)*, where each* $\rho_t^{\xi_{[t-1]}}$ *is a translation-invariant and monotone conditional risk measure, in the sense of conditions (i) and (ii) spelled out in Theorem 4.1 . Suppose also that the random vectors* ξ_t *have* finite support*. Then,* $\mathbb{F}$ *is consistent for problem* $(\mathcal{P})$.

Proof: Observe initially that, using the "tower property" of expectations, we can rewrite $\mathbb{F}$ as

$$\mathbb{F}(Z_1,\ldots,Z_T) = Z_1 + \rho_2(Z_2) + \mathbb{E}_{\xi_2}\left[\rho_3^{\xi_{[2]}}(Z_3) + \mathbb{E}_{\xi_3}^{\xi_{[2]}}\left[\rho_4^{\xi_{[3]}}(Z_4) + \ldots + \mathbb{E}_{\xi_{T-1}}^{\xi_{[T-2]}}\left[\rho_T^{\xi_{[T-1]}}(Z_T)\right]\ldots\right]\right]. \tag{5.2}$$

By the translation-invariant property of each $\rho_t^{\xi_{[t-1]}}$, we can re-write the above equation as

$$\mathbb{F}(Z_1,\ldots,Z_T) = Z_1 + \rho_2\left(Z_2 + \mathbb{E}_{\xi_2}\circ\rho_3^{\xi_{[2]}}\left(Z_3 + \mathbb{E}_{\xi_3}^{\xi_{[2]}}\circ\rho_4^{\xi_{[3]}}\left(Z_4 + \ldots + \mathbb{E}_{\xi_{T-1}}^{\xi_{[T-2]}}\circ\rho_T^{\xi_{[T-1]}}\left(Z_T\right)\ldots\right)\right)\right). \tag{5.3}$$

To simplify the notation, define now

$$\tilde{\rho}_t^{\xi_{[t-2]}} \;:=\; \mathbb{E}_{\xi_{t-1}}^{\xi_{[t-2]}}\circ\rho_t^{\xi_{[t-1]}}.$$

Then, we can write (5.3) as

$$\mathbb{F}(Z_1,\ldots,Z_T) = Z_1 + \rho_2\left(Z_2 + \;\tilde{\rho}_3^{\xi_{[1]}}\left(Z_3 + \tilde{\rho}_4^{\xi_{[2]}}\left(Z_4 + \ldots + \tilde{\rho}_T^{\xi_{[T-2]}}\left(Z_T\right)\ldots\right)\right)\right). \tag{5.4}$$

Note that, despite the apparent nested form of (5.4), the resulting formulation does not fit the framework of Theorem 4.1 since $\tilde{\rho}_t^{\xi_{[t-2]}}$ is measurable with respect to the σ-algebra generated by $\xi_{[t-2]}$ instead of that generated by $\xi_{[t-1]}$ — so we cannot write, for example, $\tilde{\rho}_t^{\xi_{[t-2]}}(Z_{t-1}+Z_t) = Z_{t-1} + \tilde{\rho}_t^{\xi_{[t-2]}}(Z_t)$. Nevertheless, we can follow similar steps as those in the proof of Theorem 4.1 and use the monotonicity and translation-invariant properties of $\tilde{\rho}_t^{\xi_{[t-2]}}$ — which follow from the fact that both $\mathbb{E}$ and $\rho_t^{\xi_{[t-1]}}$ have such properties.

To proceed, let t be an arbitrary time period such that $1 < t \le T$. Then, we have that

$$\min_{x:x_\tau\in\mathcal{X}_\tau\left(x_{[\tau-1]},\xi_{[\tau]}\right),\tau=1,\ldots,T} \mathbb{F}\Big(f_1(x_1,\xi_1)+\ldots+f_T(x_T,\xi_T)\Big) \tag{5.5}$$

$$\begin{aligned}&\min_{x:x_\tau\in\mathcal{X}_\tau\left(x_{[\tau-1]},\xi_{[\tau]}\right),\tau=1,\ldots,t-1}\Bigg\{f_1(x_1,\xi_1) + \rho_2\Bigg(f_2(x_2,\xi_2) + \tilde{\rho}_3^{\xi_{[1]}}\Bigg(f_3(x_3,\xi_3) + \tilde{\rho}_4^{\xi_{[2]}}\Bigg(\ldots\\ &\quad\ldots f_{t-1}(x_{t-1},\xi_{t-1}) + \tilde{\rho}_t^{\xi_{[t-2]}}\Bigg(\min_{x:x_\tau\in\mathcal{X}_\tau\left(x_{[\tau-1]},\xi_{[\tau]}\right),\tau=t,\ldots,T}\Bigg\{f_t(x_t,\xi_t) + \tilde{\rho}_{t+1}^{\xi_{[t-1]}}\Bigg(\ldots\\ &\quad\ldots f_{T-1}(x_{T-1},\xi_{T-1}) + \tilde{\rho}_T^{\xi_{[T-2]}}\big(f_T(x_T,\xi_T)\big)\ldots\Bigg)\Bigg\}\ldots\Bigg)\Bigg)\Bigg\}\end{aligned} \tag{5.6}$$

Let $\hat{x} := [\hat{x}_\tau : \tau = 1,\ldots,T]$ be an optimal solution of problem (5.6) (and hence of (5.5)). Note that, by construction, the solution inherited from $\hat{x}$ is an optimal solution of the inner minimization problem in (5.6).

We now show that the solution inherited from $\hat{x}$ is optimal for the subproblem $(\mathcal{P}_t)$. Note that when information up to time t is available we have that $\mathbb{E}_{\xi_t} \circ \rho_{t+1}^{\xi_{[t]}} = \rho_{t+1}^{\xi_{[t]}}$. It follows that

$$\mathbb{F}^{\xi_{[t]}}(Z_1, \dots, Z_T) \;=\; Z_1 + \dots + Z_t + \rho_{t+1}^{\xi_{[t]}}\left(Z_{t+1} + \tilde{\rho}_{t+2}^{\xi_{[t]}}\left(Z_{t+2} + \dots + \tilde{\rho}_T^{\xi_{[T-2]}}(Z_T) \dots\right)\right)$$

so problem $(\mathcal{P}_t)$ becomes

$$\begin{aligned}
&\min_{x: x_\tau \in \mathcal{X}_\tau\left(x_{[\tau-1]}, \xi_{[\tau]}\right), \tau=t,\dots,T} \mathbb{F}^{\xi_{[t]}}\left(f_1(x_1, \xi_1), \dots, f_T(x_T, \xi_T)\right) \\
&= f_1(x_1, \xi_1) + \dots + f_{t-1}(x_{t-1}, \xi_{t-1}) + \min_{x: x_\tau \in \mathcal{X}_\tau\left(x_{[\tau-1]}, \xi_{[\tau]}\right), \tau=t,\dots,T} \Big\{ f_t(x_t, \xi_t) + \\
&\qquad \rho_{t+1}^{\xi_{[t]}}\left(f_{t+1}(x_{t+1}, \xi_{t+1}) + \tilde{\rho}_{t+2}^{\xi_{[t]}}\left(\dots\dots f_{T-1}(x_{T-1}, \xi_{T-1}) + \tilde{\rho}_T^{\xi_{[T-2]}}\left(f_T(x_T, \xi_T)\right) \dots\right)\right)\Big\}.
\end{aligned} \tag{5.7}$$

It is important to observe that, unlike the case of the proof of Theorem 4.1, the inner minimization problem in (5.6) *does not* coincide with the minimization problem in (5.7), since the former optimizes the conditional risk measure $\tilde{\rho}_{t+1}^{\xi_{[t-1]}} = \mathbb{E}_{\xi_t}^{\xi_{[t-1]}} \circ \rho_{t+1}^{\xi_{[t]}}$ whereas the latter optimizes the conditional risk measure $\rho_{t+1}^{\xi_{[t]}}$ (the remaining terms are identical). We will show, however, that the solution inherited from $\hat{x}$ is an optimal solution of the minimization problem in (5.7). Indeed, suppose that this is not the case, i.e., there exists an optimal solution $\tilde{x} := [\tilde{x}_\tau \,:\, \tau = t, \dots, T]$ of the minimization problem in (5.7) (given $\hat{x}_\tau$, $\tau = 1, \dots, t-1$) such that $\tilde{x}$ is strictly better than $\hat{x}$ for some realization $\hat{\xi}_{[t]}$ of $\xi_{[t]}$, i.e.

$$\begin{aligned}
&f_t(\tilde{x}_t, \hat{\xi}_t) + \\
&\qquad \rho_{t+1}^{\hat{\xi}_{[t]}}\left(f_{t+1}(\tilde{x}_{t+1}, \xi_{t+1}) + \tilde{\rho}_{t+2}^{\hat{\xi}_{[t]}}\left(\dots + \tilde{\rho}_T^{\hat{\xi}_{[t]}, \xi_{t+1}, \dots, \xi_{T-2}}\left(f_T(\tilde{x}_T, \xi_T)\right) \dots\right)\right) \\
&< f_t(\hat{x}_t, \hat{\xi}_t) + \\
&\qquad \rho_{t+1}^{\hat{\xi}_{[t]}}\left(f_{t+1}(\hat{x}_{t+1}, \xi_{t+1}) + \tilde{\rho}_{t+2}^{\hat{\xi}_{[t]}}\left(\dots + \tilde{\rho}_T^{\hat{\xi}_{[t]}, \xi_{t+1}, \dots, \xi_{T-2}}\left(f_T(\hat{x}_T, \xi_T)\right) \dots\right)\right).
\end{aligned} \tag{5.8}$$

For any other realization of $\xi_{[t]}$, of course, the above inequality holds with $\leq$ in place of $<$, since $\tilde{x}$ is optimal. By computing the conditional expectation $\mathbb{E}_{\xi_t}^{\hat{\xi}_{[t-1]}}$ on both sides of the inequality we obtain

$$\begin{aligned}
&\mathbb{E}_{\xi_t}^{\hat{\xi}_{[t-1]}}\left[f_t(\tilde{x}_t, \xi_t)\right] + \\
&\qquad \mathbb{E}_{\xi_t}^{\hat{\xi}_{[t-1]}} \circ \rho_{t+1}^{\hat{\xi}_{[t-1]}, \xi_t}\left(f_{t+1}(\tilde{x}_{t+1}, \xi_{t+1}) + \tilde{\rho}_{t+2}^{\hat{\xi}_{[t-1]}, \xi_t}\left(\dots + \tilde{\rho}_T^{\hat{\xi}_{[t-1]}, \xi_t, \dots, \xi_{T-2}}\left(f_T(\tilde{x}_T, \xi_T)\right) \dots\right)\right) \\
&< \mathbb{E}_{\xi_t}^{\hat{\xi}_{[t-1]}}\left[f_t(\hat{x}_t, \xi_t)\right] + \\
&\qquad \mathbb{E}_{\xi_t}^{\hat{\xi}_{[t-1]}} \circ \rho_{t+1}^{\hat{\xi}_{[t-1]}, \xi_t}\left(f_{t+1}(\hat{x}_{t+1}, \xi_{t+1}) + \tilde{\rho}_{t+2}^{\hat{\xi}_{[t-1]}, \xi_t}\left(\dots + \tilde{\rho}_T^{\hat{\xi}_{[t-1]}, \xi_t, \dots, \xi_{T-2}}\left(f_T(\hat{x}_T, \xi_T)\right) \dots\right)\right).
\end{aligned} \tag{5.9}$$

In the above inequality, the assumption of finite support of each ξ_t was essential — this is what allows us to write the strict inequality in (5.9) as a consequence of the strict inequality in (5.8).

On the other hand, since $\hat{x}$ is optimal for the inner minimization problem in (5.6), it follows that for any realization $\hat{\xi}_{[t]}$ of $\xi_{[t]}$ we have

$$\begin{aligned}
& f_t(\hat{x}_t, \hat{\xi}_t) + \\
& \qquad \mathbb{E}_{\xi_t}^{\hat{\xi}_{[t-1]}} \circ \rho_{t+1}^{\hat{\xi}_{[t-1]},\xi_t} \bigg(f_{t+1}(\hat{x}_{t+1}, \xi_{t+1}) + \tilde{\rho}_{t+2}^{\hat{\xi}_{[t-1]},\xi_t} \Big(\ldots + \tilde{\rho}_T^{\hat{\xi}_{[t-1]},\xi_t,\ldots,\xi_{T-2}} \big(f_T(\hat{x}_T, \xi_T) \big) \ldots \Big) \bigg) \\
\leq\ & f_t(\tilde{x}_t, \hat{\xi}_t) + \\
& \qquad \mathbb{E}_{\xi_t}^{\hat{\xi}_{[t-1]}} \circ \rho_{t+1}^{\hat{\xi}_{[t-1]},\xi_t} \bigg(f_{t+1}(\tilde{x}_{t+1}, \xi_{t+1}) + \tilde{\rho}_{t+2}^{\hat{\xi}_{[t-1]},\xi_t} \Big(\ldots + \tilde{\rho}_T^{\hat{\xi}_{[t-1]},\xi_t,\ldots,\xi_{T-2}} \big(f_T(\tilde{x}_T, \xi_T) \big) \ldots \Big) \bigg)
\end{aligned}$$

and hence

$$\begin{aligned}
& \mathbb{E}_{\xi_t}^{\hat{\xi}_{[t-1]}} \big[f_t(\hat{x}_t, \xi_t) \big] + \\
& \qquad \mathbb{E}_{\xi_t}^{\hat{\xi}_{[t-1]}} \circ \rho_{t+1}^{\hat{\xi}_{[t-1]},\xi_t} \bigg(f_{t+1}(\hat{x}_{t+1}, \xi_{t+1}) + \tilde{\rho}_{t+2}^{\hat{\xi}_{[t-1]},\xi_t} \Big(\ldots + \tilde{\rho}_T^{\hat{\xi}_{[t-1]},\xi_t,\ldots,\xi_{T-2}} \big(f_T(\hat{x}_T, \xi_T) \big) \ldots \Big) \bigg) \\
\leq\ & \mathbb{E}_{\xi_t}^{\hat{\xi}_{[t-1]}} \big[f_t(\tilde{x}_t, \xi_t) \big] + \\
& \qquad \mathbb{E}_{\xi_t}^{\hat{\xi}_{[t-1]}} \circ \rho_{t+1}^{\hat{\xi}_{[t-1]},\xi_t} \bigg(f_{t+1}(\tilde{x}_{t+1}, \xi_{t+1}) + \tilde{\rho}_{t+2}^{\hat{\xi}_{[t-1]},\xi_t} \Big(\ldots + \tilde{\rho}_T^{\hat{\xi}_{[t-1]},\xi_t,\ldots,\xi_{T-2}} \big(f_T(\tilde{x}_T, \xi_T) \big) \ldots \Big) \bigg).
\end{aligned} \tag{5.10}$$

By putting together inequalities (5.9) and (5.10), we see that we reach a contradiction. It follows that the solution inherited from $\hat{x}$ is optimal for the subproblem $(\mathcal{P}_t)$, so the inherited optimality property holds. The above argument is valid regardless of the problem instance, so we conclude that $\mathbb{F}$ is consistent. □

5.2 Explicit $\mathbb{E}$-CVaR formulation

As mentioned above, a particular case of ECRMs defined in (5.1) is to use $\rho_t = \text{CVaR}_{\alpha_t}$. We shall denote the resulting ECRM by $\mathbb{E}$-CVaR. To the best of our knowledge, this measure was introduced in Pflug and Ruszczyński (2005) in a pension fund problem with incoming cash flows. Pflug (2006) refers to this measure as the multi-period average value-at-risk.

Theorem 5.1 shows the consistency of $\mathbb{E}$-CVaR. As we are assuming that consistency is a desirable property in applications, Theorem 5.1 establishes the $\mathbb{E}$-CVaR as a valid alternative to the nested risk measure defined in (4.1). In fact the $\mathbb{E}$-CVaR has additional properties that makes it an appealing alternative to measure risk in multistage settings. As we show below, any risk-averse MSSP defined with $\mathbb{E}$-CVaR can be written as a risk-neutral model for a modified problem that has some additional variables. Thus, moderately sized problems can be efficiently solved to optimality by using commercial solvers.

For large scale problems, existing algorithms can be readily adapted to solve the $\mathbb{E}$-CVaR case. For example the popular SDDP algorithm, developed by Pereira and Pinto (1991), can in principle be applied directly. The only changes would be the addition of some variables and a redefinition of the vectors and matrices of the problem. This is not true for the nested CVaR formulation: in Philpott and de Matos (2012) the authors show that the Bellman equations for the nested case in time t contain the recursive function at time $t+1$ in the constraints, which in general is not given in explicit form. Approximation by cutting planes can be performed, but the procedure to compute the upper bound of the optimal value needs significant changes; see Philpott and de Matos (2012), Kozmık and Morton (2013) and Philpott et al. (2013).

We will now show the explicit Bellman equations for the $\mathbb{E}$-CVaR case. From equation (5.2) in Theorem 5.1, we can write an optimization formulation for the $\mathbb{E}$-CVaR as follows :

$$\begin{aligned} \min_{x_1,\dots,x_T} & f_1(x_1,\xi_1) + \mathrm{CVaR}_{\alpha_2}(f_2(x_2,\xi_2)) + \mathbb{E}_{\xi_2}\Big[\mathrm{CVaR}_{\alpha_3}^{\xi_{[2]}}(f_3(x_3,\xi_3)) + \mathbb{E}_{\xi_3}\Big[\mathrm{CVaR}_{\alpha_4}^{\xi_{[3]}}(f_4(x_4,\xi_4)) + \dots \\ & + \mathbb{E}_{\xi_{T-1}}\left[\mathrm{CVaR}_{\alpha_T}^{\xi_{[T-1]}}(f_T(x_T,\xi_T)) \,|\, \xi_{[T-2]}\right] \dots \,|\, \xi_{[2]}\Big]\Big], \\ & \text{s.t.} \quad x_t \in \mathcal{X}_t\left(x_{[t-1]},\xi_{[t]}\right), \quad t=1,\dots,T. \end{aligned}$$

Using the CVaR representation described in Rockafellar and Uryasev (2000) we have

$$\begin{aligned} \min_{x_1,\dots,x_T} f_1(x_1,\xi_1) + \min_{\eta_2}\eta_2 + \frac{1}{1-\alpha_2}\mathbb{E}_{\xi_2}\left[(f_2(x_2,\xi_2)-\eta_2)_+\right] + \\ \mathbb{E}_{\xi_2}\Bigg[\min_{\eta_3}\eta_3 + \frac{1}{1-\alpha_3}\mathbb{E}_{\xi_3}\left[(f_3(x_3,\xi_3)-\eta_3)_+ \,|\, \xi_{[2]}\right] \\ +\mathbb{E}_{\xi_3}\Bigg[\min_{\eta_4}\eta_4 + \frac{1}{1-\alpha_4}\mathbb{E}_{\xi_4}\left[(f_4(x_4,\xi_4)-\eta_4)_+ \,|\, \xi_{[3]}\right] + \dots + \\ \mathbb{E}_{\xi_{T-1}}\left[\min_{\eta_T}\eta_T + \frac{1}{1-\alpha_T}\mathbb{E}_{\xi_T}\left[(f_T(x_T,\xi_T)-\eta_T)_+ \,|\, \xi_{[T-1]}\right] \dots \,\Big|\, \xi_{[2]}\right]\Bigg]\Bigg], \\ \text{s.t.} \quad x_t \in \mathcal{X}_t\left(x_{[t-1]},\xi_{[t]}\right), \quad t=1,\dots,T. \end{aligned}$$

Note that the auxiliary variables $\{\eta_t\}$ are such that η_{t+1} is a function of $\xi_1,\dots,\xi_t$, i.e., η_{t+1} is a "t-stage variable" just as x_t is. By interchanging the minimum and the expected value and merging the minimum on x_t with the minimum in η_{t+1} we can write Bellman equations for $t=1,\dots,T-1$ in the following way:

$$Q_t(x_{t-1},\xi_{[t]},\eta_t) = \min_{\eta_{t+1},x_t}\frac{1}{1-\alpha_t}(f_t(x_t,\xi_t)-\eta_t)_+ + \eta_{t+1} + \mathbb{E}_{\xi_{t+1}}\left[Q_{t+1}(x_t,\xi_{t+1},\eta_{t+1}) \,|\, \xi_{[t]}\right]. \tag{5.11}$$

For the last period we have

$$Q_T(x_{T-1},\xi_T,\eta_T) = \min_{x_T}\frac{1}{1-\alpha_T}(f_T(x_T,\xi_T)-\eta_T)_+. \tag{5.12}$$

The piecewise linear function in (5.11) can be linearized by the addition of a positive variable u_t such and two constraints for each $t=2,\dots,T$:

$$u_t \geq 0, \qquad u_t \geq f_t(x_t,\xi_t) - \eta_t. \tag{5.13}$$

We see then that the $\mathbb{E}$-CVaR formulation corresponds to a risk-neutral problem that modifies the original problem by adding two new variables in stages $t = 2, \ldots, T$ (u_t and η_t), and two additional constraints (5.13) as well. As mentioned before, this brings significant computational advantages, as standard algorithms for risk-neutral problems can be used in this setting.

6 Application: an ALM model for pension funds

We illustrate the application of the $\mathbb{E}$-CVaR to a pension fund problem described in Haneveld et al. (2010). The problem is a defined benefit model in which a large Dutch company owns the fund and acts as its sponsor. The company has to maintain the ratio between assets and liabilities—the *funding ratio*— above some pre-specified threshold at every time period. To achieve this goal there are three sources of income: the returns from its asset portfolio, regular contributions made by fund participants and remedial contributions performed by the company, which are essentially money injections intended to keep the fund solvent. It is worthwhile mentioning that Kilianova and Pflug (2009) also study a risk-averse multistage pension fund problem under the $\mathbb{E}$-CVaR risk measure; the fundamental differences between our problem and theirs are that (i) the problem in Kilianova and Pflug (2009) is modeled from the viewpoint of the pension fund participant, whereas ours is from the perspective of the fund manager, and (ii) the resulting optimization problem is solved as a large linear program in Kilianova and Pflug (2009); in contrast, we use this example to demonstrate the usefulness of writing the equivalent risk-neutral formulation discussed in Section 5.2, and in fact we solve the problem using standard software that implements decomposition algorithms for expectation-based MSSPs.

6.1 Problem description

The fund has three decisions to make in every stage t: the allocation strategy of its assets among the available classes i, denoted by $X_{i,t}$, the value of the contribution rate c_t of the participants of the fund and the remedial contribution Z_t. Randomness is present through the returns of the assets r_{it}, the total wages of active participants W_t, the benefit payments P_t to the participants and the liabilities L_t. In words, the constraints for every scenario are

- The total value of the assets at time t is equal to the yield of the investments at time t, plus the participants' payments and the remedial contribution minus the payments due to retirement and others.
- The value of each asset i at time $t + 1$ is the value of asset i in period t times the interests minus transaction costs due to buying and selling.
- All assets should be allocated.
- The funding ratio has to be greater or equal than a fixed threshold at every time period.

The model defines lower and upper bounds on assets classes and on the contribution rate c_t. Non-negativity of all variables completes the constraints. In the original paper the authors

α	stocks	bonds	real estate	cash	c_1	Z_1
(.95,0,0)	4948	8460	3085	0	.21	0
(0,0,.95)	7656	4952	3899	0	.21	0
(0,.95,0)	8475	4952	3080	0	.21	0
(.95,.95,.95)	7656	5470	3383	0	.21	0
(0,0,0)	7427	4951	4126	0	.13	0
(.5,.5,.5)	5777	6594	4124	0	.21	0

Table 1: First-stage solutions.

do not include any risk measure in the objective function: the objective is to minimize the expected value of the sum, for all stages, of remedial contributions and contribution rates for the participants. A full description of the model can be found in Haneveld et al. (2010).

6.2 Numerical results

We implemented the experiment using the program SLP-IOR (Kall and Mayer, 1996). The software has an interface in which one can enter a risk-neutral multistage stochastic program per stage, specifying the decision variables and the matrices as well as the scenario tree structure. A few specialized decomposition algorithms for MSSPs are implemented in SLP-IOR. Since Bellman equations for the $\mathbb{E}$-CVaR can be written using the expected value, as shown in (5.11) and (5.12), we represented the $\mathbb{E}$-CVaR as a risk neutral problem with extra variables η_t for the CVaR and u_t to linearize the max function. We see here the advantage of being able to write the problem as a risk-neutral one, a feature that is not present in other multi-period risk-averse formulations.

Using data generated according to the distributions described on page 57 of Haneveld et al. (2010), we consider a four stage problem whose scenario tree has ten bifurcations per node, for a total of 1000 scenarios. There are four types of assets, from riskier (with higher returns) to safer (lower returns): stocks, real state, bonds and cash. We can control risk at the second, third and fourth stages by using parameters α_2, α_3 and α_4. For example, the choice $(0, 0, 0)$ corresponds to the risk-neutral case, whereas $(.95, 0, 0)$ represents the situation where the decision maker is very risk-averse regarding the outcome of stage 2 but it is risk-neutral regarding stages 3 and 4. First-stage solutions are shown in Table 1.

Several comments are in order. First we observe that the remedial contribution was not necessary in the first stage for any of the six risk profiles considered. The risk neutral solution (5th row) has a contribution rate of only 13%, while in all of the other cases the contribution is at its maximum of 21%. In the first row we have the most conservative solution: only 30% of the initial wealth is invested in stocks, as opposed to an average 48% among rows 2-5. In addition, more than 50% of the initial wealth is invested in bonds.

Those numbers support the fact that the choice (.95,0,0) strongly protects the first-stage solution against volatility. The (.95,.95,.95) configuration should also offer a similar protection, but it is equally concerned with second and third stages and, as a consequence, the percentage invested in stocks and real estate is much higher than in the (.95,0,0) case. Nevertheless, among the middle four rows, this it is the one with smallest percentage of investment in stocks plus real estate. The last row is an intermediate case: the amount invested in risky assets is halfway between rows 1 and 2-5.

6.3 Beyond the first-stage solution: stochastic dominance

The vast majority of authors that report multistage experiments only show the results of the first stage. We believe there are two valid reasons for that: first, the solution for stages other than the first will probably not be implemented since the realization of the uncertainties will be different, with probability one, from any realization that is represented in the scenario tree. Second, many practical applications (energy, finance among others) use a rolling horizon procedure, that is, after implementing the first-stage solution and observing the realization of the uncertainties, the problem is re-solved and the new first-stage solution obtained is implemented, and so forth.

Despite the fact that we agree with both arguments presented, we believe that showing the solutions of the further stages increases the understanding of the problem. In fact, it helps the decision maker to understand the consequences of the first-stage solution that will be implemented. For example, a look at Table 1 would suggest that the risk neutral solution dominates all other solutions: it requires no contribution from the fund sponsor and has the smallest contribution rate of all solutions! We will use first and second order stochastic dominance (see, e.g., Müller and Stoyan 2002 for definitions) to infer about the quality of second- and third-stage solutions.

Figures 3 and 4 show parts of the second and third stage solutions that will allow to us to understand the implications of adopting one of the first-stage solutions displayed in Table 1. In each figure we plot four empirical cumulative distributions functions constructed based on the solutions of second (Table 3) and third (Table 4) stages, for each realization in the tree: investment in risky assets (top left, assets and real estate), safer assets (top right, bonds and cash), the remedial contribution and the contribution rate.

The first observation is that, for the remedial contribution, the solution (0,0,0) first-order stochastically dominates all other solutions, both in the second and third stages. That is a clearly undesirable property for a solution: it means that for every value x the probability of having losses greater than or equal to x will be larger for the risk neutral solution. Furthermore, we can see in Figure 3 that the probability of not having to include a remedial contribution in the second stage is around .6, while for the other solutions it is around .7 and .8. For the third stage we can see in Figure 4 that the probability is around .5 while for the other solutions it is between .6 and .7.

With the contribution rate the situation is the opposite: the risk neutral solution is stochastically dominated in first order by all the other solutions in the second stage, and it is dominated in first order by all but the (.95,0,0) in the third stage. In the third stage for example almost 90% of the scenarios result in the fund taking no contribution at all from the participants. While this may seem desirable, we know from the remedial contribution graphs that the fund will probably have to inject a significant amount of money in order to satisfy the funding ratio constraints.

Let us take a closer look at other solutions. The extreme risk-averse solution, (.95,.95,.95) is much better than the risk-neutral solution in terms of remedial contribution in both stages. Moreover, it dominates (0,.95,0) in the first order and (.95,0,0) in the second order. That is not surprising: the solution (.95,.95,.95) is equally risk-averse in all stages, while (0.95,0) and (.95,0,0) in some sense protect the decision maker with more weight in early stages. With respect to the contribution rate, we see that (0,.95,0) dominates in first order all the other solutions in the second stage and three other solutions in the third stage. The decision maker that implements this solution can expect lower remedial contributions but contribution rates close to the upper bound limit of .21.

The (.5, .5, .5) is an interesting case of a compromise solution. The decision maker is risk-averse in every stage, but the α value is halfway between risk neutrality and extreme risk aversion. In both stages the curves for remedial contribution and contribution rate lie roughly in the middle of all other curves, indicating that very large remedial contribution and very small contribution rates are unlikely.

Now we focus our attention on the second- and third-stage decisions with respect to investment in risky assets (stocks and real estate) and safer instruments (bonds and cash). In Figures 3 and 4 we observe that the risk neutral solution will probably invest heavy on risky assets and much less on safer instruments. Given the results obtained for the remedial contribution such findings are not surprising: the risk neutral decision maker will try to meet his obligations by betting on the possible high returns of stock and real estate. A similar analysis can be performed for the solution (0,0,.95). In fact for risky assets this solution stochastically dominates in second order all the other solutions in the second stage, while the risk neutral dominates all but (0,0,.95) and (.95,0,0). We can see from Figures 3 and 4 that the other solutions follow a more conservative pattern, with significantly smaller probabilities of investment in risky assets. For example the probability of investing less than 1.0×10^4 in risky assets is on average 45% for the more risk-averse solutions while it is only 10% for the risk-neutral solution.

For bonds and cash, the risk-neutral solution is stochastically dominated in second order by all the other solutions in both second and third stages (Figure 4). Thus, the results inform the decision maker that by implementing the risk neutral solution in the first stage he should be prepared in the subsequent stages to face the more volatile stock and real estate markets. For investments in bonds and cash, the solution (.95,.95,.95) stochastically dominates in second order all the other solutions in the third stage. Similar to what was observed for the remedial contribution and the contribution rate, the (.5,.5,.5) case is between all curves, specially in the third stage, indicating a compromise between the investment on risky and safer assets.

7 Conclusions

The incorporation of measures of risk into multistage stochastic programs has received great attention recently in the literature, as there is a clear need to model a decision maker's risk attitude in that context. There is however no consensus on how to do that; in fact, in our conversations with other researchers we have heard a variety of opinions. Our goal in this paper is three-fold: first, we wanted to illustrate by means of a simple example the pitfalls of certain types of multi-period risk measures used in practice. We study such drawbacks under the framework of a formal notion of consistency which we define in terms of an optimization problem. Providing such a formal optimization-oriented definition of consistency is our

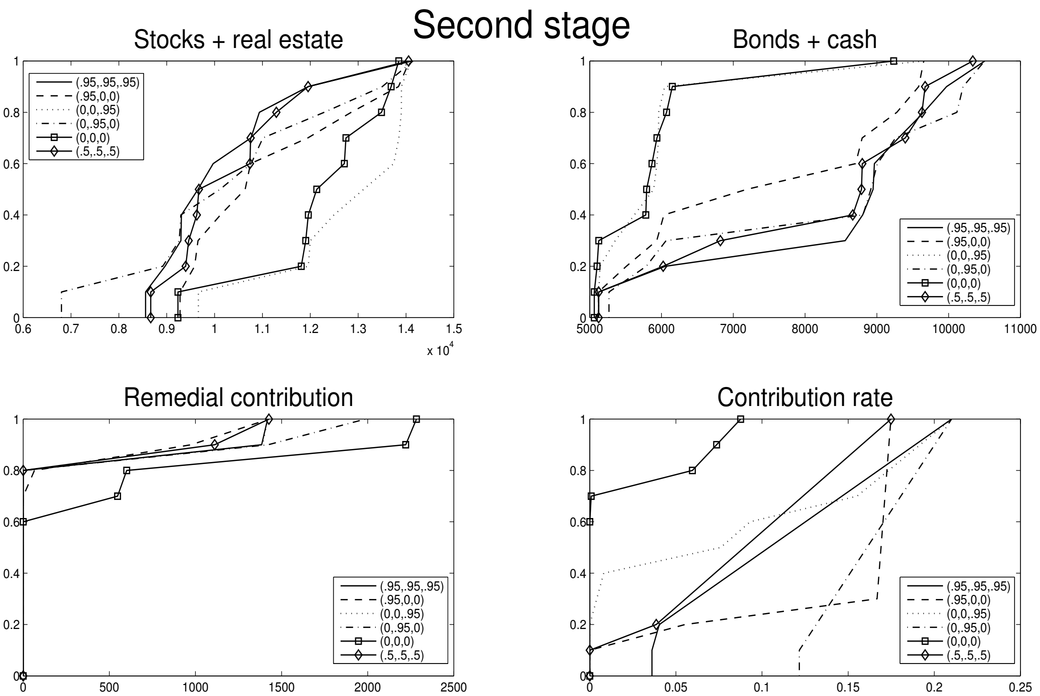

Figure 3: A representation of the second-stage solution.

second goal in this paper. The third goal is to propose a class of multi-period risk measures, which we call expected conditional risk measures (ECRMs), which extend the multi-period CVaR risk measure originally proposed by Pflug and Ruszczyński (2005). As we discuss in the paper, ECRMs have some attractive properties both from a modeling perspective as well as from an algorithmic standpoint. The modeling advantages include the fact that ECRMs give a more intuitive meaning to the decision maker about what is being measured, which in particular allows him the flexibility of choosing different risk levels in different stages. In the particular case of ECRMs with CVaR—which we denote by $\mathbb{E}$-CVaR in line with our more general definition—we have seen that the resulting model can be written as a risk-neutral problem with extra variables, thus allowing for the use of the various methods developed for multistage risk-neutral problems such as Nested Decomposition (Donohue and Birge, 2006) or the SDDP algorithm of Pereira and Pinto (1991). We believe that these features make ECRMs an attractive choice, as it overcomes some issues that arise with other alternatives proposed in the literature such as the nested CVaR. Our study of the pension fund example demonstrates the applicability of ECRMs.

Acknowledgements

The authors thank Maarten van der Vlerk for providing the source files for the pension fund experiment. They also thank Alexander Shapiro for his remarks on an earlier version of

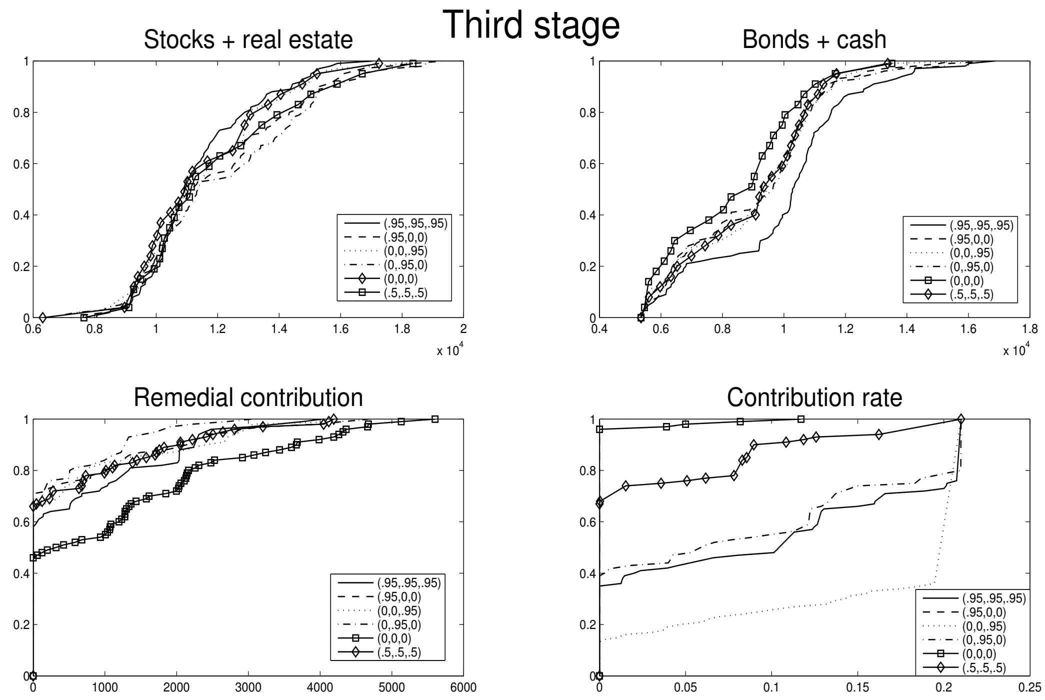

Figure 4: A representation of the third-stage solution.

this paper, and two anonymous referees for their constructive comments. This research was supported by FONDECYT projects number 1120244 and 11130056.

References

Ahmed, S., 2006. Convexity and decomposition of mean-risk stochastic programs. Mathematical Programming 106 (3), 433–446.

Artzner, P., Delbaen, F., Eber, J.-M., Heath, D., 1999. Coherent measures of risk. Mathematical finance 9 (3), 203–228.

Bion-Nadal, J., 2008. Dynamic risk measures: Time consistency and risk measures from BMO martingales. Finance and Stochastics 12 (2), 219–244.

Carpentier, P., Chancelier, J.-P., Cohen, G., De Lara, M., Girardeau, P., 2012. Dynamic consistency for stochastic optimal control problems. Annals of Operations Research 200 (1), 247–263.

Cheridito, P., Delbaen, F., Kupper, M., 2006. Dynamic monetary risk measures for bounded discrete-time processes. Electronic Journal of Probability 11 (3), 57–106.

Collado, R. A., Papp, D., Ruszczyński, A., 2012. Scenario decomposition of risk-averse multistage stochastic programming problems. Annals of Operations Research 200 (1), 147–170.

De Lara, M., Leclère, V., 2014. Time-consistency: from optimization to risk measures, manuscript, available at `https://hal.archives-ouvertes.fr/hal-00848268v2`.

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

Donohue, C., Birge, J. R., 2006. The abridged nested decomposition method for multistage stochastic programming. Algorithmic Operations Research 1 (1), 20–30.

Dupačová, J., Gröwe-Kuska, N., Römisch, W., 2003. Scenario reduction in stochastic programming: An approach using probability metrics. Math. Program. 95, 493–511.

Eichhorn, A., Römisch, W., 2005. Polyhedral risk measures in stochastic programming. SIAM Journal on Optimization 16 (1), 69–95.

Fábián, C. I., 2008. Handling CVaR objectives and constraints in two-stage stochastic models. European Journal of Operational Research 191 (3), 888–911.

Guigues, V., Sagastizábal, C., 2013. Risk-averse feasible policies for large-scale multistage stochastic linear programs. Mathematical Programming 138 (1-2), 167–198.

Haneveld, W. K. K., Streutker, M. H., Van Der Vlerk, M. H., 2010. An ALM model for pension funds using integrated chance constraints. Annals of Operations Research 177 (1), 47–62.

Heitsch, H., Römisch, W., 2009. Scenario tree modeling for multistage stochastic programs. Math. Program. 118, 371–406.

Høyland, K., Wallace, S., 2001. Generating scenario trees for multistage decision problems. Manag. Sci. 47, 295–307.

Kall, P., Mayer, J., 1996. SLP-IOR: An interactive model management system for stochastic linear programming. Mathematical Prog 75, 221–240.

Kilianova, S., Pflug, G. C., 2009. Optimal pension fund management under multi-period risk minimization. Annals of Operations Research 166, 261–270.

Kovacevic, R., Pflug, G., 2014. Are time consistent valuations information monotone? International Journal of Theoretical and Applied Finance 17 (1), 1443–1471.

Kozmık, V., Morton, D. P., 2013. Risk-averse stochastic dual dynamic programming, manuscript, available at Optimization Online (`www.optimization-online.org`).

Mehrotra, S., Papp, D., 2013. Generating moment matching scenarios using optimization techniques. SIAM J. Optim. 23 (2), 963–999.

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

Müller, A., Stoyan, D., 2002. Comparison Methods for Stochastic Models and Risks. John Wiley & Sons, Chichester.

Noyan, N., 2012. Risk-averse two-stage stochastic programming with an application to disaster management. Computers & Operations Research 39 (3), 541–559.

Pagnoncelli, B. K., Piazza, A., 2012. The optimal harvesting problem under risk aversion, manuscript, Universidad Adolfo Ibáñez, submited for publication.

Pereira, M., Pinto, L. M., 1991. Multi-stage stochastic optimization applied to energy planning. Mathematical Programming 52 (1-3), 359–375.

Pflug, G., 2006. A value-of-information approach to measuring risk in multi-period economic activity. Journal of Banking & Finance 30, 695–715.

Pflug, G., Römisch, W., 2007. Modeling, measuring and managing risk. World Scientific Pub Co Inc.

Pflug, G. C., 2000. Some remarks on the value-at-risk and the conditional value-at-risk. In: Probabilistic constrained optimization. Springer, pp. 272–281.

Pflug, G. C., 2001. Scenario tree generation for multiperiod financial optimization by optimal discretization. Math. Program., Ser. B 89 (2), 251–271.

Pflug, G. C., Pichler, A., 2011. Approximations for probability distributions and stochastic optimization problems. In: Bertocchi, M., Consigli, G., Dempster, M. A. H. (Eds.), Stochastic Optimization Methods in Finance and Energy. Springer, pp. 343–387.

Pflug, G. C., Pichler, A., 2012. A distance for multistage stochastic optimization models. SIAM Journal on Optimization 22 (1), 1–23.

Pflug, G. C., Pichler, A., 2014a. Time consistent decisions and temporal decomposition of coherent risk functionals, manuscript, available at Optimization Online (`www.optimization-online.org`).

Pflug, G. C., Pichler, A., 2014b. Time-inconsistent multistage stochastic programs: martingale bounds, manuscript, available at Optimization Online (`www.optimization-online.org`).

Pflug, G. C., Ruszczyński, A., 2005. Measuring risk for income streams. Computational Optimization and Applications 32 (1), 161–178.

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, forthcoming.

Rockafellar, R., Uryasev, S., 2002. Conditional value-at-risk for general loss distributions. Journal of Banking & Finance 26 (7), 1443–1471.

Rockafellar, R. T., Uryasev, S., 2000. Optimization of conditional value-at-risk. Journal of Risk 2, 21–42.

Rudloff, B., Street, A., Valladão, D., 2014. On the economic interpretation of time consistent dynamic stochastic programming problems, European Journal of Operational Research, forthcoming.

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. Math. of OR 31 (3), 544–561.

Schultz, R., Tiedemann, S., 2006. Conditional value-at-risk in stochastic programs with mixed-integer recourse. Mathematical Programming 105 (2-3), 365–386.

Shapiro, A., 2009. On a time consistency concept in risk averse multistage stochastic programming. Operations Research Letters 37 (3), 143–147.

Shapiro, A., 2012a. Minimax and risk averse multistage stochastic programming. European Journal of Operational Research 219 (3), 719–726.

Shapiro, A., 2012b. Time consistency of dynamic risk measures. Operations Research Letters 40 (6), 436–439.

Shapiro, A., Dentcheva, D., Ruszczyński, A., 2009. Lectures on stochastic programming: modeling and theory. Vol. 9. Society for Industrial Mathematics.

Shapiro, A., Tekaya, W., Soares, M. P., da Costa, J. P., 2013. Worst-case-expectation approach to optimization under uncertainty. Operations Research 61 (6), 1435–1449.

Wallace, S. W., Ziemba, W. T., 2005. Applications of stochastic programming. Vol. 5. SIAM Mathematical Series on Optimization.

Xin, L., Goldberg, D. A., Shapiro, A., 2013. Distributionally robust multistage inventory models with moment constraints, manuscript, available at Optimization Online (`www.optimization-online.org`).