Operations Research Letters 40 (2012) 313–318

Contents lists available at SciVerse ScienceDirect

Operations Research Letters

journal homepage: www.elsevier.com/locate/orl

SDDP for multistage stochastic linear programs based on spectral risk measures

Vincent Guigues [a,b], Werner Römisch [c,*]

[a] IMPA, Instituto de Matemática Pura e Aplicada, 110 Estrada Dona Castorina, Jardim Botanico, Rio de Janeiro, Brazil

[b] UFRJ, Escola Politécnica, Departamento de Engenharia Industrial, Ilha do Fundão, CT, Bloco F, Rio de Janeiro, Brazil

[c] Humboldt-University Berlin, Institute of Mathematics, 10099 Berlin, Germany

ARTICLE INFO

Article history:
Received 8 February 2012
Received in revised form
29 March 2012
Accepted 10 April 2012
Available online 15 May 2012

Keywords:
Spectral risk measure
Stochastic programming
Risk-averse optimization
Decomposition algorithms
Monte Carlo sampling

ABSTRACT

We consider risk-averse formulations of multistage stochastic linear programs. For these formulations, based on convex combinations of spectral risk measures, risk-averse dynamic programming equations can be written. As a result, the Stochastic Dual Dynamic Programming (SDDP) algorithm can be used to obtain approximations of the corresponding risk-averse recourse functions. This allows us to define a risk-averse nonanticipative feasible policy for the stochastic linear program. Formulas for the cuts that approximate the recourse functions are given. In particular, we show that some cut coefficients have analytic formulas.

1. Introduction

Multistage stochastic programs play a central role when developing optimization models under stochastic uncertainty in engineering, transportation, finance, and energy. Furthermore, since measuring, bounding, or minimizing the risk of decisions becomes more and more important in applications, risk-averse formulations of such optimization models are needed and have to be solved. Several risk-averse model variants allow for a reformulation as a classical multistage model, as in [6,8] and the present paper. From a mathematical point of view, multistage stochastic optimization methods represent infinite-dimensional models in spaces of random vectors satisfying certain moment conditions and contain high-dimensional integrals. Hence, their numerical solution is a challenging task. Each solution approach consists at least of two ingredients: (i) numerical integration methods for computing the expectation functionals and (ii) algorithms for solving the resulting finite-dimensional optimization models.

The favorite approach for (i) is to generate possible scenarios (i.e., realizations) of the random vector involved and to use them as 'grid points' for the numerical integration. Scenario generation can be done by Monte Carlo, quasi-Monte Carlo, or optimal quantization methods (see [5,18] for overviews and [3, Part III] for further information). Scenarios for multistage stochastic programs have to be tree structured to model the increasing chain of σ-fields. Existing stability and convergence results such as those in [11,10,12,21] provide approaches and conditions implying the convergence of such schemes, in particular, for the deterministic first-stage solutions. Hence, they justify rolling horizon approaches based on repeated solving of multistage models; see [9], for instance.

The algorithms employed for (ii) depend on structural properties of the basic optimization model and on the inherent structure induced by the scenario tree approximation (see the survey [19] on decomposition methods).

Some algorithmic approaches incorporate the scenario generation method (i) as an algorithmic step of the solution method. Such approaches are, for example, stochastic decomposition methods for multistage models (see [20]), approximate dynamic programming (see [17]), and Stochastic Dual Dynamic Programming (SDDP), initiated in [13], revisited in [16,22], and also studied in the present paper.

We consider risk-averse formulations of multistage stochastic linear programs of the form

$$\begin{aligned} &\inf_{x_1,\ldots,x_T} d_1^\top x_1 + \theta_1 \mathbb{E}\left[\sum_{t=2}^{T} d_t^\top x_t\right] + \sum_{t=2}^{T} \theta_t \rho_\phi\left(-\sum_{k=2}^{t} d_k^\top x_k\right) \\ &C_t x_t = \xi_t - D_t x_{t-1}, \\ &x_t \geq 0, \quad x_t \text{ is } \mathcal{F}_t\text{-measurable}, t = 1, \ldots, T, \end{aligned} \tag{1}$$

* Corresponding author. Tel.: +49 3020932561.
E-mail addresses: vguigues@impa.br (V. Guigues), romisch@math.hu-berlin.de (W. Römisch).

0167-6377/$ – see front matter © 2012 Elsevier B.V. All rights reserved.
doi:10.1016/j.orl.2012.04.006

where x_0 is given, parameters d_t, C_t, D_t are deterministic, $(\xi_t)_{t=1}^T$ is a stochastic process, $\mathcal{F}_t$ is the sigma-algebra $\mathcal{F}_t := \sigma(\xi_j, j \le t)$, $(\theta_t)_{t=1}^T$ are nonnegative weights summing to 1, and ρ_ϕ is a *spectral risk measure* [1] or *distortion risk measure* [14,15] depending on a risk spectrum $\phi \in L_1([0,1])$. In the above formulation, we have assumed that the (one-period) spectral risk measure takes as argument a random income and that the trajectory of the process is known until the first stage. We assume *relatively complete recourse* for (1), which means that, for any feasible sequence of decisions $(x_1, \ldots, x_t)$ to any t-stage scenario $(\xi_1, \xi_2, \ldots, \xi_t)$, there exists a sequence of feasible decisions $(x_{t+1}, \ldots, x_T)$ with probability 1. A non-risk-averse model amounts to taking $\theta_1 = 1$ and $\theta_t = 0$ for $t = 2, \ldots, T$. A more general risk-averse formulation for multistage stochastic programs is considered in [8]. For these models, dynamic programming (DP) equations are written in [8] and an SDDP algorithm is detailed to obtain approximations of the corresponding recourse functions in the form of cuts. The main contribution of this paper is to provide analytic formulas for some cut coefficients, independent of the sampled scenarios, that can be useful for implementation. We also specialize the SDDP algorithm and especially the computation of the cuts for the particular risk-averse model (1).

We start by setting down some notation.

- e will denote a column vector of all 1s;
- for $x, y \in \mathbb{R}^n$, the vector $x \circ y \in \mathbb{R}^n$ is defined by $(x \circ y)(i) = x(i)y(i),\ i = 1, \ldots, n$;
- for $x \in \mathbb{R}^n$, the vector $x^+ \in \mathbb{R}^n$ is defined by $x^+(i) = \max(x(i), 0),\ i = 1, \ldots, n$;
- the available history of the process at stage t is denoted by $\xi_{[t]} := (\xi_j, j \le t)$;
- for vectors $x_1, \ldots, x_n$, the notation $x_{n_1:n_2}$ stands for the concatenation $(x_{n_1}, x_{n_1+1}, \ldots, x_{n_2})$ for $1 \le n_1 \le n_2 \le n$;
- δ_{ij} is the Kronecker delta defined for i, j integers by $\delta_{ij} = 1$ if $i = j$ and 0 otherwise.

2. Risk-averse dynamic programming

Let $F_Z(x) = \mathbb{P}(Z \le x)$ be the cumulative distribution function of an essentially bounded random variable Z, and let $F_Z^{\leftarrow}(p) = \inf\{x : F_Z(x) \ge p\}$ be the generalized inverse of F_Z. Given a risk spectrum $\phi \in L_1([0,1])$, the spectral risk measure ρ_ϕ generated by ϕ is (see [1]):

$$\rho_\phi(Z) = -\int_0^1 F_Z^{\leftarrow}(p)\phi(p)dp.$$

Spectral risk measures have been used in various applications (portfolio selection by Acerbi and Simonetti [2]; insurance by Cotter and Dowd [4]). The conditional value-at-risk (CVaR) of level $0 < \varepsilon < 1$, denoted by $CVaR^\varepsilon$, is a particular spectral risk measure obtained taking $\phi(u) = \frac{1}{\varepsilon} 1_{0 \le u < \varepsilon}$ (see Acerbi [1]).

In what follows, we consider more generally a piecewise constant risk function $\phi(\cdot)$ with J jumps at $0 < p_1 < p_2 < \cdots < p_J < 1$. We set $\Delta\phi_k = \phi(p_k^+) - \phi(p_k^-) = \phi(p_k) - \phi(p_{k-1}),$ for $k = 1, \ldots, J$, with $p_0 = 0$, and we assume that

$$\text{(i) } \phi(\cdot) \text{ is positive}, \qquad \text{(ii) } \Delta\phi_k < 0, \quad k = 1, \ldots, J,$$
$$\text{(iii) } \int_0^1 \phi(u)du = 1.$$

In this context, ρ_ϕ can be expressed as a linear combination of conditional value-at-risk measures. With this choice of risk function ϕ, the spectral risk measure $\rho_\phi(Z)$ can be expressed as the optimal value of a linear program; see Acerbi and Simonetti [2]:

$$\rho_\phi(Z) = \inf_{w \in \mathbb{R}^J} \sum_{k=1}^J \Delta\phi_k[p_k w_k - \mathbb{E}[w_k - Z]^+] - \phi(1)\mathbb{E}[Z]. \tag{2}$$

Using this formulation for ρ_ϕ, dynamic programming equations are given in [8] for risk-averse formulation (1). More precisely, problem (1) can be expressed as

$$\begin{array}{l} \inf\limits_{x_1,\, w_{2:T}} d_1^\top x_1 + \sum\limits_{t=2}^T \theta_t c_1^\top w_t + \mathcal{Q}_2(x_1, \xi_{[1]}, z_1, w_2, \ldots, w_T), \\ C_1 x_1 = \xi_1 - D_1 x_0, \quad x_1 \ge 0,\ w_t \in \mathbb{R}^J,\ t = 2, \ldots, T, \end{array} \tag{3}$$

with $z_1 = 0$, vector $c_1 = \Delta\phi \circ p$, and where, for $t = 2, \ldots, T$,

$$\mathcal{Q}_t(x_{t-1}, \xi_{[t-1]}, z_{t-1}, w_{t:T}) = \mathbb{E}_{\xi_t|\xi_{[t-1]}} \begin{pmatrix} \inf\limits_{x_t, z_t} f_t(z_t, w_t) + \mathcal{Q}_{t+1}(x_t, \xi_{[t]}, z_t, w_{t+1:T}) \\ z_t = z_{t-1} - d_t^\top x_t,\ C_t x_t = \xi_t - D_t x_{t-1},\ x_t \ge 0 \end{pmatrix}, \tag{4}$$

with

$$f_t(z_t, w_t) = -(\delta_{tT}\theta_1 + \phi(1)\theta_t)z_t - \theta_t\ \Delta\phi^\top(w_t - z_t e)^+, \tag{5}$$

and $\mathcal{Q}_{T+1} \equiv 0$. Function $\mathcal{Q}_{t+1}$ represents at stage t a cost-to-go or recourse function which is risk averse. As shown in the next section, it can be approximated by cutting planes by some polyhedral function $\mathfrak{Q}_{t+1}$. These approximate recourse functions are useful for defining a feasible approximate policy obtained by solving

$$\begin{array}{l} \inf\limits_{x_t, z_t} f_t(z_t, w_t) + \mathfrak{Q}_{t+1}(x_t, \xi_{[t]}, z_t, w_{t+1:T}) \\ C_t x_t = \xi_t - D_t x_{t-1}, \quad x_t \ge 0,\ z_t = z_{t-1} - d_t^\top x_t, \end{array} \tag{6}$$

at stage $t = 2, \ldots, T$, knowing x_{t-1}, z_{t-1}, first-stage decision variables $w_{t:T}$, and ξ_t. First-stage decision variables x_1 and $w_{2:T}$ are the solution to (3) with $\mathcal{Q}_2$ replaced by the approximation $\mathfrak{Q}_2$.

3. Algorithmic issues

The DP equations (3)–(4) make possible the use of decomposition algorithms such as SDDP to obtain approximations of the corresponding recourse functions. When applied to DP equations (3)–(4), the convergence of this algorithm is proved in [8] under the following assumptions.

(A1) The supports of the distributions of $\xi_1, \ldots, \xi_T$ are discrete and finite.
(A2) Process (ξ_t) is interstage independent.
(A3) For $t = 1, \ldots, T$, for any feasible x_{t-1}, and for any realization $\tilde{\xi}_t$ of ξ_t, the set

$$\{x_t : x_t \ge 0,\ C_t x_t = \tilde{\xi}_t - D_t x_{t-1}\}$$

is bounded and nonempty.

In what follows, we assume that Assumptions (A1)–(A3) hold. In particular, we denote the realizations of ξ_t by $\xi_t^i,\ i = 1, \ldots, q_t < +\infty$, and set $p(t,i) = \mathbb{P}(\xi_t = \xi_t^i)$.

Since the supports of the distributions of the random vectors $\xi_2, \ldots, \xi_T$ are discrete and finite, optimization problem (1) is finite dimensional, and the evolution of the uncertain parameters over the optimization period can be represented by a scenario tree having a finite number of scenarios that can happen in the future for $\xi_2, \ldots, \xi_T$. The root node of the scenario tree corresponds to the first time step with ξ_1 deterministic.

For a given stage t, to each node of the scenario tree there corresponds an history $\xi_{[t]}$. The children nodes of a node at stage $t \ge 1$ are the nodes that can happen at stage $t+1$ if we are at this node at t. A sampled scenario $(\xi_1, \ldots, \xi_T)$ corresponds to a particular succession of nodes such that ξ_t is a possible value for the process at t and ξ_{t+1} is a child of ξ_t. A given node in the tree at stage t is identified with a scenario $(\xi_1, \ldots, \xi_t)$ going from the root node to this node.

In this context, the SDDP algorithm builds polyhedral lower bounding approximations $\mathfrak{Q}_t$ of $\mathcal{Q}_t$ for $t = 2, \ldots, T+1$. Each