

Dual SDDP for risk-averse multistage stochastic programs

Bernardo F. P. Da Costa¹, Vincent Leclère²

¹ UFRJ, Rio de Janeiro, Brazil

bernardofpc@im.ufrj.br

² CERMICS, Ecole des Ponts, Marne-la-Vallée, France

vincent.leclere@enpc.fr

Mots-clés : *Stochastic programming, SDDP, Risk measures, Duality.*

1 Introduction

Multistage stochastic programming is a powerful framework with multiple applications, *e.g.* in the finance, energy and supply chain sectors. If the uncertainty is finitely supported, those problems can be seen as large-scale deterministic problems. When there is more than 4 or 5 stages, the deterministic equivalent is usually too large to be solved directly. One of the most successful paradigms in this setting consists in leveraging time-independence assumptions to derive Bellman equations. The Stochastic Dual Dynamic Programming (SDDP) algorithm, and its numerous variants, consists in using those equations to derive approximations of the *cost-to-go* functions. It has been successfully used on a number of real-world problems, especially in the field of energy. We show how to derive a dual formulation for these problems and apply an SDDP algorithm, leading to converging and deterministic upper bounds for risk averse problems. Details can be found in [1].

While the classical formulation of a multistage program is *risk-neutral*, meaning that we minimize an expected cost, a large part of the recent literature has been devoted to efficiently introduce *risk aversion* in this framework, in particular inside the SDDP algorithm. In particular, coherent risk measures have become a usual tool to represent risk aversion in stochastic optimization problems. In multistage stochastic programming, minimizing a risk measure of the sum of costs leads to time-inconsistency. The easiest way to come up with a time-consistent risk-averse problem is to use composed Markovian risk measures, which, roughly speaking, means replacing the expectation by a risk measure inside the dynamic programming equation.

More precisely, let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space, and $\{\omega_t\}_{t \in [T]}$ be a sequence random variables. We consider the following risk-averse multistage linear program (RA-MSLP)

$$\min_{x_t, y_t} \rho_1 \left(c_1^\top y_1 + \rho_{2|\omega_1} \left(\cdots + \rho_{T|\omega_{[T-1]}} (c_T^\top y_T) \right) \right) \quad (1a)$$

$$s.t. \quad A_t x_t + B_t x_{t-1} + T_t y_t = d_t \quad \forall t \in [T] \quad (1b)$$

$$x_t \geq 0, y_t \geq 0 \quad \forall t \in [T] \quad (1c)$$

$$x_t, y_t \preceq \omega_{[t]} \quad \forall t \in [T] \quad (1d)$$

where $\rho_{t|\omega_{[t]}}$ is a coherent risk measure conditional on the past noises $\omega_{[t]} := \{\omega_1, \dots, \omega_t\}$, all equalities hold almost surely, and constraint (1d) is the *non-anticipativity* constraint, ensuring that decisions x_t, y_t are taken with only the information available at time t .

Assuming that $\{\omega_t\}_{t \in [T]}$ is a sequence of *independent* random variables, Dynamic Program-

ming leads to the following recursion :

$$V_{T+1}(x_T) = 0, \tag{2}$$

$$\begin{aligned} V_t(x) = \min_{x_t, y_t} & \rho_t [c_t^\top y_t + V_{t+1}(x_t)] \\ \text{s.t.} & \quad A_t x_t + B_t x + T_t y_t = d_t \\ & \quad x_t \geq 0, y_t \geq 0 \end{aligned} \tag{3}$$

where the value of Problem (1) is given by $V_1(x_0)$.

The classical SDDP algorithm builds outer approximations of the cost-to-go function, leading to exact lower bounds on the problem. In a risk-neutral framework, upper bounds can be estimated via Monte Carlo sampling. Unfortunately it is unclear how to extend such statistical methods to the risk-averse setting. Instead of statistical upper bounds, one can use exact upper bounds : Through backward recursion ([4]); by maintaining upper and lower bounds for all value functions; or using Fenchel duality ([3, 2]). While backward recursion can be used for risk-averse problems, it is more a one-time evaluation than an improving process. On the other hand, maintaining upper and lower bounds for the value functions can be used in a risk-averse setting to yield a converging algorithm.

In the risk-neutral case, it has been shown in [3] that we can derive a Bellman-like recursion for the Fenchel transform of V . A recursion can also be obtained by dualizing the extensive formulation of the MLSP problem, and recognizing a time-decomposition, as presented in [2]. In both cases, this provides a dual problem on which SDDP can be applied, and whose lower bounds correspond to upper bounds of (1). Here we develop these ideas in the risk-averse case : we dualize the extended formulation and show a time-decomposition for the dual problem. We also deduce a Bellman recursion on the perspective of the Fenchel transform of the value function. This leads to an algorithm providing converging upper bounds for the risk-averse setting.

Références

- [1] Bernardo Freitas Da Costa and Vincent Leclère. *Dual SDDP for risk-averse multistage stochastic programs*. arXiv preprint arXiv :2107.10930, 2021.
- [2] Vincent Guigues, Alexander Shapiro and Yi Cheng. *Duality and sensitivity analysis of multistage linear stochastic programs*. arXiv preprint arXiv :1911.07080, 2019.
- [3] Vincent Leclère, Pierre Carpentier, Jean-Philippe Chancelier, Arnaud Lenoir and François Pacaud. Exact converging bounds for stochastic dual dynamic programming via fenchel duality. *SIAM Journal on Optimization*, 2020.
- [4] Andrew Philpott, Vitor de Matos and Erlon Finardi. *On solving multistage stochastic programs with coherent risk measures*. Operations Research, 2013.