RISK-AVERSE FEASIBLE POLICIES FOR LARGE-SCALE MULTISTAGE STOCHASTIC LINEAR PROGRAMS

VINCENT GUIGUES AND CLAUDIA SAGASTIZÁBAL

Abstract. We consider risk-averse formulations of stochastic linear programs having a structure that is common in real-life applications. Specifically, the optimization problem corresponds to controlling over a certain horizon a system whose dynamics is given by a transition equation depending affinely on an interstage dependent stochastic process. We put in place a rolling horizon, time consistent, policy that for each time step defines a risk-averse problem with constraints that are deterministic for the current time step and uncertain for future times. To each uncertain constraint corresponds both a chance and a Conditional Value-at-Risk constraint. We show that the resulting risk-averse problems are numerically tractable, being at worst conic quadratic programs. For the particular case in which uncertainty appears only on the right-hand side of the constraints, such risk-averse problems are linear programs. We show how to write dynamic programming equations for these problems and define robust recourse functions. For multistage stochastic linear programs, an algorithm commonly used to obtain approximations of recourse functions is Stochastic Dual Dynamic Programming (SDDP). In our case this algorithm does not apply. However, using the convexity of our robust recourse functions, we present an algorithm to approximate them that shares some features with SDDP. To assess the methodology and compare it with SDDP, we consider a water-resource planning problem and obtain encouraging numerical results.

AMS subject classifications: 90C15, 91B30.

1. Introduction

Hedging risk is a challenging question in Stochastic Programming. Since uncertainty may appear in the objective function and/or in the constraints, risk-averse proposals can be gathered in three groups, depending on how uncertainty (and, hence, risk) is dealt with. The first two groups deal with uncertainty in both the objective function and in the constraints. More precisely, methods in the first group minimize some risk measure of the objective function, which depends on the decisions and on the underlying random variables. In this group, the so-called dynamic risk mappings [30], [25], including the different "perspectives" for polyhedral risk measures in [15], appear as powerful and useful tools. The second group, Robust Optimization, can use available data to define uncertainty sets and apply the worst case oriented robust optimization methodology [4]. Finally, general stochastic programs can always be reformulated as optimization problems with deterministic objective function and uncertain constraints. This makes possible the introduction of a last group of risk-averse models, especially designed to control risk on constraints. Among

Key words and phrases. Stochastic programming and chance constraints and CVaR and interstage dependence and dynamic programming and rolling horizon.
The first author's research was supported by CNPq grant No. 382.851/07-4. The work of the second author was partially supported by CNPq grant No. 300345, PRONEX-Optimization, and FAPERJ.

several proposals in this group, we mention the classical chance constraints [8], [27] and the more recent integrated chance constraints [19], [20], and stochastic ordering constraints [13].

Our work belongs to the last group above, with some distinctive features. Specifically, most of the works in this group replace each uncertain constraint by a risk-averse counterpart, written as a single constraint. Instead, we use a couple of constraints in the risk-averse formulation: to each uncertain constraint corresponds both a chance and a Conditional Value-at-Risk (CVaR) constraint. Our goal is twofold; while a chance constraint gives a qualitative control for constraint violation, regardless of the amount of violation, a CVaR constraint gives a quantitative measure, by keeping control of such amount.

In addition, rather than defining a single risk-averse problem, we adopt a *rolling horizon* point of view; there are as many risk-averse problems as time steps in the optimization period. At a given time step, the corresponding risk-averse problem considers constraints to be deterministic for the current time step and uncertain for future times. As a result, our approach builds *feasible* policies such that all the constraints over the optimization period are satisfied almost surely. This is a very important property in applications having some "hard" constraints that need to be satisfied with probability one.

Another crucial matter when dealing with probabilistic constraints refers to numerical tractability. In this respect, it is useful to take full advantage of the initial structure. For this reason, we consider special stochastic linear programs arising when controlling, over a multiperiod horizon, a dynamical system with a transition equation depending affinely on an interstage dependent stochastic process[1]:

$$
(1) \qquad \begin{cases} \min\limits_{x_{[1:T]},u_{[1:T]}} \sum\limits_{t=1}^{T} c_t^\top u_t & \\ x_t = A_{t-1}x_{t-1} + B_t u_t + C_t\tilde{\xi}_t + d_t & \text{for } t = 1,\dots,T, \qquad (\text{TRAN}) \\ E_t x_t + F_t(\tilde{\xi}_t)u_t \geq G_t\tilde{\xi}_t + h_t & \text{for } t = 1,\dots,T, \qquad (\text{INEQ}) \end{cases}
$$

where

- T is the number of time steps, possibly large;
- $\tilde{\xi}_t$ is a particular realization at time step t of an M-dimensional random process. Each process component $\xi_t(m)$ follows a generalized autoregressive model with time varying order; see Section 3.1 below. The realization $\tilde{\xi}_t$ becomes known at the beginning of time step t;
- $x_t \in \mathbb{R}^{N_x}$ is the state of the system at the end of time step t, with dynamics given by the transition equation (1)(TRAN) and known x_0;
- $u_t \in \mathbb{R}^{N_u}$ is the control variable, applied to the system at time step t; and
- $c_t^\top u_t$ is the cost at time step t.

[1]Observe that (1) is a problem instance, written for a given realization of the underlying process; for the moment,we do not explain how uncertainty is dealt with.

Matrices A_t, B_t, C_t, E_t, G_t are $N_x \times N_x$, $N_x \times N_u$, $N_x \times M$, $q_t \times N_x$, $q_t \times M$, respectively. The "technology" matrix in the inequality constraints (1)(INEQ), i.e.,

$$F_t(\tilde{\xi}_t) = \mathring{F}_t + \begin{bmatrix} (F_{t,1}\tilde{\xi}_t)^\top \\ \vdots \\ (F_{t,q_t}\tilde{\xi}_t)^\top \end{bmatrix}, \tag{2}$$

is an affine function of the process realization. The respective sizes of matrices $\mathring{F}_t$ and $F_{t,i}$ are $q_t \times N_u$ and $N_u \times M$. Finally, vectors d_t and h_t have dimensions N_x and q_t, respectively.

For many real-life planning problems, the evolution along the planning horizon is modelled by means of reservoirs (of a product, of water, of energy, of take-or-pay contracts, of oil). In this sense, the framework in (1) is rather comprehensive, and covers a variety of applications in inventory problems [2], [5], [21], electric energy, oil, and finance [29, Ch. 10], [16], [17], [33]; see also Section 2 below. For such applications, many of the constraints (INEQ) involve only the state x_t or only the control u_t, which corresponds to the fact that, when certain rows in E_t are not null, the corresponding rows in $F_t(\tilde{\xi}_t)$ are null, and reciprocally.

When applying our rolling horizon approach to problems of form (1), the special underlying structure -of both the stochastic process and problem constraints- yields *fully tractable risk-averse* problems for each time step. Equivalent deterministic formulations of chance-constrained problems have already been proposed for a limited number of statistical frameworks and classes of optimization problems. Our result generalizes to a somewhat broader setting similar results in [8] and [10]. More precisely, we show in Theorem 4.4 that the risk-averse problems are deterministic conic quadratic programs, that become linear programs if the technology matrix in (2) is not random, that is, if $F_{t,1} = \ldots = F_{t,q_t} = 0$.

Our methodology is specially attractive for problems with large time horizon T. In such cases, there are often so many scenarios that only sampling methods can be employed, [24], [11], [26]. In this setting, many risk-averse formulations yield huge multistage stochastic linear programs which are difficult to solve, when not intractable. Instead, with our approach we build a risk-averse feasible policy just by solving T deterministic linear or conic programs. Moreover, we show that such programs are decomposable by stages (an issue of special algorithmic interest in a large-scale context) and can be solved efficiently by Dual Dynamic Programming. Finally, we propose two algorithms to obtain approximations of the risk-averse recourse functions corresponding to our model.

Our paper is organized as follows. In Section 2 we start with an application example, the long-term planning of water reservoirs in hydro-thermal power systems. This application is used along the paper to explain and motivate the proposed approach. Section 3 gives the statistical model for the stochastic process as well as the main elements of our risk-averse rolling horizon approach. Section 4 shows that the risk-averse problems are numerically tractable. In particular, when these problems are linear programs, in Section 5 we give a stagewise decomposition and show how to approximate the robust risk-averse recourse functions mentioned above. Finally, in Section 6, the approach is assessed on the water-resource planning problem, by comparing its performance with the sampling method of Stochastic Dual Dynamic Programming from [24].

We adopt the following notation and conventions. For $t_2 \geq t_1$, the short form $v_{(t_1,t_2]}$ (resp. $v_{[t_1,t_2]}$) stands for the concatenation $(v_{t_1+1}, \ldots, v_{t_2})$ (resp. $(v_{t_1}, \ldots, v_{t_2})$), with $v_{(t,t]}$ vacuous and knowing that the concatenated objects v_j can be vectors or matrices, depending on the context. For sums and products, $\sum_{i=i_0}^{i_1} x_i = 0$ and $\prod_{i=i_0}^{i_1} x_i = 1$ whenever $i_0 > i_1$, knowing that for matrices X_i, if $i_0 > i_1$ then $\prod_{i=i_0}^{i_1} X_i = I$, the identity matrix. For a random variable X, its distribution function and its density are respectively denoted by $F_X(\cdot)$ and $f_X(\cdot)$, knowing that for $X \sim \mathcal{N}(0,1)$, we just write $F(\cdot) := F_X(\cdot)$. Finally, if X is continuous and larger values of the random variable are preferred, its Conditional Value-at-Risk of level $\varepsilon_{\mathtt{p}} \in [0,1]$ is denoted and defined by $CVaR_{\varepsilon_{\mathtt{p}}}(X) := -\mathbb{E}[X|X \leq F_X^{-1}(\varepsilon_{\mathtt{p}})]$, while the Value-at-Risk of level $\varepsilon_{\mathtt{p}}$ of X is $VaR_{\varepsilon_{\mathtt{p}}}(X) := -F_X^{-1}(\varepsilon_{\mathtt{p}})$; see [28].

2. Motivation

Our initial motivation is the long-term optimal management of water reservoirs in hydro-thermal power systems, [23]. For this problem, present operating decisions have future consequences that are difficult to quantify because the water is a commodity of unknown value and uncertain availability. Moreover, it is important to set the problem in a risk-averse framework so that indicators such as energy prices better reflect the impact of extreme events.

In this setting, it is desirable for the computational model to provide two types of output:

`(RecFun)`: approximate recourse functions for each time step $t = 1, \ldots, T$; and

`(EcInd)`: economic indicators, such as mean marginal energy prices, expected load shedding, average supplied energy.

The recourse functions from `(RecFun)`, obtained in some cases at the end of an *optimization phase*, can be seen as a pricing mechanism, giving value to water, and defining a policy. Namely, at a so-called *simulation phase*, with such prices it is possible to mimic the optimal operation of the hydro-thermal system. Recourse functions can also be used to couple planning models of different horizons and obtain a sound overall management of the system. Similar situations arise in the long-term control of other stochastic systems, as in [29, Ch. 10].

As for the indicators in `(EcInd)`, they measure the performance of the long-term management policy. They are computed by simulating the system operation over a high number of randomly generated scenarios, covering a large spectrum of foreseeable futures (including extreme droughts, floods, etc).

In Sections 3.2 and 5.2 we analyze how our robust rolling horizon approach compares to usual non-rolling horizon methodologies to fulfill requirements `(RecFun)` and `(EcInd)`.

We now give the mathematical formulation for a simplified model of the water management application that will help in clarifying our rolling horizon risk-averse methodology.

2.1. A simplified long-term energy planning problem.

The optimal management of a hydro-thermal power system in the long-term minimizes the operational cost along the period, subject to various technical constraints. Since the immediate hydro-cost is positive but negligible, operational costs are essentially related to the fuel burnt by thermal plants and penalties resulting from load shedding. In a predominantly hydro-electric system like Brazil's, the availability of (limited) amounts of hydro-power, in the form of water stored in reservoirs, makes the problem extremely

complex. There are many reservoirs in cascade, some of them with a capacity of regularization that covers several years, which are spread over geographical regions with different seasonal rainfall.

For simplicity, we consider only one reservoir with large enough capacity and suppose there is neither spillage nor upper bound on the turbines outflow. The system also has one run-of-river plant, but no thermal plants. In addition, the problem is formulated in energy variables, without entering into the issue of how to relate water to energy by explicit production functions.

In this simplified formulation, at each time step t we have a state variable x_t, the volume of the reservoir at the end of the time step; a control variable $u_t = (g_t, df_t)^\top \geq 0$, with g_t the turbines outflow, df_t the energy deficit; and $\tilde{\xi}_t$, the natural inflow of water arriving into the reservoir. Only a fraction $\gamma_t \in [0, 1)$ of this water can be stored; the remaining portion, $(1 - \gamma_t)\tilde{\xi}_t$, is immediately transformed into power by the run-of-river plant. We further assume that turbines of the run-of-river plant have enough capacity to generate power out of all of $(1 - \gamma_t)\tilde{\xi}_t$, without any losses.

In order to appropriately reflect seasonal variations, the stochastic process of water streamflows is usually represented by a periodic autoregressive model with a one year period, [22], [12]. If there are more reservoirs, the process is multivariate and there is a (nondiagonal) covariance matrix to express the dependence of inflows on neighboring geographical regions.

The objective function considers generation and shortage costs, and constraints are given below.

Water balance equation. If there is no evaporation,

$$x_t = x_{t-1} - g_t + \gamma_t \tilde{\xi}_t \tag{3}$$

corresponds to the transition equation (TRAN), with

$$A_{t-1} = 1,\ B_t = [-1 \quad 0],\ C_t = \gamma_t, \text{ and } d_t = 0.$$

In the presence of more reservoirs, A_{t-1} is an identity matrix I, B_t a concatenation of $-I$ with a zero matrix, $C_t = \gamma_t I$ is a diagonal matrix, and d_t a null vector.

Demand satisfaction. Usually this constraint has the form $g_t + df_t = \max(\text{dem}_t - (1 - \gamma_t)\tilde{\xi}_t, 0)$, where the energy deficit df_t is modelled as a (fictitious) thermal plant with large capacity and generation cost equal to the shortage cost; and dem_t is the demand at time step t, that we suppose deterministic for convenience.

Since our risk-averse formulation makes use of probabilistic constraints (that can be set only for inequality constraints), we re-write the demand satisfaction constraint as an inequality:

$$g_t + df_t \geq \text{dem}_t - (1 - \gamma_t)\tilde{\xi}_t\,. \tag{4}$$

Although, in general, modifying the feasible set may alter the optimization problem, such is not the case here when replacing the initial demand satisfaction constraint by inequality (4), thanks to the structure of the problem.

With respect to (INEQ), we see that the demand constraint corresponds to taking

$$E_t = 0,\ F_t = [1 \quad 1],\ G_t = -(1 - \gamma_t), \text{ and } h_t = \text{dem}_t,$$

with natural extensions to the multi-reservoir case.

Critical volume. Operators managing the system in real time are mostly concerned about keeping reservoirs at reasonable storage levels. In particular, they sometimes wish to keep the reservoirs

above critical values, or reference trajectories, estimated empirically by the operators or imposed by some regulatory rules. The corresponding constraints have the form

$$x_t \geq x_t^{crit}\,. \tag{5}$$

In the notation of (INEQ), constraint (5) sets

$$E_t = 1,\ F_t \text{ and } G_t \text{ null, and } h_t = x_t^{crit},$$

where the extension to more reservoirs is straightforward.

The values $x_t^{crit} \geq 0$ are given and such that,

(A1) along time steps, the critical levels are nonincreasing: $x_0 \geq x_1^{crit} \geq x_2^{crit} \geq \ldots \geq x_T^{crit}$.

This assumption, together with the inflow condition stating that $\mathbb{P}(\xi_t \geq 0) = 1$ for every t, implies that recourse in problem (1) is *relatively complete*, [6, Ch. 3, p. 92]. If Assumption (A1) did not hold or if we do not have $\xi_t \geq 0$ a.s., then the methodology described in the sequel would still be applicable and relatively complete recourse would hold adding slack variables $y_t \geq 0$ (penalized in the objective) to constraints (5): $x_t + y_t \geq x_t^{crit}$.

In this application the "technology" matrices F_t are deterministic and $F_{t,i} = 0, i = 1, \ldots, q_t$ ($q_t = 2$ for the simplified model considered in this section). If production functions are considered, and if such functions depend on uncertain factors, then matrices F_t are stochastic. Such is the case in the oil industry, for example [33], when dealing with refinery production planning with uncertain yields.

3. General setting

Consider the natural filtration $\mathcal{F}_1 \subset \ldots \subset \mathcal{F}_T$ induced by the process ξ_t, defining $\mathcal{F}_t$ as the sigma algebra $\sigma(\xi_j,\ j \leq t)$. We are interested in finding a sequence of state and control mappings $x_t(\cdot)$ and $u_t(\cdot)$ for $t = 1, \ldots, T$, i.e., an implementable *policy* satisfying the following properties:

> *Nonanticipativity:* state and control mappings $x_t(\cdot)$ and $u_t(\cdot)$ are $\mathcal{F}_t$-measurable and, hence, are functions of the available history $\tilde{\xi}_{[t]}$ of the process.
> *Feasibility:* a non-anticipative policy satisfying (INEQ) and (TRAN) for $t = 1, \ldots, T$, with probability one.

The assumption of relatively complete recourse ensures the existence of feasible policies for problem (1).

3.1. Statistical model.

Our multivariate discrete time stochastic process depends, in an affine manner, on previous values. In this context, for $m = 1, \ldots, M$, each component $\xi_t(m)$ is represented by a generalized autoregressive model, with varying orders $p_t(m) \geq 0$. Accordingly, for every integer t, there exist coefficients $\Phi_t^j(m)$ for $j = 1, \ldots, p_t(m)$, with non-null $\Phi_t^{p_t(m)}(m)$, such that

$$\xi_t(m) = \sum_{j=1}^{p_t(m)} \Phi_t^j(m)\xi_{t-j}(m) + \eta_t(m). \tag{6}$$

In this expression, $\eta_t, t = 1, \ldots, T$, are independent Gaussian vectors with $\mathbb{E}[\eta_t] = \mu_t$ and with an $M \times M$ covariance matrix $\Gamma_t := Cov(\eta_t)$. In particular, we will denote the standard deviation of $\eta_t(m)$ by $\sigma_t^\eta(m) > 0$.

In our development, we will need to express a given value of the process as a function of its past history. For this reason, it is convenient to introduce for $t = 1, \ldots, T-1$, $j = 1, \ldots, T-t$, and $m = 1, \ldots, M$, the integers

$$p_{t,j}^{\max}(m) = \max_{1 \le k \le j} \{p_{t+k}(m) - k\}, \tag{7}$$

as well as the useful past history of the process up to time step t

$$\tilde{\xi}_{[t]} = \left\{ \tilde{\xi}_{t-j}(m),\ m = 1, \ldots, M,\ j = 0, \ldots, \max(p_{t,T-t}^{\max}(m), t-1) \right\}. \tag{8}$$

The index $p_{t,j}^{\max}(m)$ specifies how much past information is needed at time step t to compute $\xi_{t+j}(m)$, for a process ξ_t modelled by (6). More precisely, as shown in [18], for $t = 1, \ldots, T-1$, $j = 1, \ldots, T-t$, and $m = 1, \ldots, M$, the relation

$$\xi_{t+j}(m) = \sum_{\ell=0}^{p_{t,j}^{\max}(m)} \alpha_{t,j}^{\ell}(m) \xi_{t-\ell}(m) + \sum_{\ell=1}^{j} \beta_{t+j}^{\ell}(m)\, \eta_{t+j-\ell+1}(m), \tag{9}$$

holds for certain coefficients $\alpha_{t,j}^{\ell}(m)$ and $\beta_{t+j}^{\ell}(m)$. Such coefficients can be derived from the model data in (6); we refer to [18] for the corresponding (recursive or explicit) formulæ, that are useful for coding the numerical implementation of the method.

3.2. **Rolling horizon versus non-rolling horizon policies.** We focus on models with recourse that, at time step t, make use of a recourse (or cost-to-go) function $\mathcal{Q}_{t+1}$. Usually, recourse functions depend only on the state variables, x_t. However, in our setting (6), the stochastic process is affinely dependent on previous values, and the state variable has to be augmented with the process history of realizations. As a result, the simulation phase, that is, the actual computation of the indicators from item `(EcInd)` in Section 2, involves solving (exactly or approximately) problems of the form

$$\begin{cases} \min\limits_{x_t, u_t} c_t^\top u_t + \mathcal{Q}_{t+1}(x_t, \tilde{\xi}_{[t]}) \\ x_t = A_{t-1} x_{t-1}(\tilde{\xi}_{[t-1]}) + B_t u_t + C_t \tilde{\xi}_t + d_t \\ E_t x_t + F_t(\tilde{\xi}_t) u_t \ge G_t \tilde{\xi}_t + h_t. \end{cases} \tag{10}$$

A usual non-rolling horizon approach approximates problem (10) by replacing $\mathcal{Q}_{t+1}(x_t, \tilde{\xi}_{[t]})$ with an approximate recourse function $\mathfrak{Q}_{t+1}(x_t, \tilde{\xi}_{[t]})$, depending on the augmented state $(x_t, \tilde{\xi}_{[t]})$. This approximation is built at the first stage, knowing $\tilde{\xi}_{[1]}$ and x_0 only.

By contrast, in a rolling horizon setting, either (a) we can determine an optimal solution to (10) or (b) we use an approximation $\mathfrak{Q}_{t+1}(\cdot, \cdot)$ of $\mathcal{Q}_{t+1}$ built at stage t, knowing the history $\tilde{\xi}_{[t]}$ and $x_{t-1}(\tilde{\xi}_{[t-1]})$.

Typically, we are in situation (a) when $\mathcal{Q}_{t+1}(x_t, \tilde{x}_{[t]})$ is given as the optimal value of a minimization problem and the corresponding minimization problem (10) is tractable. In this context, we can find an optimal solution to (10) even without an approximate representation of $\mathcal{Q}_{t+1}$. Such is the case of our rolling horizon policy (see details in Sections 3.3 and 4).

In situation (b), the algorithm used to obtain approximations of the recourse functions depends on their definitions. In a risk-neutral setting and for problems of the class we consider, SDDP can be used to obtain such approximations.

Table 1 which follows highlights the main differences of both approaches. This table shows that

TABLE 1. Usual non-rolling horizon (NRH) and rolling horizon (RH) approaches. In case (a), RH is a simulation phase whereas in case (b) it couples an optimization phase with a simulation phase.

NRH:	Optimization phase: for each $t = 1, \dots, T$, knowing the information $\tilde{\xi}_{[1]}$, compute an approximation $\mathfrak{Q}_{t+1}(\cdot, \cdot)$ of $\mathcal{Q}_{t+1}(\cdot, \cdot)$. Simulation phase: for each $i = 1, \dots, N$, for each $t = 1, \dots, T$, knowing the realization $\tilde{\xi}^i_{[t]} = (\tilde{\xi}_{[1]}, \tilde{\xi}^i_2, \dots, \tilde{\xi}^i_t)$, use the approximation $\mathfrak{Q}_{t+1}(\cdot, \tilde{\xi}^i_{[t]})$ to solve (10) and compute indicators.
RH:	Optimization/Simulation phase: for each $i = 1, \dots, N$, for each $t = 1, \dots, T$, knowing the realization $\tilde{\xi}^i_{[t]} = (\tilde{\xi}_{[1]}, \tilde{\xi}^i_2, \dots, \tilde{\xi}^i_t)$, either (a) solve directly (10) to optimality and compute indicators or (b) compute an approximation $\mathfrak{Q}_{t+1}(\cdot, \cdot)$ of $\mathcal{Q}_{t+1}(\cdot, \cdot)$ and use this approximation to solve (10) and compute indicators.

in case (a), the RH policy does not a priori provide us with approximations of the recourse functions. However, we will show in Section 5.2 that in the case of our RH policy, though we are in case (a), we can build cuts for the recourse functions as the Simulation process goes by. As a result, we also end up this process with approximate recourse functions $\mathfrak{Q}_{t+1}(\cdot, \cdot)$ available at the first stage. In this situation, we see that the simulation phase of our RH policy also contains an optimization phase (in the sense of NRH) building approximate recourse functions. As in the simulation phase of NRH, these approximate recourse functions can be used to define an alternative non-rolling horizon policy that will be denoted `RA-NRH` in the sequel. A motivation for using `RA-NRH` would be to simulate more quickly the policy that RH does. This could be possible if the approximate recourse functions are obtained using a set of $\tilde{N} << N$ scenarios.[2] However, recall that contrary to RH, `RA-NRH` does not find an optimal but an approximate solution to (10).

We see that with the usual non-rolling horizon approach, requirement `(RecFun)` is achieved in a preliminary Optimization phase while with our RH policy, approximations of the recourse functions are obtained as the simulation process goes by.

[2]Assume we have $q = q_t$ (constant in t) constraints (INEQ), and uncertainty in the right-hand side and chance constraints only. Then

- the simulation of RH policy involves solving for $t = 1, \dots, T$, N linear programs with $N_x + q(T - t + 1)$ constraints;
- the simulation of `RA-NRH` policy involves solving first for $t = 1, \dots, T$, $\tilde{N}$ linear programs with $N_x + q(T - t + 1)$ constraints and then NT linear programs with $N_x + q + \tilde{N}$ constraints.

3.3. **Building risk-averse rolling horizon policies.** If the recourse functions in (10) are risk-neutral or risk-averse, so will be the corresponding policies. Since in our motivating application the Independent System Operator is interested in finding policies that not only keep the reservoir levels above some critical value with high probability, but also avoid too big shortfalls, we use two constraints to hedge risk. More precisely, each uncertain scalar constraint is replaced by a chance constraint that prevents, with certain probability, constraint violation, and by a CVaR constraint, that prevents too big violations, if they occur.

For each $t = 1, \ldots, T$, our risk-averse recourse function $\mathcal{Q}_{t+1}$ in (10) is defined considering a t^{th} *risk-averse optimization problem*, written over the horizon $[t, T]$, considering that the past history $\tilde{\xi}_{[t]}$ is known. Accordingly, constraints at time step t are considered deterministic, while future constraints, for time steps $\tau = t+1, \ldots, T$, are dealt with as uncertain.

In this setting, given a scenario $(\tilde{\xi}_1, \ldots, \tilde{\xi}_T)$, and following [14, p. 16], given any random vector $Y(\xi_{t+1}, \ldots, \xi_\tau)$ decomposed in the form

$$Y(\xi_{t+1}, \ldots, \xi_\tau) = g_Y(\xi_{[t]}) + f_Y(\eta_{t+1}, \ldots, \eta_\tau),$$

with $\xi_{[t]}$ and $(\eta_{t+1}, \ldots, \eta_\tau)$ independent, we use the notation

$$Y(\xi_{t+1}, \ldots, \xi_\tau)|\tilde{\xi}_{[t]} := g_Y(\tilde{\xi}_{[t]}) + f_Y(\eta_{t+1}, \ldots, \eta_\tau)\,,$$

to refer to the corresponding conditional random variable.

For notational simplicity, and without loss of generality, from now on we suppose the inequality constraints in (1) are scalar, so matrices therein are row vectors (like each individual constraint in the water resource planning application from Section 2). In particular, in (2) we have that $q_t = 1$ and $F_t(\xi_t) = \mathring{F}_t + \xi_t^\top F_{t,1}^\top$.

The corresponding t^{th} risk-averse optimization problem, akin to problem (10), has the form

$$(11)\quad \begin{cases} \min\limits_{x_t, u_{[t,T]}} \sum\limits_{\tau=t}^{T} c_\tau^\top u_\tau \\ x_t = A_{t-1}x_{t-1}(\tilde{\xi}_{[t-1]}) + B_t u_t + C_t \tilde{\xi}_t + d_t, \\ E_t x_t + (\mathring{F}_t + \tilde{\xi}_t^\top F_{t,1}^\top)u_t \geq G_t \tilde{\xi}_t + h_t, \\ \text{and, for } \tau = t+1, \ldots, T: \\ \mathbb{P}\left(E_\tau x_\tau(x_t, u_{(t:\tau]}, \xi_{(t:\tau]}) + (\mathring{F}_\tau + \xi_\tau^\top F_{\tau,1}^\top)u_\tau - G_\tau \xi_\tau \geq h_\tau \middle| \tilde{\xi}_{[t]}\right) \geq 1 - \varepsilon_{\mathtt{p}}, \\ -CVaR_{\varepsilon_{\mathtt{p}}}\left(E_\tau x_\tau(x_t, u_{(t:\tau]}, \xi_{(t:\tau]}) + (\mathring{F}_\tau + \xi_\tau^\top F_{\tau,1}^\top)u_\tau - G_\tau \xi_\tau \middle| \tilde{\xi}_{[t]}\right) \geq h_\tau - \varepsilon_{\mathtt{c}}(|h_\tau| + 1), \end{cases}$$

where $\varepsilon_{\mathtt{p}},\ \varepsilon_{\mathtt{c}} \in (0, 1)$ are given confidence levels and where $x_\tau(x_t, u_{(t:\tau]}, \xi_{(t:\tau]})$ is the expression of x_τ as a function of variables $x_t, u_{(t:\tau]}$ and of random vectors $\xi_{(t:\tau]}$. This expression is obtained applying recursively transition equation (TRAN) between time steps $t+1$ and τ (see the next section).

Observe that variables for problem (11) are of the here-and-now type, i.e., they are fixed vectors. Having a solution $(x_t^{*t}, u_{[t,T]}^{*t})$ to (11), we only use its first components, (x_t^{*t}, u_t^{*t}), to define the policy. More precisely, after solving for $t = 1, \ldots, T$ all the risk-averse problems (11), the controls

$$u_{Rob} := (u_1^{*1}, \ldots, u_t^{*t}, \ldots, u_T^{*T})$$

give our rolling horizon implementable policy, which is also time consistent [31]. Note that the optimal values of controls $u_{t+1}, \dots, u_T$ in (11) are not used. Since in (11) the solution depends on $\tilde{\xi}_{[t]}$, but not on future realizations $\tilde{\xi}_{(t,T]}$, our policy is non-anticipative. Feasibility results from the fact that a solution for t^{th} risk-averse optimization problem (11) satisfies constraints (TRAN) and (INEQ) at time step t, by construction.

If uncertain inequality constraints are vectorial, each constraint component has individual chance and CVaR constraints and, hence, different confidence levels $\varepsilon_{\mathtt{p}}^i, \varepsilon_{\mathtt{c}}^i$, for $i = 1, \dots, q_t$, possibly varying with the time period $t = 1, \dots, T-1$.

On confidence levels for our RH policy. Confidence levels are parameters of our rolling horizon methodology and, as such, they need some *calibration*, to ensure that the risk-averse problems are feasible. These parameters have to be taken large enough to ensure feasibility, but small enough to guarantee constraint satisfaction with a reasonably large probability ($\varepsilon_{\mathtt{p}}$) and/or to reasonably limit the amount of constraint violation ($\varepsilon_{\mathtt{c}}$). There is not a unique setting of these parameters, since this is a problem-dependent issue. It is not possible to know in advance if, for a given choice of parameters, problems (11) will be feasible on a given scenario $\tilde{\xi}_{[1:T]}$, for all time steps t. As a result, instead of an a priori, static choice for these parameters, a dynamic tuning must be put in place. More precisely, on a given scenario and for a given time step, we can solve different problems (11), for a range of different confidence levels, increasing such levels as infeasibility arises. Once various values of $\varepsilon_{\mathtt{p}}$ and $\varepsilon_{\mathtt{c}}$ yielding feasible (11) are found, the combination with smaller $\varepsilon_{\mathtt{p}}$ should be preferred.

When uncertainty only appears in the right hand side as in our motivating application, this is not a real handicap, because each risk-averse problem (11) is a linear program (see the next section). The more general case, of uncertain matrices F_τ, may need a fine tuning for $\varepsilon_{\mathtt{p}} > \frac{1}{2}$ (as shown in the next section, for smaller $\varepsilon_{\mathtt{p}}$, the risk-averse problems are convex programs, and the trial and error process for different confidence levels can still be employed).

On confidence levels for `RA-NRH` policy. A dynamic choice for parameters $\varepsilon_{\mathtt{p}}$ and $\varepsilon_{\mathtt{c}}$ as explained in the previous paragraph, is only possible if we only wish to simulate our RH policy on a set of scenarios, and if we are not concerned about providing approximations of the recourse functions associated to our risk-averse policy. In the latter case, in particular to simulate `RA-NRH` policy, parameters $\varepsilon_{\mathtt{p}}$ and $\varepsilon_{\mathtt{c}}$ need to be fixed a priori, i.e., they can vary with constraints and time steps but for a fixed constraint and time step, they are constant along the N scenarios mentioned in Table 1.

4. Reformulating the chance and CVaR constraints

A usual concern when formulating robust or chance constraints is their tractability. We now develop some algebraic manipulations showing that our chance and CVaR constraints are conic quadratic at worst. Moreover, we show that when each $F_t(\xi_t) = \mathring{F}_t$ is deterministic (as in the water-resource planning application), problem (11) becomes a linear programming problem.

We start by applying the state transition equation (1)(TRAN) recursively to express a future state x_τ, for $\tau = t+1, \dots, T$, as a function

- of the current state x_t,
- of future controls $u_{(t,\tau]} = (u_{t+1}, \ldots, u_\tau)$, and
- of random vectors $\xi_{(t,\tau]} = (\xi_{t+1}, \ldots, \xi_\tau)$ from time step $t+1$ to time step τ.

Accordingly, the random variables in the future inequality constraints of (11) become

$$X(\xi_{(t,\tau]}) := E_\tau x_\tau(x_t, u_{(t,\tau]}, \xi_{(t,\tau]}) + (\mathring{F}_\tau + \xi_\tau^\top F_{\tau,1}^\top)u_\tau - G_\tau \xi_\tau \tag{12}$$

for $\tau = t+1, \ldots, T$. The probability and CVaR in (11) are computed with respect to the distribution of random vectors $\xi_{t+1}, \ldots, \xi_\tau$ given $\tilde{\xi}_{[t]}$, i.e., with respect to the distribution of random vectors $\xi_{t+1}|\tilde{\xi}_{[t]}, \ldots, \xi_\tau|\tilde{\xi}_{[t]}$. Keeping in mind the relation (9), the argument in both the probability and CVaR depends on the random vectors $\xi_{[t]}$ and $(\eta_{t+1}, \ldots, \eta_\tau)$. Since the history $\tilde{\xi}_{[t]}$ is known, both the conditional probability and the CVaR in (11) are finally computed with respect to the distribution of $(\eta_{t+1}, \ldots, \eta_\tau)$.

We start with some technical relations that will be useful in the sequel.

Lemma 4.1 (Affine dependence on decision variables and random process)**.** *Consider the scalar random variable $X = X(\xi_{(t,\tau]})$ defined in* (12) *using the transition equation* (1)(TRAN)*. For all $\tau = t+1, \ldots, T$ and each time step $t = 1, \ldots, T-1$, X is an affine function of $\xi_{(t,\tau]}$:*

$$X = \sum_{j=t+1}^{\tau} \nu_{j,\tau} \xi_j + \nu_{t,\tau}\,,$$

where the coefficients have the expression

$$\nu_{j,\tau} = \begin{cases} E_\tau(\prod_{k=t}^{\tau-1} A_k)x_t + \sum_{i=t+1}^{\tau} E_\tau(\prod_{k=i}^{\tau-1} A_k)(B_i u_i + d_i) + \mathring{F}_\tau u_\tau & \in \mathbb{R} \qquad \text{for } j = t \\ E_\tau(\prod_{k=j}^{\tau-1} A_k)C_j & \in \mathbb{R}^{1,M} \quad \text{for } j = t+1, \ldots, \tau-1 \\ E_\tau C_\tau + u_\tau^\top F_{\tau,1} - G_\tau & \in \mathbb{R}^{1,M} \quad \text{for } j = \tau\,. \end{cases}$$

Proof. Take fixed $t \in \{1, \ldots, T-1\}$ and $\tau \in \{t+1, \ldots, T\}$. A recursive application of the transition equation (1)(TRAN) yields for every $\ell \in \{1, \ldots, \tau\}$:

$$x_\tau = \prod_{k=\ell-1}^{\tau-1} A_k x_{\ell-1} + \sum_{j=\ell}^{\tau} (\prod_{k=j}^{\tau-1} A_k)\left[B_j u_j + C_j \xi_j + d_j\right].$$

Now for $j = 0, \ldots, \tau$, we let $\mathcal{A}_{j,\tau} = \prod_{k=j}^{\tau-1} A_k$ so that $\mathcal{A}_{\tau,\tau} = I$, the identity matrix. With this notation, the relation

$$x_\tau = \mathcal{A}_{\ell-1,\tau} x_{\ell-1} + \sum_{j=\ell}^{\tau} \mathcal{B}_{j,\tau} u_j + \sum_{j=\ell}^{\tau} \mathcal{C}_{j,\tau} \xi_j + \mathcal{d}_{\ell,\tau} \tag{13}$$

holds with

$$\mathcal{B}_{j,\tau} = \mathcal{A}_{j,\tau} B_j\,, \quad \mathcal{C}_{j,\tau} = \mathcal{A}_{j,\tau} C_j\,, \quad \text{and} \quad \mathcal{d}_{\ell,\tau} = \sum_{j=\ell}^{\tau} \mathcal{A}_{j,\tau} d_j\,.$$

Since $\xi_\tau^\top F_{\tau,1}^\top u_\tau = u_\tau^\top F_{\tau,1} \xi_\tau$, in (12) we have that

$$X = E_\tau x_\tau + \mathring{F}_\tau u_\tau + (u_\tau^\top F_{\tau,1} - G_\tau)\xi_\tau\,.$$

When using this identity and relation (13) for x_τ, written with $\ell = t+1$, we obtain that

$$\begin{aligned} X &= E_\tau\Big(\mathcal{A}_{t,\tau}x_t + \sum_{j=t+1}^{\tau} \mathcal{B}_{j,\tau}u_j + \mathcal{d}_{t+1,\tau}\Big) + \mathring{F}_\tau u_\tau \\ &\quad + \sum_{j=t+1}^{\tau} E_\tau \mathcal{C}_{j,\tau}\xi_j + (u_\tau^\top F_{\tau,1} - G_\tau)\xi_\tau, \end{aligned}$$

where we gathered together all terms depending on $\xi_{(t,\tau]}$. The desired results follow, by associating $\nu_{t,\tau}$ with the first four terms on the right-hand side above, each $\nu_{j,\tau}$, $j = t+1,\ldots,\tau-1$, equal to the corresponding term in the fifth summation, and $\nu_{\tau,\tau}$ as the remaining portion of the expression. □

From Lemma 4.1 we see that the dependence of coefficients ν on variables $(x_t, u_{(t,\tau]})$ is the following:

- $\nu_{t,\tau}$ is an affine scalar function of variables $(x_t, u_{(t,\tau]})$,
- $\nu_{j,\tau}$ is a constant M-dimensional row vector for all $j = t+1,\ldots,\tau-1$, and
- $\nu_{\tau,\tau}$ is an M-dimensional row vector, affine on u_τ if and only if $F_{\tau,1} \neq 0$, constant otherwise.

Since the relation in Lemma 4.1 is affine, the variable X from (12) is also affinely dependent on the decision variables in (11): $X = X(x_t, u_{(t,\tau]}, \xi_{(t,\tau]})$. In order to write down the chance and CVaR constraints in (11), we need to explicitly compute, for fixed $(x_t, u_{(t,\tau]})$, the mean and standard deviation of X, conditioned to $\tilde{\xi}_{[t]}$, i.e., knowing the history of realizations until time step t.

Lemma 4.2 (Conditional mean and standard deviation)**.** *Consider the scalar random variable X defined in* (12) *using the transition equation* (1)(TRAN)*, for a discrete time process ξ_t modelled by* (6)*. Consider the coefficients defined in Lemma 4.1, and let $\nu_{j,\tau}(m)$, $m = 1,\ldots,M$, denote the components of vectors $\nu_{j,\tau}$, for $j = t+1,\ldots,\tau$.*

Then for all $\tau = t+1,\ldots,T$, and $t = 1,\ldots,T-1$, the random variable $X|\tilde{\xi}_{[t]}$ is Gaussian with conditional mean

$$\mathbb{E}[X|\tilde{\xi}_{[t]}] = \nu_{t,\tau} + \sum_{j=1}^{\tau-t}\sum_{m=1}^{M} \nu_{t+j,\tau}(m)\mathbb{E}[\xi_{t+j}(m)|\tilde{\xi}_{[t]}]$$

where

$$\mathbb{E}[\xi_{t+j}(m)|\tilde{\xi}_{[t]}] = \sum_{\ell=0}^{p_{t,j}^{\max}(m)} \alpha_{t,j}^{\ell}(m)\tilde{\xi}_{t-\ell}(m) + \sum_{\ell=1}^{j} \beta_{t+j}^{\ell}(m)\,\mu_{t+j-\ell+1}(m)\,.$$

As for the conditional standard deviation, it is given by

$$\sigma(X|\tilde{\xi}_{[t]}) = \sigma(\sum_{j=1}^{\tau-t}\sum_{m=1}^{M} \nu_{t+j,\tau}(m)\xi_{t+j}(m)|\tilde{\xi}_{[t]}) = \sqrt{\sum_{\ell=1}^{\tau-t} \gamma_{t,\tau,\ell}^\top \Gamma_{t+\ell}\gamma_{t,\tau,\ell}}\,, \tag{14}$$

where $\Gamma_{t+\ell} = Cov(\eta_{t+\ell})$ and the M-dimensional vectors $\gamma_{t,\tau,\ell}$ have components

$$\gamma_{t,\tau,\ell}(m) := \sum_{j=\ell}^{\tau-t} \nu_{t+j,\tau}(m)\beta_{t+j}^{j-\ell+1}(m) \tag{15}$$

that are affine functions of u_τ if and only if $F_{\tau,1} \neq 0$, and are constant otherwise.
As a result, the conditional mean is an affine function of variables $(x_t, u_{(t,\tau]})$, while the conditional standard deviation is constant if and only if $F_{\tau,1} = 0$, otherwise it is a conic quadratic function of u_τ of the form $\|\mathcal{L}_{t,\tau}u_\tau + \ell_{t,\tau}\|_2$ for some matrix $\mathcal{L}_{t,\tau}$ and vector $\ell_{t,\tau}$ of appropriate dimensions.

Proof. Combining relation (9) with the expression for X in Lemma 4.1, we see that

$$X = \nu_{t,\tau} + \sum_{j=1}^{\tau-t}\sum_{m=1}^{M} \nu_{t+j,\tau}(m) \left(\sum_{\ell=0}^{p_{t,j}^{\max}(m)} \alpha_{t,j}^{\ell}(m)\xi_{t-\ell}(m) + \sum_{\ell=1}^{j} \beta_{t+j}^{\ell}(m)\eta_{t+j-\ell+1}(m) \right). \tag{16}$$

The desired result for $\mathbb{E}[X|\tilde{\xi}_{[t]}]$ follows, by taking the conditional expectation in both members above. Such expression depends affinely on $\nu_{t,\tau}, \ldots, \nu_{\tau,\tau}$ and, in turn, these coefficients depend affinely on $(x_t, u_{(t,\tau]})$, by Lemma 4.1.

Next, note that each component of $\gamma_{t,\tau,\ell}$ from (15) depends affinely on coefficients $\nu_{t+j,\tau}(m)$, and satisfies the relation

$$\sum_{j=1}^{\tau-t}\sum_{m=1}^{M} \nu_{t+j,\tau}(m) \sum_{\ell=1}^{j} \beta_{t+j}^{\ell}(m)\eta_{t+j-\ell+1}(m) = \sum_{\ell=1}^{\tau-t} \eta_{t+\ell}^{\top}\gamma_{t,\tau,\ell}; \tag{17}$$

with independent random vectors $\eta_{t+1}, \ldots, \eta_\tau$. The identity (14) follows, by plugging (17) into (16). Note in particular, that the expression for the standard deviation is a conic quadratic function of $\nu_{t+1,\tau}, \ldots, \nu_{\tau,\tau}$. By Lemma 4.1, coefficients $\nu_{t+1,\tau}(m), \ldots, \nu_{\tau-1,\tau}(m)$ are always constant, while $\nu_{\tau,\tau}(m)$ is either constant, or affinely depending on u_τ if and only if $F_{\tau,1} \neq 0$, as stated.

Finally, because the expression obtained for $X|\tilde{\xi}_{[t]}$ when using (17) in (16) is an affine combination of Gaussian random variables, the random variable $X|\tilde{\xi}_{[t]}$ is Gaussian too. □

Before showing that (11) is a tractable linear or conic quadratic program, we give an explicit equivalence between VaR and CVaR for Gaussian random variables; see [14, Ex. 6.3].

Lemma 4.3. *Let $X \sim \mathcal{N}(m, \sigma^2)$ be a Gaussian random variable. Then for any $\varepsilon_{\mathtt{p}} \in [0,1]$ we have $CVaR_{\varepsilon_{\mathtt{p}}}(X) = VaR_{\varphi(\varepsilon_{\mathtt{p}})}(X)$ where the (bijective) function $\varphi : [0,1] \to [0, \frac{1}{2}]$ is given by $\varphi(x) = 1 - F\left(\frac{\exp(-(F^{-1}(1-x))^2/2)}{\sqrt{2\pi}x}\right)$ for $x \in (0,1]$ and $\varphi(0) = 0$.*

Proof. The relation is trivial if $\varepsilon_{\mathtt{p}} = 1$ or $\varepsilon_{\mathtt{p}} = 0$. When $\varepsilon_{\mathtt{p}} \in (0,1)$, first note that $VaR_{\varepsilon_{\mathtt{p}}}(X) = -F_X^{-1}(\varepsilon_{\mathtt{p}}) = -m - F^{-1}(\varepsilon_{\mathtt{p}})\sigma = -m + F^{-1}(1-\varepsilon_{\mathtt{p}})\sigma$. Next, the algebraic manipulations below

$$\begin{aligned}
CVaR_{\varepsilon_{\mathtt{p}}}(X) &= -\mathbb{E}[X|X \leq F_X^{-1}(\varepsilon_{\mathtt{p}})] = -\frac{1}{\varepsilon_{\mathtt{p}}} \int_{-\infty}^{m+F^{-1}(\varepsilon_{\mathtt{p}})\sigma} x f_X(x) dx \\
&= -\frac{1}{\varepsilon_{\mathtt{p}}\sqrt{2\pi}\sigma} \int_{-\infty}^{m+F^{-1}(\varepsilon_{\mathtt{p}})\sigma} x \exp\left(-\frac{(x-m)^2}{2\sigma^2}\right) dx \\
&= -\frac{m}{\varepsilon_{\mathtt{p}}} \int_{-\infty}^{F_X^{-1}(\varepsilon_{\mathtt{p}})} f_X(x) dx - \frac{\sigma}{\varepsilon_{\mathtt{p}}\sqrt{2\pi}} \int_{-\infty}^{m+F^{-1}(\varepsilon_{\mathtt{p}})\sigma} \left(\frac{x-m}{\sigma^2}\right) \exp\left(-\frac{(x-m)^2}{2\sigma^2}\right) dx \\
&= -m + \frac{\sigma}{\varepsilon_{\mathtt{p}}\sqrt{2\pi}} \exp\left(-(F^{-1}(\varepsilon_{\mathtt{p}}))^2/2\right),
\end{aligned}$$

give the desired result, recalling that $F^{-1}(\varepsilon_{\mathtt{p}}) = -F^{-1}(1-\varepsilon_{\mathtt{p}})$. □

The statement that a chance constraint for an affine relation of a Gaussian random variable is tractable, and reduces to a conic quadratic constraint, is already known; see [8], [9], [10]. We now extend this result to a more general setting, and, more importantly, give a **completely explicit** formulation for problem (11), by means of (9), when each component $\xi_t(m),\ m = 1, \ldots, M$, is a

generalized autoregressive process with time varying order and possibly correlated noise components.

Theorem 4.4 (Tractability of problem (11))**.** *Let the scalar random variable X be defined in* (12) *using the transition equation* (1)(TRAN)*, for a discrete time process ξ_t modelled by* (6) *and let $\varphi(\cdot)$ be the bijection defined in Lemma 4.3. Then each pair of future constraints in* (11)

$$\text{(18)} \qquad \left\{ \begin{array}{rcl} \mathbb{P}\left(X \geq h_\tau \middle| \tilde{\xi}_{[t]}\right) & \geq & 1-\varepsilon_{\mathtt{p}}, \\ -CVaR_{\varepsilon_{\mathtt{p}}}\left(X \middle| \tilde{\xi}_{[t]}\right) & \geq & h_\tau - \varepsilon_{\mathtt{c}}(|h_\tau|+1), \end{array} \right.$$

is equivalent to any of the three representations below:

Robust formulation
$$\left\{ \begin{array}{rcl} \mathbb{E}[X|\tilde{\xi}_{[t]}] - F^{-1}(1-\varepsilon_{\mathtt{p}})\sigma(X|\tilde{\xi}_{[t]}) & \geq & h_\tau, \\ \mathbb{E}[X|\tilde{\xi}_{[t]}] - F^{-1}(1-\varphi(\varepsilon_{\mathtt{p}}))\sigma(X|\tilde{\xi}_{[t]}) & \geq & h_\tau - \varepsilon_{\mathtt{c}}(|h_\tau|+1), \end{array} \right.$$
where the explicit expressions for the conditional mean and standard deviation are given in Lemma 4.2;

VaR formulation
$$\left\{ \begin{array}{rcl} -VaR_{\varepsilon_{\mathtt{p}}}(X|\tilde{\xi}_{[t]}) & \geq & h_\tau, \\ -VaR_{\varphi(\varepsilon_{\mathtt{p}})}(X|\tilde{\xi}_{[t]}) & \geq & h_\tau - \varepsilon_{\mathtt{c}}(|h_\tau|+1)\,; \end{array} \right.$$
and

CVaR formulation
$$\left\{ \begin{array}{rcl} -CVaR_{\varphi^{-1}(\varepsilon_{\mathtt{p}})}(X|\tilde{\xi}_{[t]}) & \geq & h_\tau, \\ -CVaR_{\varepsilon_{\mathtt{p}}}(X|\tilde{\xi}_{[t]}) & \geq & h_\tau - \varepsilon_{\mathtt{c}}(|h_\tau|+1)\,. \end{array} \right.$$

Therefore, if $\varepsilon_{\mathtt{p}} \in (0, \frac{1}{2})$, future constraints in (11) *are conic quadratic and* (11) *is a conic quadratic program. Since future constraints* (18) *are affine in variables $(x_t, u_{(t,\tau]})$ if and only if $F_{\tau,1} = 0$, when $F_{\tau,1} = 0$ for every τ then* (11) *is a linear program for every $\varepsilon_{\mathtt{p}} \in (0,1)$.*

Proof. For a Gaussian random variable Z with expectation $\mathbb{E}[Z]$ and standard deviation $\sigma(Z)$, from the proof of Lemma 4.3 we obtain

$$\mathbb{P}(Z \geq \underline{\mathtt{Z}}) \geq 1-\varepsilon_{\mathtt{p}} \Leftrightarrow \underline{\mathtt{Z}} \leq \mathbb{E}[Z] - F^{-1}(1-\varepsilon_{\mathtt{p}})\sigma(Z) \Leftrightarrow \underline{\mathtt{Z}} \leq -VaR_{\varepsilon_{\mathtt{p}}}(Z).$$

Since the random variable $Z = X|\tilde{\xi}_{[t]}$ is Gaussian, the VaR and CVaR formulations follow from Lemma 4.3 and from the equivalences above with $\underline{\mathtt{Z}} = h_\tau$. In addition, the left-hand side equivalence yields the following expression:

$$\begin{array}{rcl} \mathbb{E}[X|\tilde{\xi}_{[t]}] - F^{-1}(1-\varepsilon_{\mathtt{p}})\sigma(X|\tilde{\xi}_{[t]}) & \geq & h_\tau, \\ \mathbb{E}[X|\tilde{\xi}_{[t]}] - F^{-1}(1-\varphi(\varepsilon_{\mathtt{p}}))\sigma(X|\tilde{\xi}_{[t]}) & \geq & h_\tau - \varepsilon_{\mathtt{c}}(|h_\tau|+1). \end{array}$$

Next, observe that when $\varepsilon_{\mathtt{p}} \in (0, \frac{1}{2})$ we have $\varphi(\varepsilon_{\mathtt{p}}) \in (0, \frac{1}{2})$, and $F^{-1}(1-\varepsilon_{\mathtt{p}})$ and $F^{-1}(1-\varphi(\varepsilon_{\mathtt{p}}))$ are both positive. Using this observation and the last assertion in Lemma 4.2, it follows that problem (11) is a conic quadratic program when $\varepsilon_{\mathtt{p}} \in (0, \frac{1}{2})$ or a linear program if $F_{\tau,1} = 0$ for every τ. □

If in (11) some element of A_t, C_t, E_t, or G_t is no longer deterministic, but uncertain and, like F_t, depending affinely on ξ_t, then both Lemma 4.2 and Theorem 4.4 fail to hold because the product of two Gaussian random variables is not a Gaussian random variable.

In Theorem 4.4, the fact that the random variable X is Gaussian leads to the equivalent VaR and CVaR formulations, by Lemma 4.3. Explicit expressions for the chance and CVaR constraints would still be available for any distribution of noises η_t, as long as the distribution of any linear combination of noises $\eta_t, t = 1, \ldots, T$, is known. For some processes, log-normal distributions are preferred to normal ones. Such is the case of our initial motivation, for which streamflows

must remain nonnegative for all realizations. In this case, one needs to consider approximations of the generalized inverse of the cumulative distribution function of the (no longer normal) random variable $X|\tilde{\xi}_{[t]}$; see [32] and [1]. By suitably adapting Theorem 4.4, we would once more obtain a deterministic (linear or conic quadratic) program, that *approximates* problem (11).

5. Defining robust recourse functions

Even if robust optimization usually refers to worst-case oriented methodologies, our chance and CVaR constraints can also be considered as a robust way of dealing with uncertainty (in the sense of [3]), because noises η_t are Gaussian. More precisely, the probabilistic constraint is nothing but its robust version, taking as uncertainty set for η_t the ellipsoid

$$\{x_t \ : \ (x_t-\mu_t)^\top \Gamma_t^{-1}(x_t-\mu_t) \leq (F^{-1}(1-\varepsilon_{\mathtt{p}}))^2\}\,. \tag{19}$$

Similarly, the robustified version of the constraint $X|\tilde{\xi}_{[t]} \geq h_\tau - \varepsilon_{\mathtt{c}}(|h_\tau|+1)$ for $\tau = t+1,\ldots,T$, using for η_ℓ ($\ell = t+1,\ldots,\tau$) the uncertainty ellipsoidal set

$$\{x_\ell \ : \ (x_\ell-\mu_\ell)^\top \Gamma_\ell^{-1}(x_\ell-\mu_\ell) \leq (F^{-1}(1-\varphi(\varepsilon_{\mathtt{p}})))^2\}\,, \tag{20}$$

coincides with the CVaR constraint in (11). These relations justify the naming "robust" for the recourse function $\mathcal{Q}_{t+1}(x_t,\tilde{\xi}_{[t]})$ given by

$$\left\{\begin{array}{l} \min\limits_{u_{[t+1,T]}} \sum\limits_{\tau=t+1}^{T} c_\tau^\top u_\tau \\ \text{s.t., for } \tau = t+1,\ldots,T: \\ \mathbb{P}\left(E_\tau x_\tau(x_t,u_{(t:\tau]},\xi_{(t:\tau]}) + (\mathring{F}_\tau + \xi_\tau^\top F_{\tau,1}^\top)u_\tau - G_\tau\xi_\tau \geq h_\tau \Big| \tilde{\xi}_{[t]}\right) \geq 1-\varepsilon_{\mathtt{p}}, \\ -CVaR_{\varepsilon_{\mathtt{p}}}\left(E_\tau x_\tau(x_t,u_{(t:\tau]},\xi_{(t:\tau]}) + (\mathring{F}_\tau + \xi_\tau^\top F_{\tau,1}^\top)u_\tau - G_\tau\xi_\tau \Big| \tilde{\xi}_{[t]}\right) \geq h_\tau - \varepsilon_{\mathtt{c}}(|h_\tau|+1), \end{array}\right. \tag{21}$$

representing the cost-to-go function in (11). We now intend to describe two algorithms to obtain approximations of these recourse functions (requirement `(RecFun)` in Section 2). To this end, we start in the next section with a useful rewriting of problem (11).

5.1. Stagewise decomposition of (11).

Our goal is to write dynamic programming equations for (11), by decomposing its feasible set by stages. To explain this rewriting, we go back to our motivating application. In this context, to work out the calculations, and without loss of generality, consider in (11) a feasible set of the form

$$\mathcal{S} := \left\{\begin{array}{ll} (x_t,(g_\tau,df_\tau)_{\tau=t}^T): & \\ x_t = x_{t-1} - g_t + \gamma_t\tilde{\xi}_t,\ g_t \geq 0,\ df_t \geq 0 & \\ g_t + df_t \geq \mathrm{dem}_t - (1-\gamma_t)\tilde{\xi}_t & \\ x_t \geq x_t^{crit} & \\ \text{and, for } \tau = t+1,\ldots,T: & \\ g_\tau \geq 0,\ df_\tau \geq 0 & \\ -CVaR_{\varepsilon_{\mathtt{p}}}\left(x_\tau(x_t,g_{(t:\tau]},\xi_{(t:\tau]})|\tilde{\xi}_{[t]}\right) \geq x_\tau^{crit}(1-\varepsilon_{\mathtt{c}}) - \varepsilon_{\mathtt{c}} & (a) \\ \mathbb{P}\left(g_\tau + df_\tau \geq \mathrm{dem}_\tau - (1-\gamma_\tau)\xi_\tau|\tilde{\xi}_{[t]}\right) \geq 1-\varepsilon_{\mathtt{p}} & (b) \end{array}\right\},$$

with only one type of risk-averse formulation for each future inequality constraint, and noting that the set depends on x_{t-1} and on the realization history: $\mathcal{S} = \mathcal{S}(x_{t-1}, \tilde{\xi}_{[t]})$.

The first important observation is that, by taking the conditional expectation in the (stochastic) flow balance transition equation for the reservoir, we obtain a (deterministic) transition equation for the mean reservoir volumes:

$$\bar{x}_\tau := \mathbb{E}[x_\tau|\tilde{\xi}_{[t]}] \quad \text{satisfies} \quad \bar{x}_\tau = \left[\begin{array}{ll} x_{t-1} - g_t + \gamma_t \tilde{\xi}_t & \tau = t \\ \bar{x}_{\tau-1} - g_\tau + \gamma_\tau \mathbb{E}[\xi_\tau|\tilde{\xi}_{[t]}] & \tau = t+1, \ldots, T\,. \end{array} \right.$$

In the above relations, we have used that decision variables g_τ are of the here-and-now type for (11), i.e., they are deterministic.

With these new variables $\bar{x}_\tau$, the robust formulation in Theorem 4.4 for the critical volume (affine) constraints (a) is written down as follows:

$$\bar{x}_\tau \geq x_\tau^{R\ crit} := x_\tau^{crit}(1 - \varepsilon_{\mathtt{c}}) - \varepsilon_{\mathtt{c}} + F^{-1}(1 - \varphi(\varepsilon_{\mathtt{p}}))\sigma(x_\tau|\tilde{\xi}_{[t]})\,.$$

The explicit expression for the conditional standard deviation above is obtained from (12), written with all the matrices therein null, except for E_τ, set to 1. The corresponding affine parameters in Lemma 4.1 are $\nu_{j,\tau} = \gamma_j$ for $j = t+1, \ldots, \tau$, so that

$$x_\tau = x_t - \sum_{j=t+1}^{\tau} g_j + \sum_{j=t+1}^{\tau} \gamma_j \xi_j \text{ and, hence, } \sigma(x_\tau|\tilde{\xi}_{[t]}) = \sigma(\sum_{j=t+1}^{\tau} \gamma_j \xi_j|\tilde{\xi}_{[t]})\,,$$

which can be computed using (9).

A similar reasoning, applying Theorem 4.4 to affine constraints (b), gives the following reformulation for future demand constraints:

$$g_\tau + df_\tau \geq \mathrm{dem}_\tau^R - (1 - \gamma_\tau)\mathbb{E}[\xi_\tau|\tilde{\xi}_{[t]}] \text{ where } \mathrm{dem}_\tau^R := \mathrm{dem}_\tau + (1 - \gamma_\tau)F^{-1}(1 - \varepsilon_{\mathtt{p}})\sigma(\xi_\tau|\tilde{\xi}_{[t]})\,.$$

In order to obtain a stagewise decomposition of the feasible set in (11), we use the subset

$$\mathcal{S}_t^t(\bar{x}_{t-1}) := \left\{ (\bar{x}_t, g_t, df_t) : \begin{array}{l} \bar{x}_t = \bar{x}_{t-1} - g_t + \gamma_t \tilde{\xi}_t,\ g_t \geq 0,\ df_t \geq 0 \\ g_t + df_t \geq \mathrm{dem}_t - (1 - \gamma_t)\tilde{\xi}_t \\ \bar{x}_t \geq x_t^{crit} \end{array} \right\}$$

and, for $\tau = t+1, \ldots, T$, the subsets $\mathcal{S}_\tau^t(\bar{x}_{\tau-1})$, given by

$$\left\{ (\bar{x}_\tau, g_\tau, df_\tau) : \begin{array}{l} \bar{x}_\tau = \bar{x}_{\tau-1} - g_\tau + \gamma_\tau \mathbb{E}[\xi_\tau|\tilde{\xi}_{[t]}],\ g_\tau \geq 0,\ df_\tau \geq 0 \\ g_\tau + df_\tau \geq \mathrm{dem}_\tau^R - (1 - \gamma_\tau)\mathbb{E}[\xi_\tau|\tilde{\xi}_{[t]}] \\ \bar{x}_\tau \geq x_\tau^{R\ crit} \end{array} \right\}.$$

For our motivating application, the objective function is also separable by stages, making problem (11) solvable by dynamic programming, introducing the functions

$$\mathcal{Q}_\tau^t(\bar{x}_{\tau-1}, \tilde{\xi}_{[t]}) = \min \left\{ c_\tau^\top u_\tau + \mathcal{Q}_{\tau+1}^t(\bar{x}_\tau, \tilde{\xi}_{[t]}) : (\bar{x}_\tau, u_\tau) \in \mathcal{S}_\tau^t(\bar{x}_{\tau-1}) \right\}, \tag{22}$$

where we used the short notation $u_\tau = (g_\tau, df_\tau)$ for the controls. The functions above are defined recursively for $\tau = T, T-1, \ldots, t$, starting from $\mathcal{Q}_{T+1}^t(\cdot) \equiv 0$. The super index t in the notation reflects the fact that functions $\mathcal{Q}_\tau^t$ correspond to the stagewise decomposition of the t^{th} risk-averse optimization problem.

Following the developments above, risk-averse problem (11) written for a given history $\tilde{\xi}^i_{[t]}$ has the following stagewise decomposition:

$$
\left\{
\begin{array}{ll}
\min\limits_{\bar{x}_{[t:T]},u_{[t:T]}} \sum\limits_{\tau=t}^{T} c_\tau^\top u_\tau & \\
\bar{x}_\tau = A_{\tau-1}\bar{x}_{\tau-1} + B_\tau u_\tau + C_\tau \mathbb{E}[\xi_\tau|\tilde{\xi}^i_{[t]}] + d_\tau & (\pi^{t,i}_\tau) \\
E_\tau \bar{x}_\tau + \mathring{F}_\tau^\top u_\tau \geq G_\tau \mathbb{E}[\xi_\tau|\tilde{\xi}^i_{[t]}] + h^t_\tau & (\lambda^{t,i}_\tau) \\
\text{for } \tau = t, t+1, \ldots, T, &
\end{array}
\right. \tag{23}
$$

where

$$
h^t_\tau := h_\tau + \max\left(F^{-1}(1-\varepsilon_{\mathtt{p}})\sigma(X|\tilde{\xi}^i_{[t]}), -\varepsilon_{\mathtt{c}}(|h_\tau|+1) + F^{-1}(1-\varphi(\varepsilon_{\mathtt{p}}))\sigma(X|\tilde{\xi}^i_{[t]})\right). \tag{24}
$$

From the above decomposition of the risk-averse problems, we can make two interesting observations:

(A) **Interpretation of risk-averse problem** (11)**.** The above representation (23) shows that problem (11) is obtained from a problem instance replacing
- future realizations $\tilde{\xi}_\tau$ by their conditional means $\mathbb{E}[\xi_\tau|\tilde{\xi}_{[t]}]$ and
- right-hand sides h_τ by some "robust" counterparts h^t_τ strictly greater than h_τ.

In the case of the hydro-thermal problem from Section 2, looking at subsets $\mathcal{S}^t_\tau(\bar{x}_{\tau-1})$ above, this interpretation can be specialized as follows: each risk averse problem is obtained replacing in a problem instance
- the inflows by their conditional means, given the history $\tilde{\xi}_{[t]}$;
- the nominal demands dem_τ by "robust" demands dem^R_τ which are strictly greater than the nominal demands;
- the critical levels x^{crit}_τ by "robust" critical levels $x^{R\ crit}_\tau$ which, for $\varepsilon_{\mathtt{c}}$ sufficiently small, are strictly greater than the critical levels.

(B) **Solution method for** (11)**.** Since we can write (deterministic) dynamic programming equations for our risk-averse problems, each one of these problems can be solved efficiently by Dual Dynamic Programming (DDP). This can be interesting when T is large. Moreover, as explained in the next section, using DDP to solve (11), we can build approximations for our recourse functions.

To finish, we mention that it is still possible to write down for (11) Dynamic Programming relations similar to the above even if $F_{\tau,i} \neq 0$ for some $\tau \in \{t+1, \ldots, T\}$ and $i \leq q_\tau$.

5.2. **Approximate recourse functions.** We now give two algorithms to approximate the recourse functions (21) associated to our model. For both of them, the first step consists in sampling N scenarios $(\tilde{\xi}^i_2, \ldots, \tilde{\xi}^i_T)$, $i = 1, \ldots, N$, for $(\xi_2, \ldots, \xi_T)$ (recall that realization $\tilde{\xi}_1$ is available at $t = 1$). Then our rolling horizon policy is used on each scenario. The first algorithm solves directly each risk-averse problem and builds for each recourse function one cut per scenario using a dual solution. The second algorithm constructs various cuts per scenario for each recourse function solving each risk-averse problem by DDP.

First algorithm: using a primal and a dual solution to each risk-averse problem. For scenario i and time step t, risk-averse problem (11) is solved with $\tilde{\xi}_{[t]}$ replaced by $\tilde{\xi}^i_{[t]}$. Since

this risk-averse problem can be written as (23), for stage t the recourse function $\mathcal{Q}_{t+1}$, defined in (21), has the expression

$$\mathcal{Q}_{t+1}(\bar{x}_t, \tilde{\xi}_{[t]}) := \left\{ \begin{array}{l} \min\limits_{\bar{x}_{[t+1:T]}, u_{[t+1:T]}} \sum\limits_{\tau=t+1}^{T} c_\tau^\top u_\tau \\ \bar{x}_\tau = A_{\tau-1}\bar{x}_{\tau-1} + B_\tau u_\tau + C_\tau \mathbb{E}[\xi_\tau|\tilde{\xi}_{[t]}] + d_\tau,\ \tau = t+1, \ldots, T, \\ E_\tau \bar{x}_\tau + \mathring{F}_\tau^\top u_\tau \geq G_\tau \mathbb{E}[\xi_\tau|\tilde{\xi}_{[t]}] + h_\tau^t,\ \tau = t+1, \ldots, T. \end{array} \right. \tag{25}$$

Using Lemmas 4.1 and 4.2 and the above representation, we easily check that these recourse functions are convex (see the proof of Proposition 5.1). As a result, letting $\bar{x}_t^i$ be an optimal solution of $\bar{x}_t$ in (23), it is possible to define a cut (an hyperplane lying below the function) for $\mathcal{Q}_{t+1}$ at $(\bar{x}_t^i, \tilde{\xi}_{[t]}^i)$. Collecting these cuts, after a simulation phase over a set of i scenarios, we have for $\mathcal{Q}_{t+1}$ the approximation

$$\mathfrak{Q}_{t+1}^i(\bar{x}_t, \tilde{\xi}_{[t]}) = \max_{1 \leq k \leq i} \ \bar{x}_t^\top a_t^k + \tilde{\xi}_{[t]}^\top b_t^k + c_t^k$$

for each $t = 1, 2, \ldots, T-1,$, whereas $\mathfrak{Q}_{T+1}^i = \mathcal{Q}_{T+1} \equiv 0$. The formulas for cut coefficients (a_t^k, b_t^k, c_t^k) are given in the following proposition:

Proposition 5.1 (Cuts computation for robust recourse functions)**.** *Consider risk averse problem* (23) *for some* $t \in \{1, \ldots, T-1\}$ *and let* $(\pi_t^{t,i}, \pi_{t+1}^{t,i}, \ldots, \pi_T^{t,i}, \lambda_t^{t,i}, \lambda_{t+1}^{t,i}, \ldots, \lambda_T^{t,i})$ *be a dual solution of this problem (see* (23)*). Let* $\mathcal{Q}_{t+1}$ *be the corresponding recourse function defined in* (25)*. Then valid cuts for* $\mathcal{Q}_{t+1}$ *are given by*

$$\begin{array}{l} a_t^i = A_t^\top \pi_{t+1}^{t,i}, \\ c_t^i = \sum\limits_{\tau=t+1}^{T} \left[d_\tau^\top \pi_\tau^{t,i} + \lambda_\tau^{t,i} h_\tau^t + \sum\limits_{m=1}^{M} \sum\limits_{\ell=1}^{\tau-t} \left(\sum\limits_{p=1}^{N_x} \pi_\tau^{t,i}(p) C_\tau(p,m) + \lambda_\tau^{t,i} G_\tau(1,m) \right) \beta_\tau^\ell(m) \mu_{\tau-\ell+1}(m) \right], \\ b_t^i = [b_t^{i,1}; b_t^{i,2}; \ldots; b_t^{i,M}] \text{ where } b_t^{i,m} \text{ is the } (p_{t,T-t}^{\max}(m)+1)\text{-vector given by} \\ b_t^{i,m}(\ell+1) = \sum\limits_{\tau \in I_{t,\ell}(m)} \left[\sum\limits_{p=1}^{N_x} \pi_\tau^{t,i}(p) C_\tau(p,m) + \lambda_\tau^{t,i} G_\tau(1,m) \right] \alpha_{t,\tau-t}^\ell(m),\ \ell = 0, 1, \ldots, p_{t,T-t}^{\max}(m), \end{array}$$

with $I_t^\ell(m) := \left\{ \tau : t+1 \leq \tau \leq T,\ \ell \leq p_{t,\tau-t}^{\max}(m) \right\}$ *and with* h_τ^t *given by* (24)*.*

Proof. Taking the dual of (25), $\mathcal{Q}_{t+1}(\bar{x}_t, \tilde{\xi}_{[t]})$ can be seen as the optimal value of the linear program:

$$\left\{ \begin{array}{l} \max\limits_{\pi_{[t+1:T]}, \lambda_{[t+1:T]}} \pi_{t+1}^\top A_t \bar{x}_t + \sum\limits_{\tau=t+1}^{T} \left[\pi_\tau^\top (C_\tau \mathbb{E}[\xi_\tau|\tilde{\xi}_{[t]}] + d_\tau) + \lambda_\tau (G_\tau \mathbb{E}[\xi_\tau|\tilde{\xi}_{[t]}] + h_\tau^t) \right] \\ \pi_\tau - A_\tau^\top \pi_{\tau+1} + E_\tau^\top \lambda_\tau = 0,\ \ \tau = t+1, \ldots, T-1, \\ -B_\tau^\top \pi_\tau + \mathring{F}_\tau \lambda_\tau = c_\tau,\ \lambda_\tau \geq 0,\ \tau = t+1, \ldots, T, \\ \pi_T + E_T^\top \lambda_T = 0. \end{array} \right. \tag{26}$$

Next, observe that $(\pi_{t+1}^{t,i}, \pi_{t+2}^{t,i}, \ldots, \pi_T^{t,i}, \lambda_{t+1}^{t,i}, \lambda_{t+2}^{t,i}, \ldots, \lambda_T^{t,i})$ is feasible for problem (26). As a result, $\mathcal{Q}_{t+1}(\bar{x}_t, \tilde{\xi}_{[t]})$ is bounded from below by

$$\bar{x}_t^\top A_t^\top \pi_{t+1}^{t,i} + \sum_{\tau=t+1}^{T} \left[(C_\tau \mathbb{E}[\xi_\tau|\tilde{\xi}_{[t]}] + d_\tau)^\top \pi_\tau^{t,i} + \lambda_\tau^{t,i} (G_\tau \mathbb{E}[\xi_\tau|\tilde{\xi}_{[t]}] + h_\tau^t) \right]. \tag{27}$$

Using Lemmas 4.1 and 4.2, we see that $\sigma(X|\tilde{\xi}_{[t]})$ appearing in the definition of h_τ^t is in fact not a function of $\tilde{\xi}_{[t]}$ but only of problem data between time steps $t+1$ and τ. Therefore, recalling definition (24) of h_τ^t, the only terms in the lower bound (27) that are functions of $\tilde{\xi}_{[t]}$ are $\mathbb{E}[\xi_\tau|\tilde{\xi}_{[t]}]$, $\tau = t+1,\ldots,T$. Using (9), we have for these terms the expression $\mathbb{E}[\xi_\tau(m)|\tilde{\xi}_{[t]}] = \sum_{\ell=0}^{p_{t,\tau-t}^{\max}(m)} \alpha_{t,\tau-t}^\ell(m)\tilde{\xi}_{t-\ell}(m) + \sum_{\ell=1}^{\tau-t} \beta_\tau^\ell(m)\mu_{\tau-\ell+1}(m)$, for $m = 1,\ldots,M$, and $\tau = t+1,\ldots,T$. Plugging this expression of $\mathbb{E}[\xi_\tau(m)|\tilde{\xi}_{[t]}]$ into the lower bound (27), we see that this latter can be split into constant terms, a term depending on $\bar{x}_t$, and terms depending on $\tilde{\xi}_{[t]}$. Identifying these terms with respectively c_t^i, $\bar{x}_t^\top a_t^i$, and $\tilde{\xi}_{[t]}^\top b_t^i$, we obtain the desired expression for a_t^i and c_t^i if

$$\begin{aligned}\tilde{\xi}_{[t]}^\top b_t^i &= \sum_{m=1}^{M}\sum_{\tau=t+1}^{T}\sum_{\ell=0}^{p_{t,\tau-t}^{\max}(m)}\left[\sum_{p=1}^{N_x}\pi_\tau^{t,i}(p)C_\tau(p,m)+\lambda_\tau^{t,i}G_\tau(1,m)\right]\alpha_{t,\tau-t}^\ell(m)\tilde{\xi}_{t-\ell}(m)\\ &= \sum_{m=1}^{M}\sum_{\ell=0}^{p_{t,T-t}^{\max}(m)}\sum_{\tau\in I_t^\ell(m)}\left[\sum_{p=1}^{N_x}\pi_\tau^{t,i}(p)C_\tau(p,m)+\lambda_\tau^{t,i}G_\tau(1,m)\right]\alpha_{t,\tau-t}^\ell(m)\tilde{\xi}_{t-\ell}(m)\end{aligned}$$

which gives the desired formula for b_t^i (in the last equality, we have used that the sequence $(p_{t,k}^{\max}(m))_k$ is nondecreasing). □

As a result, after simulating the rolling horizon policy on a set of N scenarios, we obtain N cuts for each recourse function.

At this stage, it is interesting to underline the differences between the algorithm presented above for obtaining approximations of our recourse functions with the usual non-rolling horizon approach. To be more specific, for the latter approach, we will assume that SDDP is used to estimate risk-neutral recourse functions (both approaches are also compared numerically on a real-life problem in Section 6).

(i) In our rolling horizon setting, since

$$\mathcal{Q}_t = \mathcal{Q}_t^{t-1} \quad \text{while} \quad \mathcal{Q}_{t+1} = \mathcal{Q}_{t+1}^t\,,$$

there is no Dynamic Programming relation linking two consecutive recourse functions. On the contrary, SDDP relies on the DP equations linking the successive recourse functions.

(ii) In both cases, at iteration i, a forward pass computes a set of points $(x_t^i, \tilde{\xi}_{[t]}^i)$, $t = 1,\ldots,T-1$, where cuts are built at this iteration. With SDDP, x_t^i is obtained using the available approximation of recourse function $\mathcal{Q}_{t+1}$. On the contrary, in our case, since each risk-averse problem is solved exactly, we obtain the solution that would be obtained using recourse function $\mathcal{Q}_{t+1}$ (though we only have an approximate representation of this function).

(iii) In our case the cuts are built on the basis of dual solutions of the problems solved in the forward pass whereas SDDP needs a backward pass requiring much more computational bulk.

Second algorithm: solving each risk-averse problem by DDP. We have shown in the previous section how a nested decomposition of (11), written for a given history $\tilde{\xi}_{[t]}$, results in the recursion (22). In particular, the "2nd-stage" recourse function computed with $\tau = t$ plays an

important role:

$$\mathcal{Q}_{t+1}(x_t, \tilde{\xi}_{[t]}) \text{ from (10) is given by (21)}$$
$$\text{and corresponds to}$$
$$\mathcal{Q}^t_{t+1}(x_t, \tilde{\xi}_{[t]}), \text{ defined by (22) written for } \tau = t+1.$$

By virtue of (11), Lemma 4.2, and Theorem 4.4, the functions $\mathcal{Q}^t_\tau$ are convex polyhedral functions of $(\bar{x}_{\tau-1}, \tilde{\xi}_{[t]})$ for all $\tau = t, \dots, T$.

On each scenario i and time step t, problem (11) can be solved by applying dual dynamic programming to (22), using a Benders' method capable of handling infeasible states in the forward step; see [6] and also [7]. Along iterations, a set of cuts is built for $\mathcal{Q}^t_\tau$ at some points $(\bar{x}^{t,k}_{\tau-1}, \tilde{\xi}^i_{[t]})$. In particular, when $\tau = t+1$, cuts are generated for the recourse function $\mathcal{Q}_{t+1} = \mathcal{Q}^t_{t+1}$.

To be more specific, we now provide the formulas for the cuts built along the simulated scenarios. In these developments, we assume that all risk-averse problems (11) are feasible. If the DDP algorithm applied for time step t and scenario j requires $K_{t,j}$ iterations then after simulating our RH policy on a set of i scenarios we have built $\sum_{j=1}^i K_{t,j}$ cuts for $\mathcal{Q}^t_\tau$. To alleviate notation, we place ourselves on a given scenario i and time step t and assume that NC cuts have been determined for each $\mathcal{Q}^t_\tau$, $\tau = t+1, \dots, T+1$. Let us see how the next cut is built for these functions. The available approximations $\mathfrak{Q}^{t,NC}_{\tau+1}$ for $\mathcal{Q}^t_{\tau+1}$, $\tau = t, \dots, T$, have the form

$$\mathfrak{Q}^{t,NC}_{\tau+1}(\bar{x}_\tau, \tilde{\xi}_{[t]}) = \max_{1 \le k \le NC} \ \bar{x}_\tau^\top a^{t,k}_\tau + \tilde{\xi}_{[t]}^\top b^{t,k}_\tau + c^{t,k}_\tau. \tag{28}$$

Note that for $t \ge 2$, DDP was used for the previous time step $t-1$ to obtain an optimal solution $\bar{x}^i_{t-1}$ of $\bar{x}_{t-1}$ for risk-averse problem (11) written with t replaced by $t-1$. With this notation, approximations (28) are used in a forward pass to determine a set of points $\bar{x}^{t,NC}_\tau$, $\tau = t-1, t, \dots, T$, with

$$\bar{x}^{t,NC}_{t-1} = \begin{cases} \bar{x}^i_{t-1} & \text{if } t \neq 1, \\ x_0 & \text{if } t = 1, \end{cases} \quad \text{and for } \tau = t, \dots, T,$$

$\bar{x}^{t,NC}_\tau$ an optimal solution of $\bar{x}_\tau$ for the linear program

$$\begin{cases} \min\limits_{\bar{x}_\tau, u_\tau, \theta^t_{\tau+1}} c_\tau^\top u_\tau + \theta^t_{\tau+1} \\ \bar{x}_\tau = A_{\tau-1}\bar{x}^{t,NC}_{\tau-1} + B_\tau u_\tau + C_\tau \mathbb{E}[\xi_\tau | \tilde{\xi}^i_{[t]}] + d_\tau, \\ E_\tau \bar{x}_\tau + \mathring{F}_\tau^\top u_\tau \ge G_\tau \mathbb{E}[\xi_\tau | \tilde{\xi}^i_{[t]}] + h^t_\tau, \\ \theta^t_{\tau+1} \ge \bar{x}_\tau^\top a^{t,k}_\tau + (\tilde{\xi}^i_{[t]})^\top b^{t,k}_\tau + c^{t,k}_\tau, \ k = 1, \dots, NC. \end{cases} \tag{29}$$

New cuts are built for $\mathcal{Q}^t_\tau$, $\tau = t+1, \dots, T+1$, in a backward manner, solving problems of form (29) with one more cut for each $\mathcal{Q}^t_\tau$ (NC is replaced by $NC+1$). These problems are solved from time step $\tau = T$ down to $\tau = t+1$. This is possible since $\mathcal{Q}^t_{T+1} \equiv 0$ is known.

Proposition 5.2 (Alternative cuts computation for robust recourse functions)**.** *Consider risk averse problem* (23) *for some* $t \in \{1, \dots, T-1\}$. *Let* $\mathcal{Q}_{t+1}$ *be the corresponding recourse function defined in* (25) *and let* $\mathcal{Q}^t_\tau$, $\tau = t+1, \dots, T+1$, *be the functions given in* (22). *Let* $(\pi^{t,NC}_\tau, \lambda^{t,NC}_\tau, (\rho^{t,NC,k}_\tau)_{k=1}^{NC+1})$ *be a dual solution of problem* (29) *written with* NC *replaced by* $NC+1$. *Then valid cuts for* $\mathcal{Q}_{T+1}$ *are obtained taking null values for* $a^{t,NC}_T, b^{t,NC}_T$, *and* $c^{t,NC}_T$,

while valid cuts for $\mathcal{Q}_\tau^t$, $\tau = t+1, \ldots, T$, *are given by*

$$\begin{array}{lcl}
a_{\tau-1}^{t,NC+1} & = & A_{\tau-1}^\top \pi_\tau^{t,NC}, \quad c_{\tau-1}^{t,NC+1} = d_\tau^\top \pi_\tau^{t,NC} + \lambda_\tau^{t,NC} h_\tau^t + \sum_{k=1}^{NC+1} \rho_\tau^{t,NC,k} c_\tau^{t,k}, \\
b_{\tau-1}^{t,NC+1} & = & [b_{\tau-1}^{t,NC+1,1}; b_{\tau-1}^{t,NC+1,2}; \ldots; b_{\tau-1}^{t,NC+1,M}] \text{ where } b_{\tau-1}^{t,NC+1,m} \text{ is the vector given by} \\
b_{\tau-1}^{t,NC+1,m}(\ell+1) & = & \left[\sum_{p=1}^{N_x} \pi_\tau^{t,NC}(p) C_\tau(p,m) + \lambda_\tau^{t,NC} G_\tau(1,m)\right] \alpha_{t,\tau-t}^\ell(m) \\
& & + \sum_{k=1}^{NC+1} \rho_\tau^{t,NC,k} b_\tau^{t,k,m}(\ell+1) \text{ for } \ell = 0, 1, \ldots, p_{t,\tau-t}^{\max}(m), \text{ while} \\
b_{\tau-1}^{t,NC+1,m}(\ell+1) & = & \sum_{k=1}^{NC+1} \rho_\tau^{t,NC,k} b_\tau^{t,k,m}(\ell+1) \text{ for } \ell = p_{t,\tau-t}^{\max}(m)+1, \ldots, p_{t,T-t}^{\max}(m),
\end{array}$$

with h_τ^t *given by* (24). *In particular,* $a_t^{t,NC+1}, b_t^{t,NC+1}$, *and* $c_t^{t,NC+1}$ *define a valid cut for* $\mathcal{Q}_{t+1} = \mathcal{Q}_{t+1}^t$.

Proof. The dual of optimization problem (29) written with NC, $\bar{x}_{\tau-1}^{t,NC}$, and $\tilde{\xi}_{[t]}^i$ respectively replaced by $NC+1$, $\bar{x}_{\tau-1}$, and $\tilde{\xi}_{[t]}$ is given by

(30)
$$\begin{cases}
\max\limits_{\pi_\tau, \lambda_\tau, \rho_\tau^k} \pi_\tau^\top \left[A_{\tau-1}\bar{x}_{\tau-1} + C_\tau \mathbb{E}[\xi_\tau | \tilde{\xi}_{[t]}] + d_\tau\right] + \lambda_\tau \left[G_\tau \mathbb{E}[\xi_\tau | \tilde{\xi}_{[t]}] + h_\tau^t\right] + \sum\limits_{k=1}^{NC+1} \rho_\tau^k \left[\tilde{\xi}_{[t]}^\top b_\tau^{t,k} + c_\tau^{t,k}\right] \\
-B_\tau^\top \pi_\tau + \mathring{F}_\tau \lambda_\tau = c_\tau, \ \sum_{k=1}^{NC+1} \rho_\tau^k = 1, \ \ \pi_\tau + E_\tau^\top \lambda_\tau - \sum_{k=1}^{NC+1} \rho_\tau^k a_\tau^{t,k} = 0, \\
\lambda_\tau \geq 0, \ \rho_\tau^k \geq 0, \ k = 1, \ldots, NC+1.
\end{cases}$$

We then conclude as in Proposition 5.1, using (9) and observing that $(\pi_\tau^{t,NC}, \lambda_\tau^{t,NC}, (\rho_\tau^{t,NC,k})_{k=1}^{NC+1})$ is feasible for problem (30). □

On the first scenario, when no approximations of functions $\mathcal{Q}_\tau^t$ are available, we start the first forward passes using known lower bounding functions (for instance constant functions).

At each iteration of the DDP algorithm (applied on each scenario at each time step), an upper bound on the optimal value of the corresponding risk-averse problem can be computed using a feasible solution. A lower bound can be built solving an approximate problem of form (29) for time step $\tau = t$. The DDP algorithm stops when the difference between these upper and lower bounds is smaller than a given confidence level.

6. Numerical experience

We assess the rolling horizon approach on a large-scale instance of the water-planning problem described in Section 2, similar to Brazil's power system, except for the standard deviations of noises η_t; see Section 6.1 below.

6.1. **Power system data.** We consider a hydro-thermal power system operating over a horizon of 4 years, discretized in $T = 48$ time steps, from January 2005 to December 2008. Most of the data was made available by CEPEL[3] and corresponds to part of Brazil's power system, represented by 4 different subsystems that can trade energy in the form of import-export exchanges. Each subsystem, South-East (SE), South (S), North-East (NE), and North (N), corresponds to a geographical region; some energy exchanges between the N, NE, and SE subsystems make use of a fifth, fictitious, node (F). In a specific subsystem, a single reservoir aggregates all the hydro-power, while thermal generation is considered individually: there are 24, 14, 6, and 0 thermal plants in the SE (the largest one), S, NE, and N subsystems, respectively.

[3]The authors specially acknowledge the good will and availability of Débora Dias Jardim Penna.

With respect to the simplified model given in Section 2, there are additional control variables (spillage and import-export exchanges), as well as additional lower and upper bounds for hydro and thermal generation, for subsystem exchanges, and for reservoir volumes. The water balance equation (3) now considers spillage, while demand constraints (4) consider the subsystem thermal generation as well as import-export exchanges, with a total monthly demand of 54804 MWMonth [4], taken constant over the horizon. The water balance relations (3) are started by setting the initial reservoir levels $x_0(m)$ at full capacity. Each reservoir critical level $x_t^{crit}(m)$ in (5) was set to 20% of the maximum level of the reservoir, for all time steps, so that assumption (A1) holds (recall that if (A1) did not hold, our methodology would still be applicable, by adding slack variables to the critical minimal volume constraints). With this assumption, relatively complete recourse is guaranteed if $\mathbb{P}(\xi_t(m) \geq 0) = 1$, and we checked numerically that $\max_{t,m} \mathbb{P}(\xi_t(m) < 0) = 2.7 \times 10^{-75}$.

The objective function is given by the total thermal operating cost (ranging between R$ 6.27 per MWh and R$ 1047 per MWh) plus the cost of load shedding (the shortage cost is set at R$ 4170.44 per MWh). Unnecessary spillage and exchanges are discouraged by introducing small penalties and trading costs between subsystems.

Following the lines of [22], the inflows in each reservoir are modelled by a periodic autoregressive model of the form (6). The parameters of each model were estimated based on historical data from 1931 to 2005, with one important modification, relative to standard deviations. Namely, we reduced the estimated value of $\sigma_t^\eta(m)$ because, with the original estimations, the model generated too many negative water inflows that have no meaningful physical interpretation. A more sophisticated implementation that keeps the original standard deviations should use log-normal distributions for the noises $\eta_t(m)$. We adopted the former approach not only for simplicity, but also for noises to remain Gaussian, so that we can use exact expressions for the inverse of the cumulative distribution functions in Theorem 4.4. As for the conditional expectations and standard deviations in Theorem 4.4, they are computed exactly, using the recursive expressions derived in [18] for the coefficients in (9) and model (6) parameters.

Due to the modified standard deviations, our results should be interpreted as an illustration of our methodology, rather than reflecting the real behavior of the Brazilian power system.

To obtain the indicators referred to in item `(EcInd)` of Section 2, the simulation phase uses $N = 10000$ streamflow scenarios (we displayed 500 of them in Figure 1 for the S and SE subsystems).

The implementation was done in Matlab, using Mosek's optimization library to solve linear programming problems (`http://www.matlab.com` and `http://www.mosek.com`).

6.2. **Impact of confidence levels chosen a priori on the RH policy.** In this section, we illustrate the influence of the confidence levels on the behavior of `RH` policy. For illustration, we consider only CVaR constraints and different values for parameters $\varepsilon_{\mathtt{p}}$ and $\varepsilon_{\mathtt{c}}$, chosen a priori and taken constant for the different time steps and reservoirs.

The following set of confidence levels $\varepsilon_{\mathtt{c}}$ is considered: 0.01, 0.05, 0.1, 0.2, 0.4, 0.6, and 0.8. For each confidence level, the corresponding `RH` policy is simulated. We report in Figure 3 the mean and standard deviation (s.d) of the total cost, and the total hydro productions.

[4]We adopt the convention 1 MWMonth$=\frac{365.25\times 24}{12}$ MWh$=$ 730.5 MWh

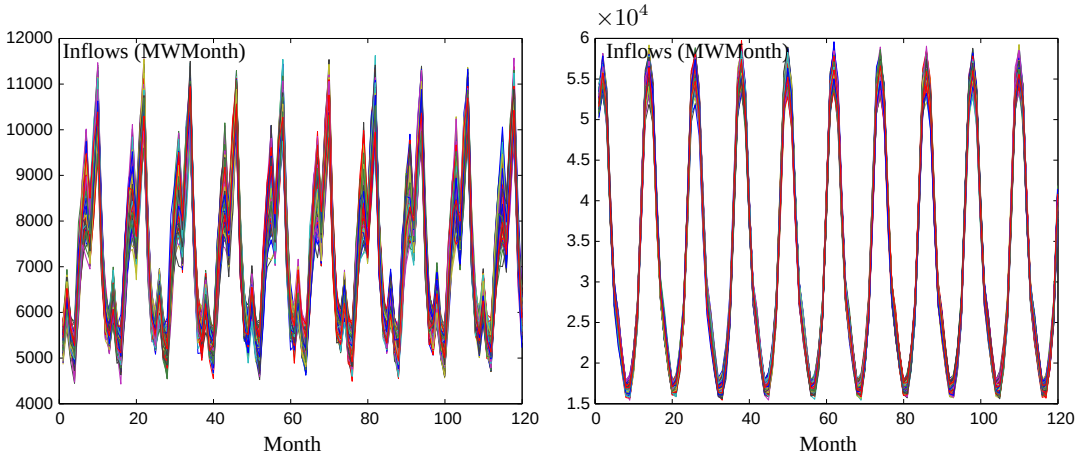

FIGURE 1. 500 inflow scenarios for the S and SE subsystems (left and right, respectively).

In both cases, we observe that the mean and the variability (s.d) of the total cost increase with confidence levels. The increase in these confidence levels Indeed, when confidence levels increase, risk-averse problem (11) foresees lower future "robust" demands and crtical levels; respectively referred to as dem_τ^R and $x_\tau^{R\ crit}$ in Section 5.1. As a result, increasing confidence levels forces to use more water on the first part of the optimization period. This prevents the impact of the confidence levels on water reservoirs management: the more we are risk-averse (the smaller the confidence levels reservoir)

Finally, load shedding

In the sequel, we compare three different policies:

- `RH`: our robust rolling horizon policy using a dyamic choice of confidence levels, as explained in Section (3.3). Each risk-averse problem (11) was solved directly as a linear program, without using the stagewise decomposition in Section 5.1.
- `RA-NRH`: the non-rolling risk-averse policy referred to in Section 3.2 that uses the approximate robust recourse functions built with the first algorithm from Section 5. This algorithm uses a preliminary optimization phase of 1000 scenarios for approximating the recourse functions.
- `SDDP`: a non-rolling horizon risk-neutral policy, that approximates recourse functions by Stochastic Dual Dynamic Programming as in [24], using 200 fixed scenarios for the forward pass and 20 discrete realizations for each noise η_t in the backward pass.

6.3. **Comparison between SDDP and RH using a dynamic choice of confidence levels.** Table 2 reports on cost-related results obtained with the variants, referred to as `RH CVaR`, `RH CC`, and `RH CC-CVaR`. Each table contains the empirical mean and standard deviation (s.d.) of the whole system generation and exchange cost over the 500 scenarios, as well as the corresponding VaR $p\%$, for $p = 1, 5$, and 90, where VaR $p\%$ is the $(1-p/100)$-quantile of the empirical distribution of the cost.[5]

[5]If X is a continuous random variable for which lower values are preferred, its Value-at-Risk of level ε_{p} is given by $VaR_{\varepsilon_{\mathrm{p}}}(X) := F_X^{-1}(1-\varepsilon_{\mathrm{p}})$ for any $\varepsilon_{\mathrm{p}} \in [0,1]$.

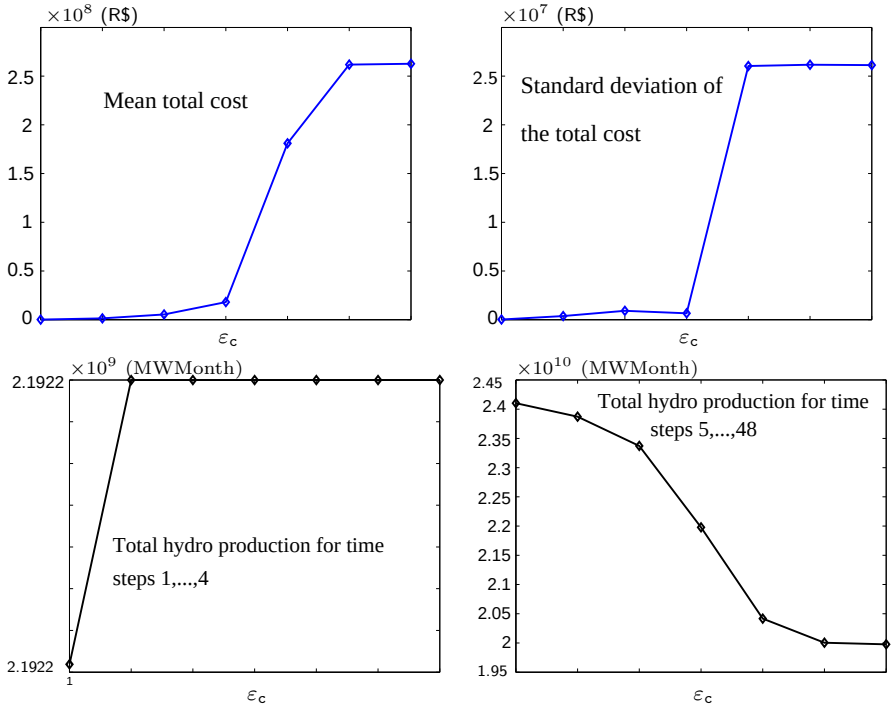

FIGURE 2. RH policy with only chance constraints and a priori choice of confidence levels ε_p. Evolution of the mean (upper left) and s.d (upper right) of the total cost for different values of these confidence levels. Evolution of the equivalent reservoir mean level (bottom left) and of mean hydro productions (bottom right).

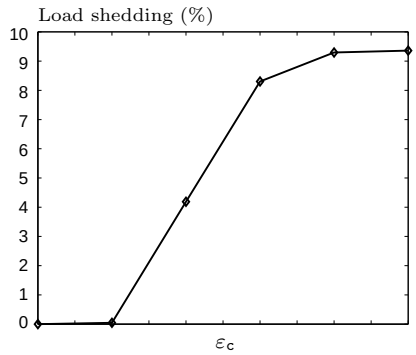

FIGURE 3. RH policy with only CVaR constraints and a priori choice of confidence levels ε_c. Evolution of the mean (upper left) and s.d (upper right) of the total cost for different values of these confidence levels. Evolution of the equivalent reservoir mean level (bottom left) and of mean hydro productions (bottom right).

TABLE 2. Central and dispersion characteristics of the generation and exchange cost (R$)

Output $(\varepsilon_{\mathtt{p}}, \varepsilon_{\mathtt{c}})$	RH CVaR $(0.4, 0.05)$	RH CVaR $(0.45, 0.01)$	RH CVaR $(0.45, 0.005)$	RH CC $(0.19, --)$	RH CC-CVaR $(0.45, 0.01)$	RH CC-CVaR $(0.45, 0.005)$
Mean	5.192×10^9	6.865×10^8	7.532×10^8	8.206×10^8	7.947×10^8	7.954×10^8
s.d.	6.144×10^8	1.744×10^8	1.517×10^8	1.403×10^8	1.425×10^8	1.438×10^8
VaR 1%	6.625×10^9	1.115×10^9	1.137×10^9	1.165×10^9	1.154×10^9	1.152×10^9
VaR 5%	6.264×10^9	9.911×10^8	1.015×10^9	1.063×10^9	1.042×10^9	1.043×10^9
VaR 90%	4.387×10^9	4.639×10^8	5.556×10^8	6.355×10^8	6.066×10^8	6.050×10^8

On this set of scenarios and for our data, the lowest cost-related values are obtained with RH CVaR policy, for $(\varepsilon_{\mathtt{p}}, \varepsilon_{\mathtt{c}}) = (0.45, 0.01)$, while the highest ones are given by RH CVaR variant with $(\varepsilon_{\mathtt{p}}, \varepsilon_{\mathtt{c}}) = (0.4, 0.05)$.

It is important to keep in mind that problems like (1) are simply not tractable for large time horizons (recall that $T = 120$ in our application). For this reason, there are very few policies that can be implemented, even in the risk-neutral setting. An exception is the risk-neutral policy estimated by sampling decomposition methods such as SDDP. However, as confirmed by our numerical results below, SDDP approximate policies are less accurate than the ones obtained with the rolling horizon approach. This feature is partly explained by the fact that the RH approach can fully exploit the probabilistic structure of uncertainty, while SDDP needs to resort to a discretization. More precisely, because the inverse of the (Gaussian) $\mathcal{N}(0,1)$ distribution function, F^{-1}, is available with arbitrary accuracy, RH needs no approximation or discretization for the distributions of $\xi_2, \ldots, \xi_T$ in Theorem 4.4. The situation is quite different with SDDP, since it samples forward paths from an approximate discrete distribution for $\xi_2, \ldots, \xi_T$, obtained from discrete noises $\eta_2, \ldots, \eta_T$, each one taking 20 equiprobable different values in our implementation. As a result, the corresponding scenario tree has 20^{119} paths, out of which 200 are traversed by SDDP at the end of the optimization process. It could be argued that our SDDP implementation does not resample scenarios in the forward pass and uses the stopping test from [24], deemed too precocious in [**?**]; see also [**?**]. Notwithstanding, neither increasing the number of sampled forward paths (say by a factor 10^{10}), nor tightening the stopping test can substantially change the fact that only a negligible proportion of the scenario tree is explored by the method. Also, scenarios used in the simulation phase are drawn from the true, continuous, distributions of $\xi_2, \ldots, \xi_T$ given by model (6) and not from the approximate discrete distributions used in the optimization phase of SDDP.

An additional comparison of both policies is given in Table 7, showing how severe is SDDP computational load, already in our "loose" implementation, when compared to RH.

We refer to the end of Section 6.7, for a more specific analysis of SDDP in our test-case.

6.4. **Critical levels, operational costs, and load shedding.** In our runs, future constraints (5) in problem (11), on the reservoir minimal levels, were not active in general. In fact, these constraints were satisfied with a probability much larger than $1 - \varepsilon_{\mathtt{p}}$, for all of our choices of $\varepsilon_{\mathtt{p}}$. We made a set of unreported tests to evaluate the sensitivity of both WS and RH to variations in

TABLE 3. Computational load of the implemented policies for the test-case

SDDP optimization phase (one iteration)	Number of LPs	Each LP has 56 variables and box constraints, 8 equality/inequality constraints, and for $t < T$, a number of cuts depending on iteration k
Forward	24 000	200(k-1)+1
Backward	480 000	200k+1

Policy	Computational Effort
SDDP optimization phase	192 000 Forward LPs
(8 iterations)	3.84×10^6 Backward LPs
SDDP simulation phase	60 000 LPs with 1601 cuts
SDDP total CPU time	3 weeks
RH, for each $t = 1, \ldots, 120$	500 LPs with $56(T - t + 1)$ variables and box constraints and $8(T - t + 1)$ equality/inequality constraints
RH total CPU time	6 h

the critical levels $x_t^{crit} \in \{0.1, 0.2, 0.5, 0.8\}x_t^{\max}$ and we observed a natural behavior for a hydro-dominated system: costs became prohibitive for the higher values of the critical levels, due to the appearance of load shedding for many time steps.

Our assessment of the different policies starts with Table 8, reporting the corresponding statistical indicators for operational costs, that is, generation and exchange costs, excluding shortage.

TABLE 4. Central and dispersion characteristics of the generation and exchange cost (R$)

Output	WS	RH	SDDP
Mean	5.752×10^8	6.865×10^8	6.809×10^8
s.d.	2.059×10^8	1.744×10^8	3.682×10^8
VaR 1%	1.026×10^9	1.115×10^9	2.146×10^9
VaR 5%	9.273×10^8	9.911×10^8	1.272×10^9
VaR 90%	3.115×10^8	4.639×10^8	3.299×10^8

From a statistical point of view, the p-value of a Student's paired t-test applied on the samples of the cost of policies WS and SDDP (resp. RH and SDDP) is $10^{-21} \simeq 0$ (resp. 0.61)[6]. Therefore, RH and SDDP have comparable mean costs, different from the optimal policy WS. Since, in addition, the p-value of an F-test applied on the samples of the cost of policies RH and SDDP is $1.5\times10^{-57} \simeq 0$, the respective standard deviations can be considered different.

Figure 10 compares the empirical distribution of the generation and exchange cost over the 500 scenarios for the three policies, where $C_{\texttt{WS}}, C_{\texttt{RH}}$, and $C_{\texttt{SDDP}}$ are respectively the generation and exchange cost obtained with WS, RH, and SDDP. Due to the different magnitude of the costs in Table 8, to ease the comparisons the graphs show the distribution of the corresponding ratios.

The cost distribution for SDDP exhibits a fatter tail, because with this policy some scenarios are very costly. This is a natural result, because RH is risk-averse while SDDP uses a risk-neutral

[6]RH mean cost is 19.3% higher than WS mean cost but only corresponds to a 0.8% increase with respect to SDDP mean cost

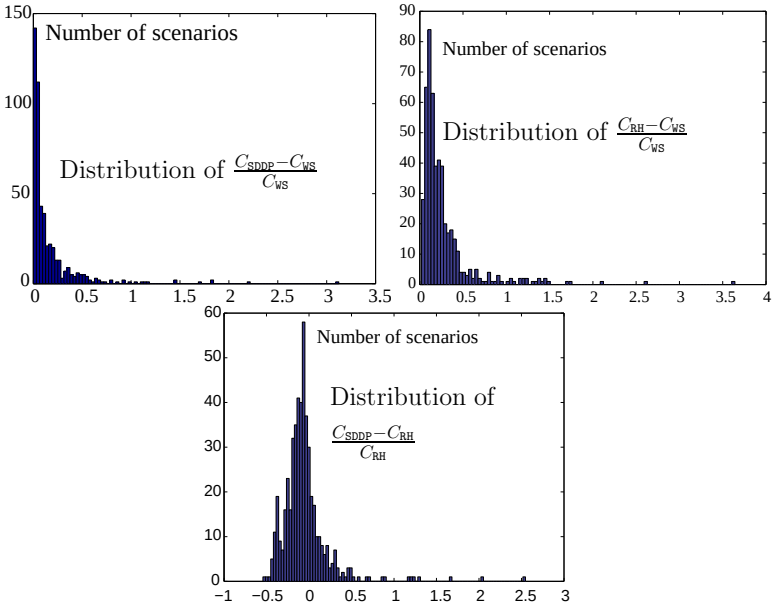

FIGURE 4. Empirical distribution of cost ratios obtained with the three policies.

formulation. This is also consistent with Table 8, showing larger values for the quantiles VaR 1% and VaR 5% for the SDDP policy.

It is also important to look at the portion of demand left unsatisfied by each policy, when compared to the total cost, considering all scenarios and time steps. SDDP incurs a very low level of load shedding (0.001%), at the expense of a significant increase in the mean cost (shortage cost is high: R$ 4170.44 per MWh). All the rolling horizon runs satisfied the demand, excepting RH CVaR with $(\varepsilon_{\texttt{p}}, \varepsilon_{\texttt{c}}) = (0.4, 0.05)$, with a load shedding level of 0.153%.

Table 9 reports the total cost for policies WS, SDDP, and the "best" and "worst" RH policies. Except for the "worst" one, that leaves demand unsatisfied, all the rolling horizon policies had total mean cost smaller than the one obtained with SDDP (SDDP total mean cost is about 24% above the corresponding value for the "best" RH policy).

TABLE 5. Central and dispersion characteristics of the total cost (R$)

Output	WS	Best RH CVaR $(\varepsilon_{\texttt{p}}, \varepsilon_{\texttt{c}}) = (0.45, 0.01)$	Worst RH CC-CVaR $(\varepsilon_{\texttt{p}}, \varepsilon_{\texttt{c}}) = (0.4, 0.05)$	SDDP
Mean	5.752×10^{8}	6.865×10^{8}	3.302×10^{10}	8.543×10^{8}
s.d.	2.059×10^{8}	1.744×10^{8}	8.146×10^{9}	2.185×10^{9}
VaR 1%	1.026×10^{9}	1.115×10^{9}	5.132×10^{10}	2.991×10^{9}
VaR 5%	9.273×10^{8}	9.911×10^{8}	4.672×10^{10}	1.272×10^{9}
VaR 90%	3.115×10^{8}	4.639×10^{8}	2.269×10^{10}	3.299×10^{8}

The mean generation cost per subsystem for the three policies is given in Table 10.

From these tables, and as already announced in Section **??**, for our data set and in terms of cost distribution, SDDP policy is far from an optimal (or a close to optimal) risk-neutral policy. Because the stochastic process is interstage dependent, the process history (7) gives in (1) a state vector

TABLE 6. Mean generation cost (over the scenarios) in the subsystems (R$).

Policy	South-East	South	North-East	North
WS	5.726×10^{8}	0	0	0
RH	6.823×10^{8}	2.056×10^{5}	1.527×10^{6}	0
SDDP	6.428×10^{8}	2.522×10^{7}	1.029×10^{7}	0

with 24 components. Such state dimension is practically unmanageable by Dynamic Programming techniques for large horizons T, even in a deterministic setting. Uncertainty only worsens the situation, making SDDP simulation phase use an approximation of the recourse function that is likely to be poor at the considered states. Better approximations could be obtained by setting up an SDDP rolling horizon methodology, but, in view of the figures in Table 7, it is clear that this alternative is computationally impossible. Likewise, a risk-averse formulation, for example along the lines of [**?**] or [**?**], needs more variables and constraints, and would increase the already heavy computational effort. Moreover, even if distributing calculations in many different processors made the computational load acceptable, implementing a risk-averse SDDP as in [**?**] would still require to find at each stage suitable weights defining a trade-off between the mean and the CVaR of the future cost. At first sight, choosing sound values for these weights for a specific problem is not a simpler task than finding adequate confidence levels $\varepsilon_{\mathtt{p}}$ and $\varepsilon_{\mathtt{c}}$ for the rolling horizon methodology.

6.5. **Exchanges, hydro-thermal generation, and marginal costs.** The most significant import-export exchanges go from the South-East to the South, for all policies. This is due to the fact that while the South has the highest demand at each time step, the South-East has both the highest hydro- (more than 70% of the total) and thermal-capacity. Other exchanges appear from the North to the North-East and the South-East, and from the North-East to the South-East. With WS and RH policies, and contrary to SDDP, the North receives no energy from the other subsystems. This subsystem, with no thermal plants and only one small reservoir, but also small demand, has a somehow "isolated" operation with both WS and RH. This is not the case with SDDP which, as shown in Figure 11, uses more water than the other two policies in the first half of the optimization period and, hence, ends up the needing to transfer some energy to the North.

We report in Figure 11 the average evolution for the South-East and South subsystems. Figure 12 displays the hydro-generation evolution of the whole system.

We see that, during the first five years, SDDP uses more water at nearly all time steps and scenarios. In particular, for SDDP, hydro-generation does not fluctuate much, likely because of the use of approximate recourse functions that are optimistic for the first years. This explains why SDDP's reservoir levels and generation costs tend to be lower than the other policies over the first half of the optimization period. However, since this initial optimism also prevents SDDP from using as much water as RH and WS later on, eventually SDDP becomes more costly. With all the policies, both the S and SE reach (or nearly reach) their critical levels on all scenarios at the end of the optimization period.

The right graph in Figure 12, with the 0.05-quantile for the hydro-generation, shows a relatively low variability from scenario to scenario. All policies make a similar management of the hydro-power. However, over the first half of the optimization period, RH tends to be a bit more conservative

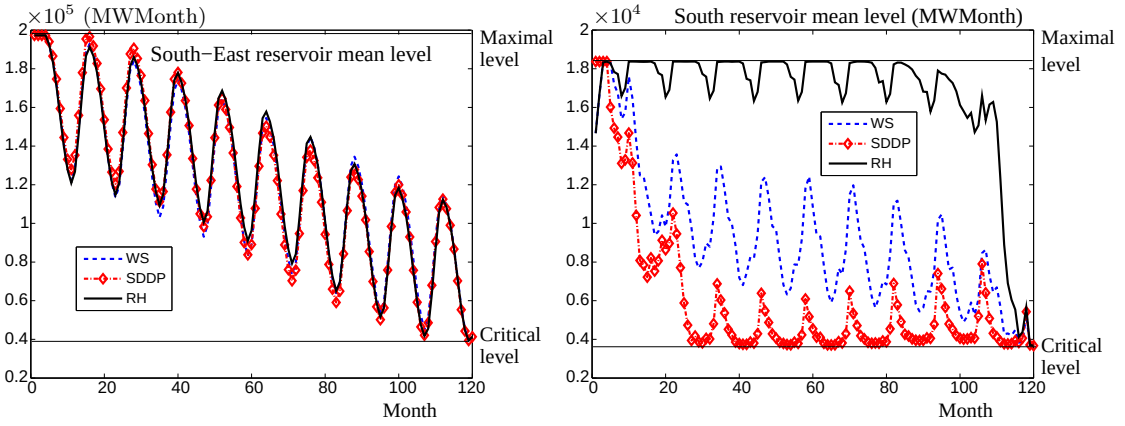

FIGURE 5. Mean reservoir level for South-East and South subsystems.

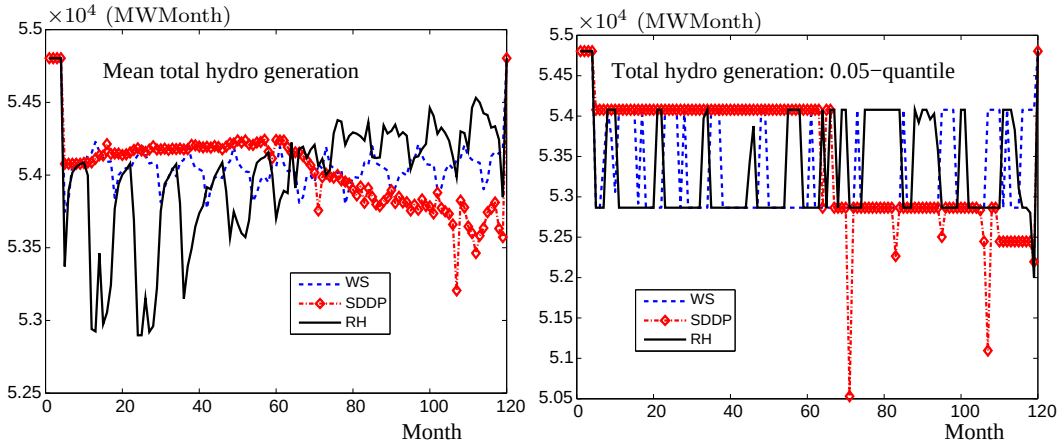

FIGURE 6. Mean hydro-generation (left) and 0.05-quantile (right).

with respect to the use of water and uses slightly less water than the other two policies. The conservatism of RH over the first years may be explained by the fact that, at the beginning of the process ($t \approx 1$), (11) needs to foresee a far away future ($\tau \in [2, 120]$), using robust uncertainty sets (19) and (20) that are overly pessimistic with respect to the final stages ($\tau \approx 120$). By contrast, over the last five years, as the number of uncertain parameters in (11) decreases, uncertainty sets become tighter, RH "realizes" that more water than necessary has been stored and it uses slightly more water than the other two policies.

Figure 13 plots the mean level and 0.05- and 0.95-quantiles for the total thermal generation. With both RH and WS, since the demand is satisfied for every time step and scenario, the thermal generation merely complements the hydro-generation to attain the demand level, a natural feature for a predominantly hydroelectric system.

It appears that RH policy uses more (resp. less) thermal power over the first half (resp. second half) of the optimization period; a phenomenon which is consistent with previous observations. For this reason, RH is more (resp. less) costly than WS over the first half (resp. second half) of the optimization period. Since SDDP uses less water in the second half of the optimization period, it

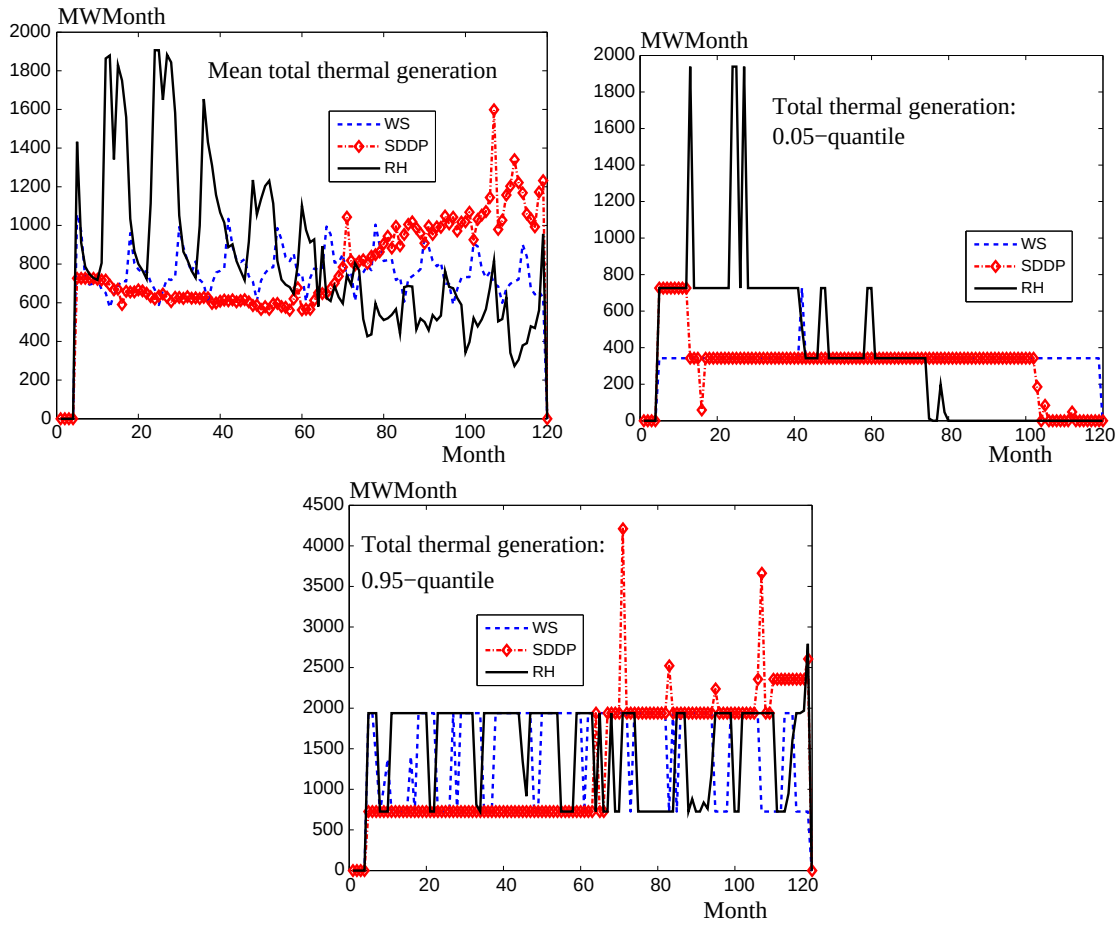

FIGURE 7. Total thermal generation: mean, 0.05-, and 0.95-quantiles.

needs to call on more thermal plants. It can also be seen (by comparing Figures 12 and 13) that the thermal generation is comparatively much smaller than the hydro-generation, for all policies, as expected in a hydro-dominated system (in our configuration, thermal power can cover at most 12.8% of the demand at each time step). When needed, thermal plants are committed in ascending order of their operational cost, to prevent load shedding. Due to the fact that only some thermal plants enter the system, the marginal cost (mean optimal Lagrange multipliers for the demand satisfaction constraint) in all subsystems is low for all the policies, except for some time steps with `SDDP`, as shown in Figure 8.

More precisely, at some time steps `SDDP` needs to call on more thermal plants than `RH` and `WS`, because `SDDP` uses less water, in particular in the second half of the optimization period. For such time steps, `SDDP`'s mean marginal cost can climb from around R$10-15 per MWh to R$ 38.7 per MWh, with peaks to R$ 4170.44 per MWh for some realizations. In fact, `SDDP`'s variability of marginal cost (from one scenario to another) is quite significant, as reported in Figure 14.

Stability appears as another advantage of `RH` over `SDDP`: we obtain lower marginal costs that are less volatile. Notwithstanding, on average, both policies give relatively low values: the mean

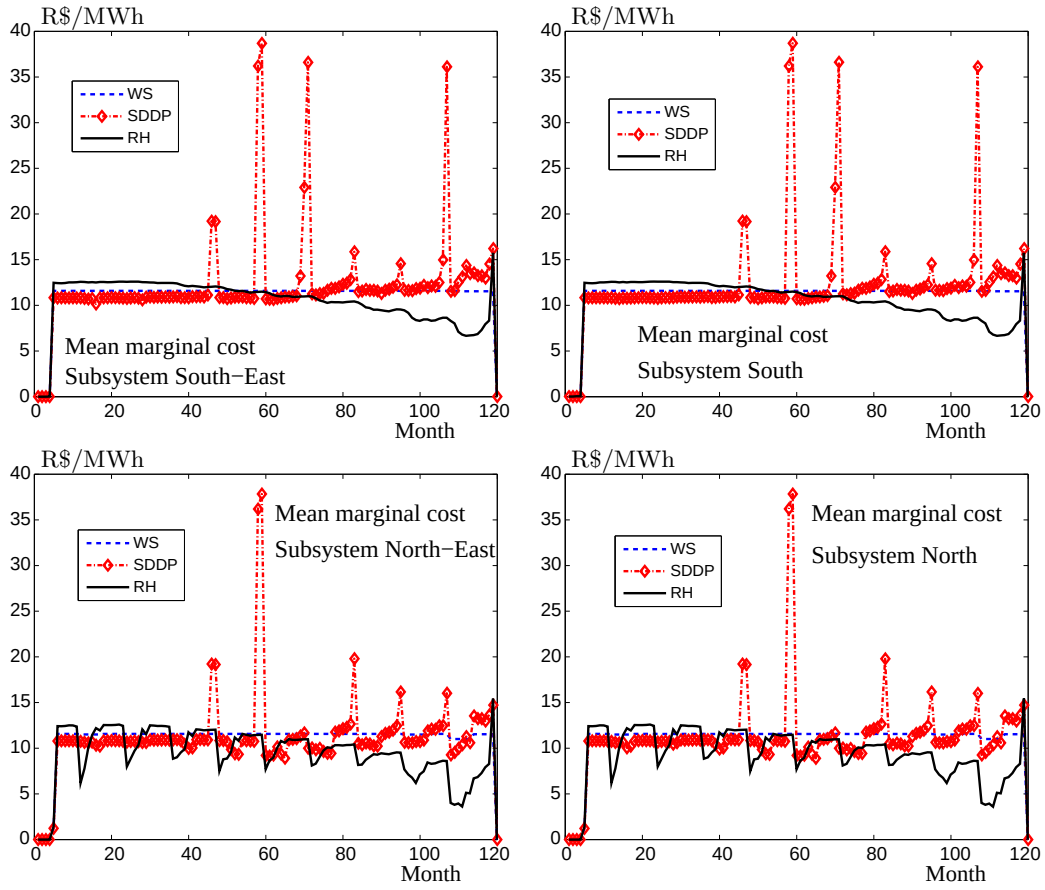

FIGURE 8. Mean marginal costs per subsystem.

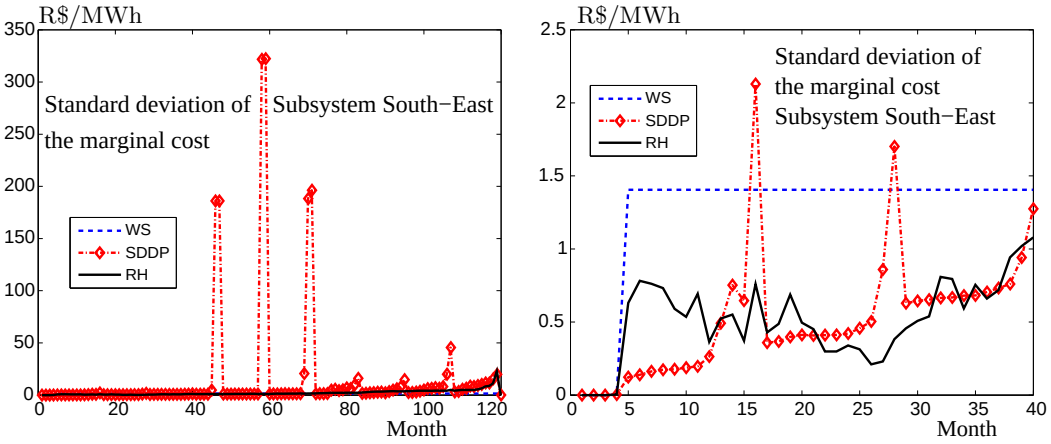

FIGURE 9. Marginal cost standard deviations for SE subsystem, with a zoom for the first 40 time steps.

marginal costs in R$ per MWh obtained with RH (SDDP) policy amount to 10.39 (12.13), 10.39 (12.14), 9.37 (11.07), and 9.37 (11.08) for the SE, S, NE, and N subsystems, respectively.

6.6. **Comparison between SDDP and our approximate robust recourse functions.** It is important to keep in mind that problems like (1) are simply not tractable for large time horizons (recall that $T = 120$ in our application). For this reason, there are very few policies that can be implemented, even in the risk-neutral setting. An outstanding exception is the risk-neutral policy estimated by sampling decomposition methods such as SDDP. However, as confirmed by our numerical results below, SDDP approximate policies are less accurate than the ones obtained with the rolling horizon approach. This feature is partly explained by the fact that the RH approach can fully exploit the probabilistic structure of uncertainty, while SDDP needs to resort to a discretization. More precisely, because the inverse of the (Gaussian) $\mathcal{N}(0,1)$ distribution function, F^{-1}, is available with arbitrary accuracy, RH needs no approximation or discretization for the distributions of $\xi_2, \ldots, \xi_T$ in Theorem 4.4. The situation is quite different with SDDP, since it samples forward paths from an approximate discrete distribution for $\xi_2, \ldots, \xi_T$, obtained from discrete noises $\eta_2, \ldots, \eta_T$, each one taking 20 equiprobable different values in our implementation. As a result, the corresponding scenario tree has 20^{119} paths, out of which 200 are traversed by SDDP at the end of the optimization process. It could be argued that our SDDP implementation does not resample scenarios in the forward pass and uses the stopping test from [24], deemed too precocious in [**?**]; see also [**?**]. Notwithstanding, neither increasing the number of sampled forward paths (say by a factor 10^{10}), nor tightening the stopping test can substantially change the fact that only a negligible proportion of the scenario tree is explored by the method. Also, scenarios used in the simulation phase are drawn from the true, continuous, distributions of $\xi_2, \ldots, \xi_T$ given by model (6) and not from the approximate discrete distributions used in the optimization phase of SDDP.

An additional comparison of both policies is given in Table 7, showing how severe is SDDP computational load, already in our "loose" implementation, when compared to RH.

TABLE 7. Computational load of the implemented policies for the test-case

SDDP optimization phase (one iteration)	Number of LPs	Each LP has 56 variables and box constraints, 8 equality/inequality constraints, and for $t < T$, a number of cuts depending on iteration k
Forward	24 000	$200(k\text{-}1)+1$
Backward	480 000	$200k+1$

Policy	Computational Effort
SDDP optimization phase	192 000 Forward LPs
(8 iterations)	3.84×10^6 Backward LPs
SDDP simulation phase	60 000 LPs with 1601 cuts
SDDP total CPU time	3 weeks
RH, for each $t = 1, \ldots, 120$	500 LPs with $56(T-t+1)$ variables and box constraints and $8(T-t+1)$ equality/inequality constraints
RH total CPU time	6 h

We refer to the end of Section 6.7, for a more specific analysis of SDDP in our test-case.

6.7. **Critical levels, operational costs, and load shedding.** In our runs, future constraints (5) in problem (11), on the reservoir minimal levels, were not active in general. In fact, these constraints were satisfied with a probability much larger than $1 - \varepsilon_{\mathtt{p}}$, for all of our choices of $\varepsilon_{\mathtt{p}}$. We made a set of unreported tests to evaluate the sensitivity of both `WS` and `RH` to variations in the critical levels $x_t^{crit} \in \{0.1, 0.2, 0.5, 0.8\} x_t^{\max}$ and we observed a natural behavior for a hydro-dominated system: costs became prohibitive for the higher values of the critical levels, due to the appearance of load shedding for many time steps.

Our assessment of the different policies starts with Table 8, reporting the corresponding statistical indicators for operational costs, that is, generation and exchange costs, excluding shortage.

TABLE 8. Central and dispersion characteristics of the generation and exchange cost (R$)

Output	`WS`	`RH`	`SDDP`
Mean	5.752×10^8	6.865×10^8	6.809×10^8
s.d.	2.059×10^8	1.744×10^8	3.682×10^8
VaR 1%	1.026×10^9	1.115×10^9	2.146×10^9
VaR 5%	9.273×10^8	9.911×10^8	1.272×10^9
VaR 90%	3.115×10^8	4.639×10^8	3.299×10^8

From a statistical point of view, the p-value of a Student's paired t-test applied on the samples of the cost of policies `WS` and `SDDP` (resp. `RH` and `SDDP`) is $10^{-21} \simeq 0$ (resp. 0.61)[7]. Therefore, `RH` and `SDDP` have comparable mean costs, different from the optimal policy `WS`. Since, in addition, the p-value of an F-test applied on the samples of the cost of policies `RH` and `SDDP` is $1.5\times10^{-57} \simeq 0$, the respective standard deviations can be considered different.

Figure 10 compares the empirical distribution of the generation and exchange cost over the 500 scenarios for the three policies, where $C_{\mathtt{WS}}, C_{\mathtt{RH}}$, and $C_{\mathtt{SDDP}}$ are respectively the generation and exchange cost obtained with `WS`, `RH`, and `SDDP`. Due to the different magnitude of the costs in Table 8, to ease the comparisons the graphs show the distribution of the corresponding ratios.

The cost distribution for `SDDP` exhibits a fatter tail, because with this policy some scenarios are very costly. This is a natural result, because `RH` is risk-averse while `SDDP` uses a risk-neutral formulation. This is also consistent with Table 8, showing larger values for the quantiles VaR 1% and VaR 5% for the `SDDP` policy.

It is also important to look at the portion of demand left unsatisfied by each policy, when compared to the total cost, considering all scenarios and time steps. `SDDP` incurs a very low level of load shedding (0.001%), at the expense of a significant increase in the mean cost (shortage cost is high: R$ 4170.44 per MWh). All the rolling horizon runs satisfied the demand, excepting `RH CVaR` with $(\varepsilon_{\mathtt{p}}, \varepsilon_{\mathtt{c}}) = (0.4, 0.05)$, with a load shedding level of 0.153%.

Table 9 reports the total cost for policies `WS`, `SDDP`, and the "best" and "worst" `RH` policies. Except for the "worst" one, that leaves demand unsatisfied, all the rolling horizon policies had total mean cost smaller than the one obtained with `SDDP` (`SDDP` total mean cost is about 24% above the corresponding value for the "best" `RH` policy).

[7] `RH` mean cost is 19.3% higher than `WS` mean cost but only corresponds to a 0.8% increase with respect to `SDDP` mean cost

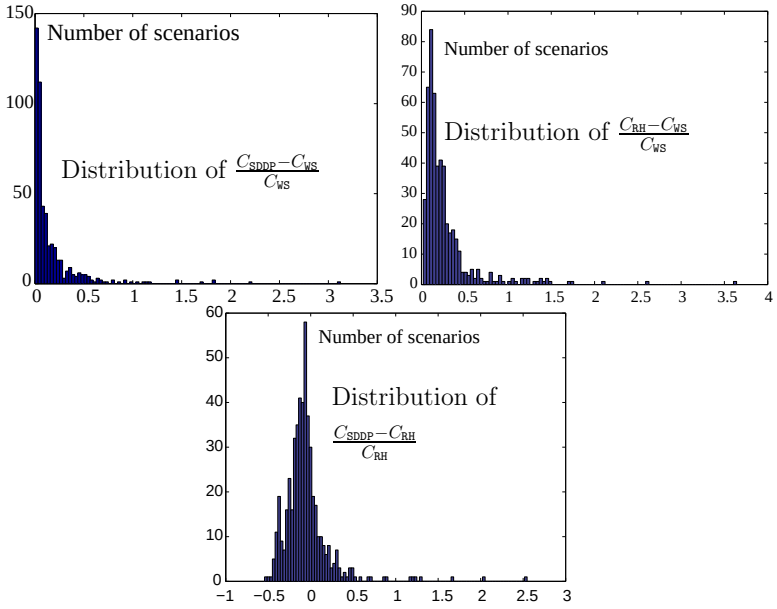

FIGURE 10. Empirical distribution of cost ratios obtained with the three policies.

TABLE 9. Central and dispersion characteristics of the total cost (R$)

Output	WS	Best RH CVaR $(\varepsilon_{\texttt{p}}, \varepsilon_{\texttt{c}}) = (0.45, 0.01)$	Worst RH CC-CVaR $(\varepsilon_{\texttt{p}}, \varepsilon_{\texttt{c}}) = (0.4, 0.05)$	SDDP
Mean	5.752×10^{8}	6.865×10^{8}	3.302×10^{10}	8.543×10^{8}
s.d.	2.059×10^{8}	1.744×10^{8}	8.146×10^{9}	2.185×10^{9}
VaR 1%	1.026×10^{9}	1.115×10^{9}	5.132×10^{10}	2.991×10^{9}
VaR 5%	9.273×10^{8}	9.911×10^{8}	4.672×10^{10}	1.272×10^{9}
VaR 90%	3.115×10^{8}	4.639×10^{8}	2.269×10^{10}	3.299×10^{8}

The mean generation cost per subsystem for the three policies is given in Table 10.

TABLE 10. Mean generation cost (over the scenarios) in the subsystems (R$).

Policy	South-East	South	North-East	North
WS	5.726×10^{8}	0	0	0
RH	6.823×10^{8}	2.056×10^{5}	1.527×10^{6}	0
SDDP	6.428×10^{8}	2.522×10^{7}	1.029×10^{7}	0

From these tables, and as already announced in Section **??**, for our data set and in terms of cost distribution, SDDP policy is far from an optimal (or a close to optimal) risk-neutral policy. Because the stochastic process is interstage dependent, the process history (7) gives in (1) a state vector with 24 components. Such state dimension is practically unmanageable by Dynamic Programming techniques for large horizons T, even in a deterministic setting. Uncertainty only worsens the situation, making SDDP simulation phase use an approximation of the recourse function that is likely to be poor at the considered states. Better approximations could be obtained by setting up an SDDP rolling horizon methodology, but, in view of the figures in Table 7, it is clear that this alternative is computationally impossible. Likewise, a risk-averse formulation, for example along

the lines of [**?**] or [**?**], needs more variables and constraints, and would increase the already heavy computational effort. Moreover, even if distributing calculations in many different processors made the computational load acceptable, implementing a risk-averse SDDP as in [**?**] would still require to find at each stage suitable weights defining a trade-off between the mean and the CVaR of the future cost. At first sight, choosing sound values for these weights for a specific problem is not a simpler task than finding adequate confidence levels $\varepsilon_{\mathtt{p}}$ and $\varepsilon_{\mathtt{c}}$ for the rolling horizon methodology.

6.8. **Exchanges, hydro-thermal generation, and marginal costs.** The most significant import-export exchanges go from the South-East to the South, for all policies. This is due to the fact that while the South has the highest demand at each time step, the South-East has both the highest hydro- (more than 70% of the total) and thermal-capacity. Other exchanges appear from the North to the North-East and the South-East, and from the North-East to the South-East. With `WS` and `RH` policies, and contrary to `SDDP`, the North receives no energy from the other subsystems. This subsystem, with no thermal plants and only one small reservoir, but also small demand, has a somehow "isolated" operation with both `WS` and `RH`. This is not the case with `SDDP` which, as shown in Figure 11, uses more water than the other two policies in the first half of the optimization period and, hence, ends up the needing to transfer some energy to the North.

We report in Figure 11 the average evolution for the South-East and South subsystems. Figure 12 displays the hydro-generation evolution of the whole system.

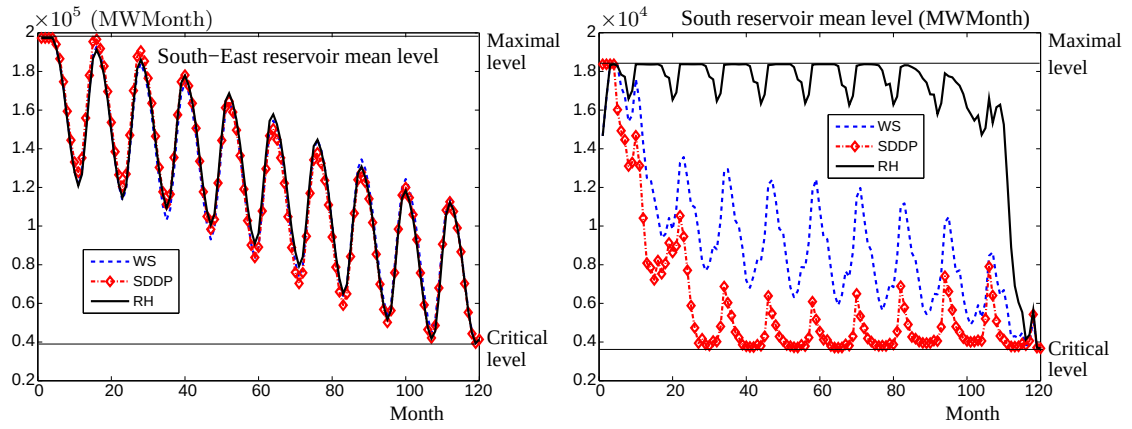

FIGURE 11. Mean reservoir level for South-East and South subsystems.

We see that, during the first five years, `SDDP` uses more water at nearly all time steps and scenarios. In particular, for `SDDP`, hydro-generation does not fluctuate much, likely because of the use of approximate recourse functions that are optimistic for the first years. This explains why `SDDP`'s reservoir levels and generation costs tend to be lower than the other policies over the first half of the optimization period. However, since this initial optimism also prevents `SDDP` from using as much water as `RH` and `WS` later on, eventually `SDDP` becomes more costly. With all the policies, both the S and SE reach (or nearly reach) their critical levels on all scenarios at the end of the optimization period.

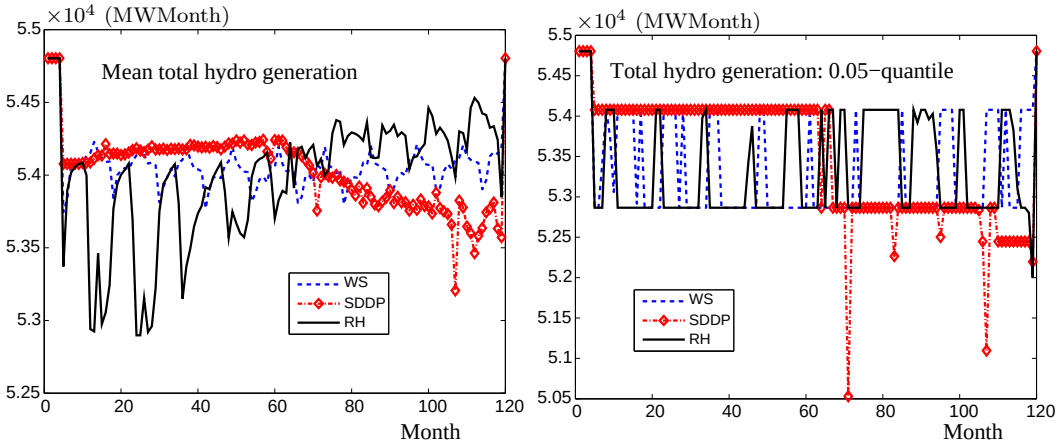

FIGURE 12. Mean hydro-generation (left) and 0.05-quantile (right).

The right graph in Figure 12, with the 0.05-quantile for the hydro-generation, shows a relatively low variability from scenario to scenario. All policies make a similar management of the hydro-power. However, over the first half of the optimization period, `RH` tends to be a bit more conservative with respect to the use of water and uses slightly less water than the other two policies. The conservatism of `RH` over the first years may be explained by the fact that, at the beginning of the process ($t \approx 1$), (11) needs to foresee a far away future ($\tau \in [2, 120]$), using robust uncertainty sets (19) and (20) that are overly pessimistic with respect to the final stages ($\tau \approx 120$). By contrast, over the last five years, as the number of uncertain parameters in (11) decreases, uncertainty sets become tighter, `RH` "realizes" that more water than necessary has been stored and it uses slightly more water than the other two policies.

Figure 13 plots the mean level and 0.05- and 0.95-quantiles for the total thermal generation. With both `RH` and `WS`, since the demand is satisfied for every time step and scenario, the thermal generation merely complements the hydro-generation to attain the demand level, a natural feature for a predominantly hydroelectric system.

It appears that `RH` policy uses more (resp. less) thermal power over the first half (resp. second half) of the optimization period; a phenomenon which is consistent with previous observations. For this reason, `RH` is more (resp. less) costly than `WS` over the first half (resp. second half) of the optimization period. Since `SDDP` uses less water in the second half of the optimization period, it needs to call on more thermal plants. It can also be seen (by comparing Figures 12 and 13) that the thermal generation is comparatively much smaller than the hydro-generation, for all policies, as expected in a hydro-dominated system (in our configuration, thermal power can cover at most 12.8% of the demand at each time step). When needed, thermal plants are committed in ascending order of their operational cost, to prevent load shedding. Due to the fact that only some thermal plants enter the system, the marginal cost (mean optimal Lagrange multipliers for the demand satisfaction constraint) in all subsystems is low for all the policies, except for some time steps with `SDDP`, as shown in Figure 8.

More precisely, at some time steps `SDDP` needs to call on more thermal plants than `RH` and `WS`, because `SDDP` uses less water, in particular in the second half of the optimization period. For such

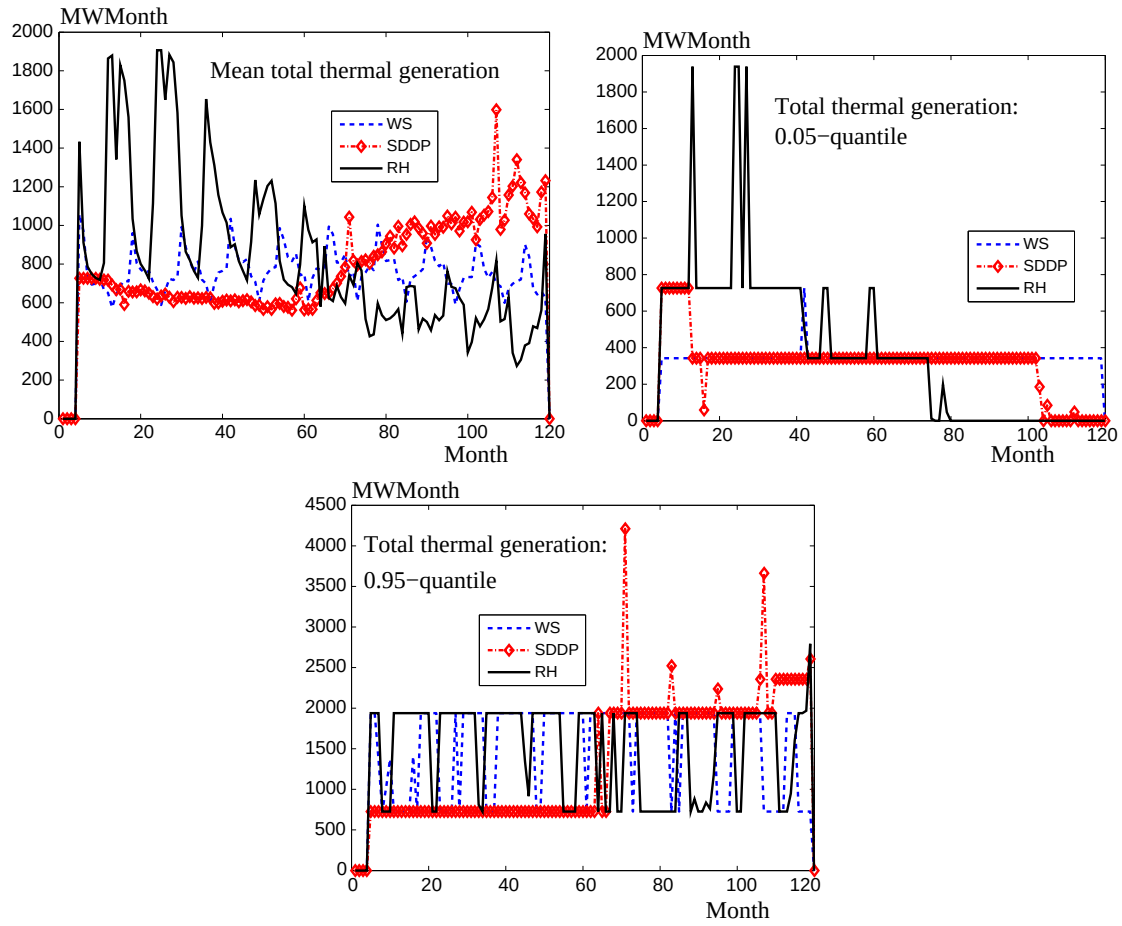

FIGURE 13. Total thermal generation: mean, 0.05-, and 0.95-quantiles.

time steps, SDDP's mean marginal cost can climb from around R$10-15 per MWh to R$ 38.7 per MWh, with peaks to R$ 4170.44 per MWh for some realizations. In fact, SDDP's variability of marginal cost (from one scenario to another) is quite significant, as reported in Figure 14.

Stability appears as another advantage of RH over SDDP: we obtain lower marginal costs that are less volatile. Notwithstanding, on average, both policies give relatively low values: the mean marginal costs in R$ per MWh obtained with RH (SDDP) policy amount to 10.39 (12.13), 10.39 (12.14), 9.37 (11.07), and 9.37 (11.08) for the SE, S, NE, and N subsystems, respectively.

FINAL CONSIDERATIONS

The novelty of this paper is in the specific integration of topics from risk modelling, dynamic programming and multistage stochastic optimization to arrive at mathematically sound solution methods working efficiently for large-scale models.

When the optimization horizon is large, (1) becomes intractable and very few policies can be implemented, even in a risk-neutral formulation. One of the rare exceptions is SDDP, which traverses a scenario tree by randomly choosing some paths. When T is large, such a scenario tree (generated by discretizing the continuous distribution of the stochastic process defining the uncertainty in (1))

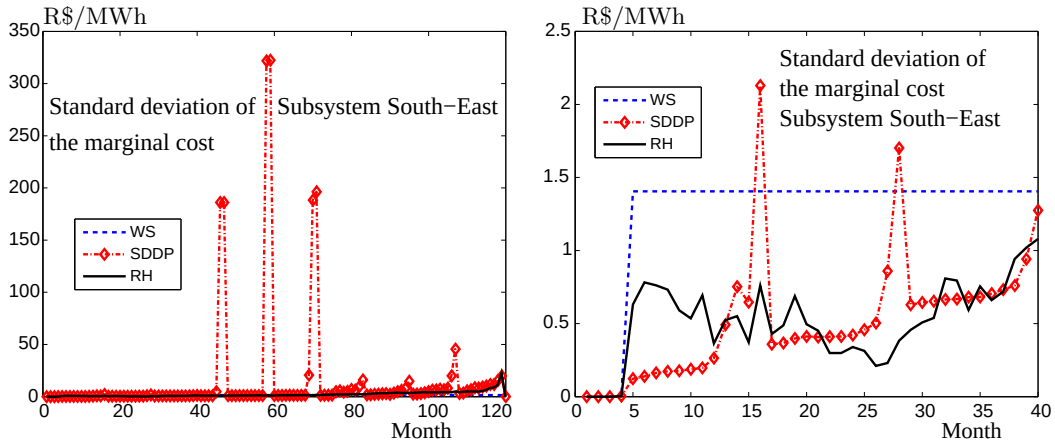

FIGURE 14. Marginal cost standard deviations for SE subsystem, with a zoom for the first 40 time steps.

has an astronomical number of paths. Since the method can only explore an infimal proportion of the tree (at the sake of a substantial computational effort), sometimes SDDP approximate policies can be poor. We observed indications of these features for some of the indicators analyzed in our case study, in particular those referring to volatility. For our test-case, the rolling horizon policy provided generation profiles quite close to the optimal ones. When compared to SDDP, average results are similar, but obtained with much less computational effort: from Table 7, we see that RH's overall CPU time was about 80 times smaller than SDDP's. With the rolling horizon approach, there is also a reduction in volatility, not only because the optimal decision at time step t depends *only* on the information provided by the t^{th} risk-averse problem, but also because our SDDP implementation was risk-neutral and is not re-solved at each time stage (a risk-averse rolling horizon formulation would have made SDDP computational times unacceptable). An interesting observation is that the gain in stability provided by RH did not increase the mean marginal cost.

Since we assume the model parameters in (6) to be known and since values of the inverse distribution F^{-1} in Theorem 4.4 can be obtained with arbitrary accuracy, the rolling horizon approach can fully exploit the probabilistic structure of uncertainty, without resorting to discretizations or approximations. If parameters in (6) are only estimations, or just approximate generalized inverse distributions are available, then the rolling horizon policy would be an approximation. The question of how to control the quality of the approximate policy in this case is an interesting topic, combining statistical inference and sensitivity analysis.

Since we can always suppose, at least from a theoretical point of view, that the objective function of a stochastic optimization problem is uncertainty independent (problems $\min\{f(x,\tilde{\xi}) : x \in X\}$ and $\min\{t : t \geq f(x,\tilde{\xi}), x \in X\}$ are equivalent), our analysis also applies to stochastic optimization problems with feasible set as in (1), but with objective function depending affinely on the uncertain parameters. Furthermore, the methodology would still be applicable if in (1) a linear term $H_t x_{t-1}$ was introduced in the right hand side of (INEQ), or for matrices B_t similar to the technology matrices F_t in (2). For the water-resource application, in particular, these extensions would make

possible to consider uncertain demand and generation costs, as in [**?**], [**?**], as long as the rolling horizon risk-averse subproblems remain tractable.

Likewise, since both the mean and the standard deviation of the cost were close to the optimal values in the numerical experiment, it would be interesting to find theoretical upper bounds for the bias and variance of the total cost.

Finally, instead of CVaR constraints, one could use integrated chance constraints (ICC), as in [20]:

$$\mathbb{E}[\min(X, h_\tau)] \geq h_\tau - \varepsilon_{\mathtt{c}}(|h_\tau| + 1),$$

because the set of ICC controls remains closed and convex. For comparison, our CVaR constraints can be written in the form $\mathbb{E}[X|X \leq \tilde{h}_\tau] \geq h_\tau - \varepsilon_{\mathtt{c}}(|h_\tau| + 1)$, by taking $\tilde{h}_\tau = -VaR_{\varepsilon_{\mathtt{p}}}(X)$ (the connection with ICC is given by the relation $\mathbb{E}[\min(X, x)] = x\mathbb{P}(X > x) + \mathbb{E}[X|X \leq x]\mathbb{P}(X \leq x)$). However, even for null technology matrices, ICC would result in convex non-linear optimization problems (still tractable, but no longer linear programs).

Acknowledgments. We would like to thank the three reviewers and the Associate Editor for beneficial comments and suggestions.

References

[1] N. Beaulieu, A. Abu-Dayya, and P. MacLane. Estimating the distribution of a sum of independent lognormal random variables. *IEEE Transactions on Communications*, 43(12):2869–2873, 1995.

[2] A. Ben-Tal, A. Goryashko, E. Guslitzer, and A. Nemirovski. Adjustable robust solutions of uncertain linear programs. *Mathematical Programming*, 99:351–376, 2004.

[3] A. Ben-Tal and A. Nemirovski. Robust solutions of uncertain linear programs. *Oper. Res. Letters*, 25:1–13, 1999.

[4] A. Ben-Tal and A. Nemirovski. Selected topics in robust convex optimization. *Math. Program.*, 112(1):125–158 (electronic), 2008.

[5] D. Bertsimas and A. Thiele. A Robust Optimization Approach to Inventory Theory. *Operations Research*, 54(1):150–168, 2006.

[6] J. Birge and F. Louveaux. *Introduction to Stochastic Programming*. Springer-Verlag, New York, 1997.

[7] H. Chao, A. Poncet, and C. Yuen. Handling infeasibilities when applying Benders decomposition to scheduling optimization. *WSEAS Transactions on Power Systems*, 1(7):1209–1216, 2006.

[8] A. Charnes and W. Cooper. Chance-constrained programming. *Management Sci.*, 6:73–79, 1959/1960.

[9] A. Charnes and W. Cooper. Chance constraints and normal deviates. *J. Amer. Statist. Assoc.*, 57:134–148, 1962.

[10] A. Charnes and W. Cooper. Deterministic equivalents for optimizing and satisficing under chance constraints. *Operations Res.*, 11:18–39, 1963.

[11] Z.-L. Chen and W.B. Powell. A convergent cutting-plane and partial-sampling algorithm for multistage stochastic linear programs with recourse. *Journal of Optimization Theory and Applications*, 102:497–524, 1999.

[12] T. Croley. A manual for hydrologic time series deseasonalization and serial dependence reduction. *Iowa Institute of Hydraulic Research report, The University of Iowa*, 199, 1977. `http://www.iihr.uiowa.edu/products/pubvid/publications.php`.

[13] D. Dentcheva and A. Ruszczyński. Optimization with stochastic dominance constraints. *SIAM J. Optim.*, 14(2):548–566 (electronic), 2003.

[14] D. Dentcheva, A. Ruszczyński, and A. Shapiro. *Lectures on Stochastic Programming*. SIAM, Philadelphia, 2009.

[15] A. Eichhorn and W. Römisch. Dynamic risk management in electricity portfolio optimization via polyhedral risk functionals. *Power and Energy Society General Meeting - Conversion and Delivery of Electrical Energy in the 21st Century, 2008 IEEE*, pages 1–8, July 2008.

[16] M. Fodstad, K. Midthun, F. Romo, and A. Tomasgard. *Geometric Modelling, Numerical Simulation, and Optimization*, chapter Optimizatin Models for the Natural Gas Value Chain, pages 521–558. SINTEF, 2007.

[17] V. Guigues. Robust production management. *Optimization and Engineering*, 10:505–532, 2009.

[18] V. Guigues and C. Sagastizábal. Exploiting structure of autoregressive processes in risk-averse multistage stochastic linear programs. Technical report, IMPA, 2010. `http://www.preprint.impa.br/Shadows/SERIE_D/2010/72.html`.

[19] W. Klein Haneveld. *Duality in stochastic linear and dynamic programming, volume 274 of Lecture notes in Economics and Mathematical Systems*. Springer-Verlag,Berlin, 1986.

[20] W. Klein Haneveld and M. Vlerk. Integrated Chance Constraints: Reduced Forms and Algorithm. *Computational Management Science*, 3(4):245–269, 2006.

[21] D. Kuhn, W. Wiesemann, and A. Georghiou. Primal and dual linear decision rules in stochastic and robust optimization. *Mathematical Programming*, 2011.

[22] M.E. Maceira and J. Damázio. The use of PAR(p) model in the stochastic dual dynamic programming optimization scheme used in the operation planning of the Brazilian hydropower system. *VIII International Conference on Probabilistic Methods Applied to Power Systems*, pages 397–402, Sept. 2004.

[23] M.E. Maceira, J. Damázio, F. Costa, and A. Melo. Chain of Models for Setting the Energy Dispatch and Spot Price in the Brazilian Power System. *System Computation Conference - PSCC'02, Sevilla, Spain*, pages 1–4, June 2002.

[24] M. Pereira and L. Pinto. Multi-stage stochastic optimization applied to energy planning. *Math. Program.*, 52(2):359–375, 1991.

[25] G. Pflug and W. Römisch. *Modeling, measuring and managing risk*. World Scientific Publishing Co. Pte. Ltd., Hackensack, NJ, 2007.

[26] A. Philpott and Z. Guan. On the convergence of stochastic dual dynamic programming and related methods. *Oper. Res. Letters*, 36(4):450–455, 2008.

[27] A. Prékopa. *Stochastic programming*, volume 324 of *Mathematics and its Applications*. Kluwer Academic Publishers Group, Dordrecht, 1995.

[28] R.T. Rockafellar and S. Uryasev. Conditional Value-at-Risk for General Loss Distributions. *Journal of Banking and Finance*, 26(7):1443–1471, 2002.

[29] A. Ruszczyński and A. Shapiro, editors. *Stochastic Programming*, volume 10 of *Handbooks in Operations Research and Management Science*. Elsevier Science B.V., Amsterdam, 2003.

[30] A. Ruszczyński and A. Shapiro. Conditional risk mappings. *Math. Oper. Res.*, 31(3):544–561, 2006.

[31] A. Shapiro. On a time consistency concept in risk averse multi-stage stochastic programming. *Oper. Res. Letters*, 37:143–147, 2009.

[32] J. Wu, N. Mehta, and J. Zhang. A flexible lognormal sum approximation method. *Proc. IEEE Global Telecommunications Conference GLOBECOM'05*, 6:3413–3417, 2005.

[33] M.K. Zanjani, M. Nourelfath, and D. Aït-Kadi. A multi-stage stochastic programming approach for production planning with uncertainty in the quality of raw materials and demand. Technical report, CIRRELT, 2009.

Vincent Guigues: PUC-Rio, Departamento de Engenharia Industrial, Rua Marquês de São Vicente, 225, Gávea, Rio de Janeiro, Brazil, `vguigues@puc-rio.br` and IMPA, Instituto de Matemática Pura e Aplicada, 110 Estrada Dona Castorina, Jardim Botanico, Rio de Janeiro, Brazil, `vguigues@impa.br` and Claudia Sagastizábal: CEPEL, Electric Energy Research Cent er, Eletrobrás Group, On leave from INRIA Rocquencourt, France, `sagastiz@impa.br`.