THE VALUE OF ROLLING-HORIZON POLICIES FOR RISK-AVERSE HYDRO-THERMAL PLANNING

VINCENT GUIGUES[A] AND CLAUDIA SAGASTIZÁBAL[B]

[A] PUC-Rio, Departamento de Engenharia Industrial
Rua Marquês de São Vicente, 225, Gávea, Rio de Janeiro, Brazil,
Tel: 55 21 3527 1284, Fax: 55 21 3527 1288, 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 [B] CEPEL, Electric Energy Research Center, Eletrobrás Group, On leave from INRIA
Rocquencourt, France, sagastiz@impa.br.

Abstract. We consider the optimal management of a hydro-thermal power system in the mid and long terms. From the optimization point of view, this amounts to a large-scale multistage stochastic linear program, often solved by combining sampling with decomposition algorithms, like stochastic dual dynamic programming. Such methodologies, however, may entail prohibitive computational time, especially when applied to a risk-averse formulation of the problem. We propose instead a risk-averse rolling-horizon policy that is nonanticipative, feasible, and time consistent. The policy is obtained by solving a sequence of multi-stage problems with deterministic constraints for the current time step and future chance and CVaR constraints.

The considered hydro-thermal model takes into account losses resulting from run-of-river plants efficiencies as well as uncertain demand and streamflows. Constraints aim at satisfying demand while keeping reservoir levels above minzones almost surely. We show that if the problem uncertainty is represented by a periodic autoregressive stochastic process with lag one, then the probabilistic constraints can be computed explicitly. As a result, each one of the aforementioned multi-stage problems is an easy to solve medium-size linear program.

For a real-life power system we compare our approach with three alternative policies. Namely, a robust nonrolling-horizon policy and two risk-neutral policies obtained by stochastic dual dynamic programming, implemented in nonrolling- and rolling-horizon modes, respectively. Our numerical assessment confirms the superiority of the risk-averse rolling-horizon policy that yields comparable average indicators, but with reduced volatility and with substantially less computational effort.

Stochastic programming and Chance constraints and Interstage dependence and Rolling horizon and Hydro-thermal planning

AMS subject classification: 90C15, 91B30.

*Corresponding author. Email: vguigues@puc-rio.br. The first author's research was supported by CNPq grant number 382.851/07-4. Reasearch of the second author partially supported by CNPq and Faperj.

1. Introduction

The optimal operation of a hydrothermal system in the mid and long terms usually minimizes the expected value of the operating cost, essentially composed of fuel costs plus penalties for load shedding. This type of problem is of interest not only in centralized systems, but also for ISO and private and institutional agents acting on vertically integrated systems.

When the planning horizon covers several years, the optimal operation problem (OOP) is formulated as a stochastic program with recourse; see Birge and Louveaux (1997). The problem is often large scale because there are many power plants and many time periods need to be considered. For a study over 5 years, a typical time step of one month -needed to describe suitably the hydrological uncertainty- makes up a total of 60 months. To eliminate “boundary” effects, see Section 5.1.1 below, in practice the time horizon is in fact doubled, yielding $T = 120$ stages. For systems subject to cold winters, uncertainty in the seasonal demand needs to be represented too, increasing even more the problem complexity. For these reasons, a linear modeling is often adopted for the optimization problem.

A prototypical example is Brazil’s power system, predominantly hydroelectric, for which the availability of limited amounts of hydro-power in the form of water stored in reservoirs makes the OOP very complex. There are many reservoirs in cascade, some of them with a capacity of regularization that covers several years, and spread over geographical regions with different seasonal rainfall. Water is a commodity of unknown value and uncertain availability and present operating decisions have future consequences that are difficult to quantify. In this setting, important indicators obtained when solving the OOP, like mean marginal prices or the average future unsupplied energy, need to reflect well the impact of extreme events such as extended droughts. This is why it is interesting to develop OOP variants that not only consider average phenomena, well represented by the cost in expected value, but also take into account the underlying risk factors. Since the *hydrological risk* is mostly associated to low volumes of stored water, operators struggle to keep the reservoirs over some critical minimum values or *minzones* that trigger an alert for the system as a whole, and require special actions to avoid future blackouts.

The question of how to measure risk in financial applications has been intensively developed over the past years, Artzner et al. (1999), Rockafellar and Uryasev (2000). In the energy sector, similar ideas have been applied for hedging financial risk of price-taker companies operating in deregulated markets (we refer to Mo et al. (2001), Wallace and Fleten (2003), Conejo et al. (2004), Cabero et al. (2005), Eichhorn et al. (2004), Eichhorn and Römisch (2008), and references therein). The problem of measuring *hydraulic risk* for the OOP was first explored in Brignol and Rippault (1999), and more recently in Guigues et al. (2009), Guigues and Sagastizábal (2009), Andrieu et al. (2010), Guigues (2011), Philpott and de Matos (2010), Shapiro (2011).

For problems with uncertain data, two classes of models can be proposed: Robust Optimization or Stochastic Programming. In a multistage setting, Stochastic Programming may lead to intractable problems. Already in a risk-neutral formulation, the huge number of time stages of the OOP can only be dealt with by sampling, as in Pereira and Pinto (1991), Chen and Powell (1999), Philpott and Guan (2008). Risk aversion makes the situation only worse, due to the need of more variables and/or constraints to model the problem; see Guigues and Römisch (2010), Philpott and de Matos (2010), Shapiro (2011). By contrast, the Robust Optimization techniques from Ben-Tal and Nemirovski (1998) yield tractable problems; for instance, the adjustable robust counterparts Ben-Tal et al. (2003) give a tractable OOP model in Guigues (2009). But in general the quality of a robust policy is only good for a problem with a moderate number of stages, like in Guigues (2009). As confirmed by our numerical results in Section 5.2, since uncertainty sets become progressively larger as the time stages increase, a robust policy may end up being exceedingly conservative.

To circumvent the curse of intractability and/or conservatism, inherent to multistage uncertain programs, we employ a risk-averse approach in a *rolling-horizon* setting, as in Guigues and Sagastizábal (2009), described in Section 2. Essentially, we consider solving successively $T-1$ multi-stage problems with shorter and shorter horizons, for $t = 1, \ldots, T$. For the t^{th} multi-stage problem, constraints are considered deterministic at time step t, and in a probabilistic sense for $[t+1, T]$. The approach builds *feasible* policies: all the constraints over the optimization period hold almost surely. This is a very important property for the OOP, for which the minzones constraints need to be satisfied with probability one. Moreover, the policy is *time consistent* in the sense of Shapiro (2009).

The considered hydro-thermal model considers uncertain streamflows and includes run-of-river plants, that can either spill or generate; see Section 3. Since turbines have a maximum capacity, some of the uncertain streamflow, meant to be immediately transformed into energy by the run-of-river plants, may be lost. As explained in Section 3.1, the corresponding plant efficiency induces a loss of energy that depends on the whole system configuration and that is represented by a nondecreasing function of the streamflow. Constraints aim at satisfying demand while keeping reservoir levels above minzones almost surely. In Section 4 we show that if the streamflow is represented by a periodic autoregressive stochastic process with lag one, then the probabilistic constraints can be computed explicitly. As a result, each one of the aforementioned multi-stage problems is an easy to solve medium-size linear program. The model also takes into account uncertain demand. In Section 4.4 we address the general case, when demand is also a stochastic autoregressive process with lag one, for which calculations are no longer explicit, but involve estimations.

With respect to Guigues and Sagastizábal (2009), this work includes the modeling of both run-of-river plant efficiencies and uncertain demand. Although realistic, this more general framework complicates the recursive computation of the chance constraint coefficients derived in Guigues and Sagastizábal (2009). In this paper we show how to make those computations explicit, in a direct manner, for autoregressive processes with lag equal to one. In addition, the theoretical development is supported by the final Section 5, with numerical results organized in three parts. First, to emphasize the fact that rolling horizons are indeed beneficial in a risk-averse setting, we (favorably) compare our policy with a robust one, similar to the one by Bertsimas and Thiele (2006). The second set of tests evaluates the impact of modeling the run-of-river efficiencies, and compares our policy with two stochastic dynamic dual programming (SDDP) policies, implemented in nonrolling- and rolling-horizons, but risk neutral. The third and final test-case refers to a real-size OOP, and confirms the superiority of the risk-averse rolling-horizon policy which yields average indicators that are comparable to the ones obtained by a risk-neutral SDDP, but with reduced volatility and with substantially less computational effort.

We now set down some notation. For a random variable ξ, $\tilde{\xi}$ denotes a particular realization, whereas $\mathbb{E}(\xi)$ and $\sigma(\xi)$ are the expected value and the standard deviation, respectively. For the process ξ, $\xi_{[t]} = (\xi_j, j \le t)$ denotes its history up to time t. Conditional expectations and probabilities are denoted by $\mathbb{E}(\xi_1|\tilde{\xi}_2) := \mathbb{E}(\xi_1|\xi_2 = \tilde{\xi}_2)$ and $\mathbb{P}(\xi_1 \in A|\tilde{\xi}_2) := \mathbb{P}(\xi_1 \in A|\xi_2 = \tilde{\xi}_2)$. The cumulative distribution function is denoted by $F_\xi(\cdot)$, knowing that for $\xi \sim \mathcal{N}(0,1)$, we just write $F(\cdot) := F_\xi(\cdot)$. The generalized inverse of a nondecreasing function F is given by $F^{\leftarrow}(\varepsilon_{\mathtt{p}}) = \inf\{x : F(x) \ge \varepsilon_{\mathtt{p}}\}$. For a continuous random variable X for which higher values are preferred, the Conditional Value-at-Risk of level $\varepsilon_{\mathtt{p}} \in (0,1)$ of X is defined by $CVaR_{\varepsilon_{\mathtt{p}}}(X) = -\mathbb{E}[X|X \le F_X^{-1}(\varepsilon_{\mathtt{p}})]$ while the Value-at-Risk of level $\varepsilon_{\mathtt{p}}$ of X is $VaR_{\varepsilon_{\mathtt{p}}}(X) = -F_X^{-1}(\varepsilon_{\mathtt{p}})$.

2. Main features of a rolling-horizon approach

As explained in Wets (2000), in a *here-and-now* approach, a decision must be selected before a realization of the random data becomes available. In a problem of the *wait-and-see* type, one is allowed to wait before making the decision until realizations of the random variables can be observed; see also Ch. I, Birge and Louveaux (1997). The risk-averse rolling-horizon methodology lies in-between and could be labeled "here-and-now looking-forward".

In order to explain the rolling-horizon (RH) approach, we first consider how a sampling method, like SDDP, attempts to solve a stochastic program over a multi-stage scenario tree. Since solving the nested optimization problem for all possible scenarios is not possible, SDDP samples a small subset of scenarios to build a policy. The policy is defined by a piecewise linear lower approximation of the future cost for each stage, which is improved along iterations. Each SDDP iteration consists of a *forward* and a *backward* step, illustrated by Figure 1 below.

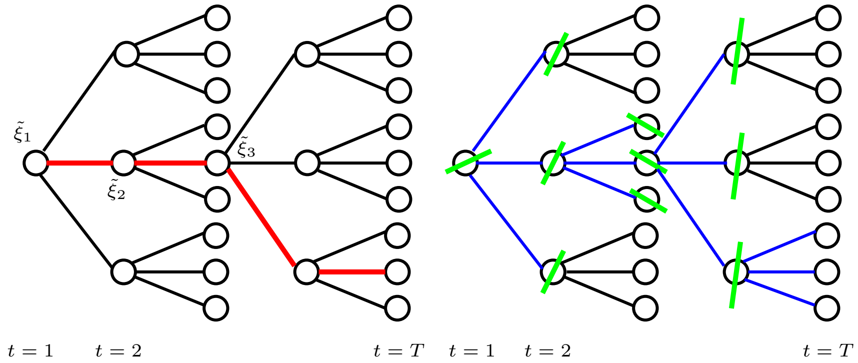

FIGURE 1. Forward and Backward steps in SDDP.

The forward step gives new states and costs at each node of the sampled scenarios, drawn in red and denoted by $(\tilde{\xi}_1, \tilde{\xi}_2, \tilde{\xi}_3, \ldots, \tilde{\xi}_T)$ in the graph on the left of the figure. The backward step adds cuts to each future stage cost function approximation. The graph on the right in Figure 1 shows these cuts in green, built using all the uncertainty information (marked with blue lines) that is available along the red path (cuts are shared between nodes of the same stage).

To avoid the dimensionality explosion typical of multi-stage problems, SDDP randomly traverses the tree in the forward step (that is, taking different red lines) along iterations. However, when uncertainty is interstage dependent, the curse of dimensionality remains an issue because state variables need to be augmented with the uncertainty history. The inclusion of risk measures can only increase the number of state variables, so the situation is even worse for risk averse variants of SDDP, like Guigues and Römisch (2010), Philpott and de Matos (2010), or Shapiro (2011). Because of its computational complexity, even when set in a risk-neutral formulation, in practice SDDP is often stopped after a few iterations, by using a loose criterion. For example, terminating when certain lower bound for the cost, computed at the forward pass, appears as if having stabilized around some level. Rather than seeking for optimality, SDDP emphasizes a rich representation for uncertainty: under the assumption of relatively complete recourse, SDDP policies are always feasible, but also sub-optimal.

The very rich information on uncertainty is not fully exploited by SDDP from the optimization point of view. For this reason, we consider an alternative approach that makes use of chance- and CVaR-constrained problems, defined in a *rolling-horizon* setting. This is done by defining a sequence of multi-stage programs with shorter and shorter horizons, illustrated by Figure 2.

An RH model considers $T-1$ successive risk-averse problems, each one defined for a time $t = 1, 2, \ldots, T-1$. In the figure, the top left and right graphs correspond to the first and second problems, respectively, while the bottom graphs represent the $(T-1)^{th}$ and T^{th} ones (noting that the last problem is just a one stage deterministic problem). For $t < T$, the t^{th} problem is a multi-stage program defined over the horizon $[t, T]$, with the first stage given by present time t (with

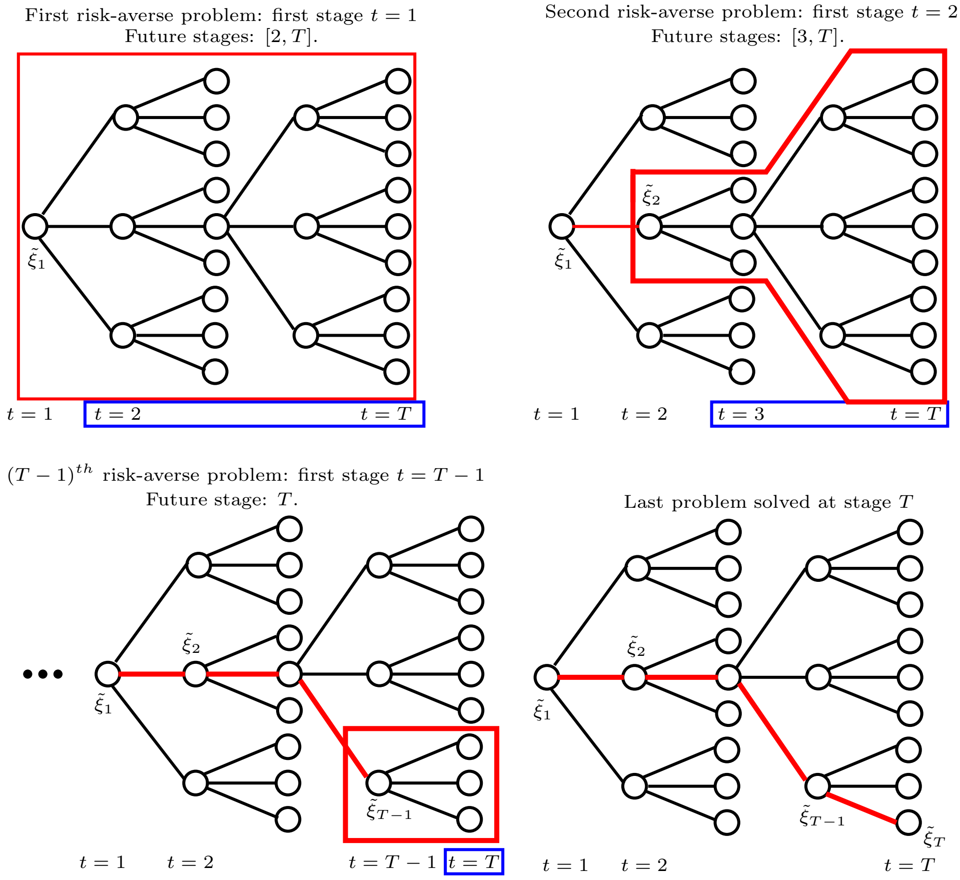

FIGURE 2. Rolling Horizon approach

deterministic constraints), and subsequent stages covering the future $[t+1, T]$ (with uncertain constraints). In the figure, time steps corresponding to future stages are shown with a blue box. Similarly, the considered subtrees are surrounded by a red polyhedron. Each subtree is defined by the portion of the full scenario tree, "descending" from the first stage node until final stage T.

Here arises an important difference: an RH implementation of SDDP would sample only *some forward paths* over such a subtree. By contrast, our model considers the *full subtree* uncertainty, by setting chance and CVaR constraints for future stage constraints. When uncertainty depends on the past in an affine manner, like in autoregressive models, Section 4 shows how to reformulate future constraints explicitly so that each risk-averse problem is just a deterministic linear program.

Since this sequence of $T-1$ multi-stage problems makes a poorer representation of uncertainty than a T-stage program, our risk-averse RH policy is still sub-optimal. Nevertheless, thanks to the explicit formulation of chance constraints that represent uncertainty of each full subtree (instead of just parsing a few forward paths, as in SDDP), our risk-averse policy is "closer" to the wait-and-see one. This is confirmed by our numerical results in Section 5, which also show the substantial difference in computational effort: for a typical real-life case, over a set of 500 scenarios, our approach takes 6h while a nonrolling-horizon SDDP takes 3 weeks. As a result, even though if in principle embedding SDDP in a rolling-horizon setting might produce a good policy, when T is large the computational complexity of SDDP makes such an approach impossible, at least with the computational capacity available nowadays.

3. A simplified optimal operation problem with energy loss representation

Our approach is general and applies to real-life hydro-thermal systems, like the Brazilian one. However, for the sake of clarity, the mathematical formulation leading to the chance-constrained problem is first given for a simplified hydro-thermal system. We consider only one reservoir, one run-of-river plant, and 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. However, the run-of-river plant has turbines whose capacity may not suffice to convert all the streamflow into power.

3.1. **Energy losses and hydro-thermal complement.** In our simplified formulation, at time step t ($1 \leq t \leq T$) we have a state variable x_t, the volume of the reservoir at the beginning of the time step; a nonnegative control variable u_t, the turbine outflow; and ξ_t, the natural streamflow of water arriving into the reservoir. Only a fraction $\gamma_t \in [0,1]$ of this water can be stored, the remaining portion, $(1-\gamma_t)\xi_t$, is in principle immediately transformed into power by the run-of-river plant. However, due to capacity limits, it may not be possible to turbine all of $(1-\gamma_t)\xi_t$, resulting in some loss of energy. We represent such losses by a convex function, $\mathcal{L}_t(\cdot)$. Several representations are possible; here we consider a piecewise-defined function, depending on certain parameters a, b, and $L > 0$, satisfying $0 < a < \frac{1}{L}$ and $b = \frac{1+aL}{2L}$. The special choice of b makes the loss function differentiable at this point.

These parameters define three regions, R_1, R_2, R_3 such that

$$\mathcal{L}_t(x) = \begin{cases} 0 & \text{if } x \in R_1 := (-\infty, a], \\ Lx(x-a) & \text{if } x \in R_2 := [a,b], \\ Lb(b-a) + x - b & \text{if } x \in R_3 := [b, +\infty)\,. \end{cases}$$

Figure 3 represents a typical loss function for $(a,b,L) := (5, 7.5, 0.1)$.

The loss function determines the difference $(1-\gamma_t)\xi_t - \mathcal{L}_t((1-\gamma_t)\xi_t)$ that will effectively be converted into power by the run-of-river plant. When the run-of-river energy is greater than or equal to the demand, no additional generation is necessary. In this sense, denoting by dem_t the

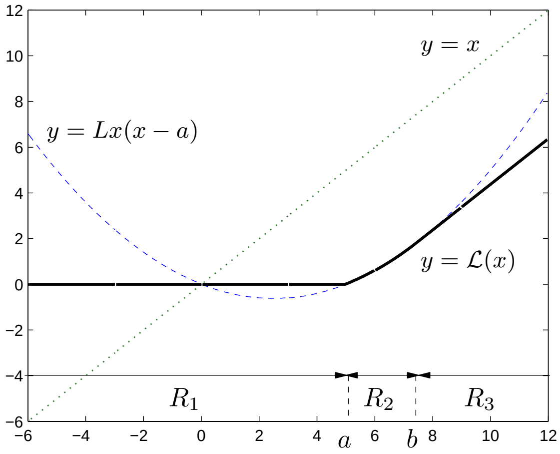

FIGURE 3. A typical loss function for a run-of-river power plant

demand arriving at the system at time step t, the function

$$\tilde{\mathcal{L}}_t(x) := \mathrm{dem}_t - (x - \mathcal{L}_t(x)) = \begin{cases} \mathrm{dem}_t - x & \text{if } x \le a, \\ \mathrm{dem}_t + Lx^2 - (aL+1)x & \text{if } a \le x \le b, \\ \mathrm{dem}_t + Lb(b-a) - b & \text{if } x \ge b\,, \end{cases} \tag{1}$$

gives the *hydro-thermal complement*:

$$\mathtt{HTcomp}(\mathrm{dem}_t, \xi_t) := \max(0, \mathrm{dem}_t - (1-\gamma_t)\xi_t + \mathcal{L}_t((1-\gamma_t)\xi_t)) = \max(0, \tilde{\mathcal{L}}_t((1-\gamma_t)\xi_t))\,, \tag{2}$$

to be produced by the remaining power plants.

3.2. **Problem formulation.** For $t = 1, \ldots, T$ the main constraints for the optimal operation problem (OOP) are given below.
Water balance equation:

$$x_{t+1} = x_t - u_t + \gamma_t \xi_t. \tag{3}$$

Demand satisfaction: At time step t, let df_t denote the nonnegative energy deficit, modeled as a (fictitious) thermal plant with large enough capacity and generation cost equal to the cost of deficit, a known data denoted by cd_t. Then the identity $u_t + df_t = \max\Big(\mathrm{dem}_t - (1-\gamma_t)\xi_t + \mathcal{L}_t((1-\gamma_t)\xi_t), 0\Big)$ implies satisfaction of the relations

$$\begin{cases} u_t + df_t \ge \mathrm{dem}_t - (1-\gamma_t)\xi_t + \mathcal{L}_t((1-\gamma_t)\xi_t) \\ u_t + df_t \ge 0\,. \end{cases}$$

Since, by definition, generation variables are nonnegative, writing demand constraints as

$$u_t + df_t + (1-\gamma_t)\xi_t - \mathcal{L}_t((1-\gamma_t)\xi_t) \ge \mathrm{dem}_t \tag{4}$$

will not change the optimal value of the OOP. Moreover, the particular structure of the optimal operation planning problem (6) below is such that, for a given time step, if the run-of-river production is less than the demand, the corresponding demand constraint is active.

Critical minimal volume: Operators managing the system in real time are mostly concerned with keeping reservoirs at reasonable storage levels. In particular, they sometimes wish to keep the reservoirs above critical values, or reference trajectories, estimated empirically, and denoted by $x_{t+1}^{crit} \geq 0$. These trajectories, also called *minzones*, are sometimes imposed by some regulatory rules fixed by the ISO, and define the constraint:

$$x_{t+1} \geq x_{t+1}^{crit}. \tag{5}$$

For this simplified system, and over a time horizon of T steps, the deterministic OOP problem has the form

$$\begin{cases} \min \sum_{t=1}^{T} [cu_t^\top u_t + cd_t^\top df_t] \\ (3), (4), (5) \quad t = 1, \ldots, T \\ (u_t, df_t) \in \mathcal{S}_t, \text{ a polyhedral bounded set, independent of } \xi_t, \text{ dem}_t \text{ for } t = 1, \ldots, T. \end{cases} \tag{6}$$

In the objective function above, cu_t is the hydro generation cost, assumed to satisfy $0 < cu_t < cd_t$. For the general case of a power mix with many plants, the vector cu_t (resp. u_t) stores both the unit costs for thermal plants and for energy transfers between subsystems (resp. thermal generation and import/export exchanges).

3.3. **Simplified statistical setting.** Uncertainty in (6) appears through ξ_t and dem_t, the natural streamflow and demand. In order to appropriately reflect seasonal variations, both processes are represented by periodic autoregressive models.

For the sake of simplicity again, we consider here that the twelve orders of the PAR model for the streamflow are equal to one: each month's streamflow depends only on the previous month's rains.

More precisely, letting μ_t and σ_t denote the mean and the (finite and positive) standard deviation of ξ_t, consider the standardized random variable

$$Z_t = \frac{\xi_t - \mu_t}{\sigma_t}. \tag{7}$$

Then there exists a coefficient Φ_{t+1} such that

$$Z_{t+1} = \Phi_{t+1} Z_t + \eta_{t+1}, \tag{8}$$

where $\eta_1, \ldots, \eta_T$ are independent Gaussian random variables with standard deviations $\sigma(\eta_t) = \sigma_t^\eta > 0$. For $t = 1$, ξ_1 is a known data, often referred to as *hydrological tendency* and, because the model is periodic, the functions μ_t, σ_t, Φ_t and σ_t^η are 12-periodic (one year period).

The recursive application of (8) gives for any time step t and $j \geq 1$:

$$Z_{t+j} = \left(\prod_{\ell=t+1}^{t+j} \Phi_\ell \right) Z_t + \sum_{\ell=t+1}^{t+j} \left(\prod_{k=\ell+1}^{t+j} \Phi_k \right) \eta_\ell \, ,$$

with the convention that $\prod_{k=k_1}^{k_2} \Phi_k = 1$ whenever $k_1 > k_2$. This relation, together with (7), gives the following expression for the streamflows

$$\xi_{t+j} = \alpha_{t,j} \xi_t + \sum_{\ell=t+1}^{t+j} \beta_{t+j}^\ell \eta_\ell + \theta_{t,j} \tag{9}$$

for $t = 1, \ldots, T-1$, $j = 1, \ldots, T-t$, and with

$$\alpha_{t,j} := \frac{\sigma_{t+j}}{\sigma_t} \prod_{\ell=t+1}^{t+j} \Phi_\ell, \quad \beta_{t+j}^\ell := \sigma_{t+j} \prod_{k=\ell+1}^{t+j} \Phi_k, \quad \theta_{t,j} := \mu_{t+j} - \alpha_{t,j} \mu_t. \tag{10}$$

The fact of considering lags greater than one complicates the coefficient expressions, but the corresponding relation (9) remains affine, on both the past values $(\xi_j, j \leq t)$ and the noises $(\eta_{t+1:t+j})$. Recursive relations for computing recursively coefficients α, β, and θ in the general case can be found in Guigues and Sagastizábal (2010). In this work we give instead explicit (not too involved) relations that can be derived directly for processes with lag equal to one.

The demand process being also periodic autoregressive, dem_{t+j} has an expression similar to the one in (9), *mutatis mutandis*. To ease the presentation, first we consider the demand to be deterministic, leaving to Section 4.4 the explanation of how to handle uncertainty in the demand.

4. Risk-averse rolling-horizon feasible policy

As illustrated by Figure 2, at a given stage t, the risk-averse RH policy is based on a here-and-now looking-forward approach, in which the data for the current time stage is deterministic while the future is considered uncertain, but depending on the history of realizations until time step t.

4.1. **Policy definition.** Each multi-stage problem in Figure 2 has uncertain future constraints dealt with by using probabilistic and CVaR constraints. For this reason, we express future states x_{t+j+1} in terms of the current state and past uncertainty, by applying recursively (3):

$$x_{t+j+1} = x_t - \sum_{i=0}^{j} u_{t+i} + \sum_{i=0}^{j} \gamma_{t+i} \xi_{t+i}, \; j = 1, \ldots, T-t. \tag{11}$$

With the above expression, knowing the first state variable x_t, at time stage t the feasible set of our t^{th}-stage risk-averse problem is given by

$$
(12)\qquad \begin{cases} (u_t, df_t, \dots, u_T, df_T) \in \prod_{j=0}^{T-t} \mathcal{S}_{t+j} : \\ u_t + df_t + (1-\gamma_t)\tilde{\xi}_t - \mathcal{L}_t((1-\gamma_t)\tilde{\xi}_t) \geq \mathrm{dem}_t \\ x_{t+1} \geq x_{t+1}^{crit} \\ x_{t+1} = x_t - u_t + \gamma_t \tilde{\xi}_t \qquad\qquad \text{and, for } j = 1, \dots, T-t : \\ \mathbb{P}\left(u_{t+j} + df_{t+j} + (1-\gamma_{t+j})\xi_{t+j} - \mathcal{L}_{t+j}((1-\gamma_{t+j})\xi_{t+j}) \geq \mathrm{dem}_{t+j} \middle| \tilde{\xi}_{[t]} \right) \geq 1 - \varepsilon_{\mathtt{p}} \\ \mathbb{P}\left(x_{t+j+1} \geq x_{t+j+1}^{crit} \middle| \tilde{\xi}_{[t]} \right) \geq 1 - \varepsilon_{\mathtt{p}} \,, \end{cases}
$$

for a given confidence level $\varepsilon_{\mathtt{p}} \in (0,1)$, possibly varying with t and $j \in \{1, \dots T-t\}$. The probabilistic constraints in (12) are conditioned only to the realization history of streamflows $\tilde{\xi}_{[t]}$ because the demand is assumed to be deterministic for now. Otherwise, the demand history also conditions the calculation of moments, cf. (24) in Section 4.4.

The optimal state x_{t+1} and control u_t are obtained by solving the following chance-constrained linear program:

$$
(13)\qquad \min \sum_{\tau=t}^{T} [cu_\tau^\top u_\tau + cd_\tau^\top df_\tau] \ \text{ subject to (12).}
$$

As explained below, see (22) in Section 4.3, the risk-averse problem is deterministic, hence the use of a deterministic objective function in (13).

If $(x_{t+1}^*, u_t^*, df_t^*, \dots, u_T^*, df_T^*)$ is an optimal solution to (13), then decisions $(x_{t+1}^*, u_t^*, df_t^*)$ are taken at time step t and satisfy the constraints (3), (4), (5). As a result, the policy formed by all these sub-vectors, obtained after solving (13) for $t = 1, \dots, T$, is *feasible*. Moreover, for each time step, controls (u_t^*, df_t^*) depend on the reservoir levels at the beginning of this time step and on the history of streamflows $\tilde{\xi}_1, \dots, \tilde{\xi}_t$, but not on future scenarios of streamflows. Consequently, our policy is *time consistent* in the sense of Shapiro (2009). Notwithstanding, the family of optimization problems (13) for $t = 1, \dots, T$ is not time consistent in the sense of Carpentier et al. (2010).

4.2. **On the satisfaction of relatively complete recourse.** Under Assumptions (A1) and (A2) below, relatively complete recourse holds for the OOP:

(A1) along time steps the critical levels are non increasing: $x_1 \geq x_2^{crit} \geq x_3^{crit} \geq \dots \geq x_{T+1}^{crit}$.

(A2) $\mathbb{P}(\xi_t \geq 0) = 1$ for every t.

In this context, it is possible to choose a confidence level $\varepsilon_{\mathtt{p}}$ such that the feasible set (12) of problem (13) is nonempty. In order to keep the streamflows modeled by (8) nonnegative, i.e., for Assumption (A2) to hold, sometimes log normal noises are preferred, see Wu et al. (2005), Beaulieu et al. (1995) (log normal noises make chance constraints approximate, instead of explicit and exact). For Gaussian noises, relatively complete recourse is ensured by adding slack variables z_t (penalized in the objective function) for the minzone constraints (5).

In our numerical experience in Section 5 noises are Gaussian and we consider a real-life power mix with 4 reservoirs to be managed over 10 years, corresponding to 120 stages. For this case, we obtained empirically that $\max_{t,m} \mathbb{P}(\xi_t(m) < 0) = 2.7 \times 10^{-75}$ for $t = 1, \ldots, T = 120$ and $m = 1, \ldots, 4$.

4.3. **Making chance constraints explicit.** Since the PAR model has lag one, the history conditioning the chance constraints in (12) is the last realization, $\tilde{\xi}_t$. Indeed, plugging (9) into the rightmost term of (11) results in

$$x_{t+j+1} = x_t - \sum_{i=0}^{j} u_{t+i} + A_{t,j}\xi_t + \sum_{\ell=1}^{j} B_{t,j}^{\ell}\eta_{t+\ell} + \Theta_{t,j}, \; j = 1, \ldots, T-t, \tag{14}$$

an expression affine on the noises $(\eta_{t+1}, \ldots, \eta_{t+j})$. Once more, lag one assumption makes simple the explicit calculation of scalar coefficients $A_{t,j}, \Theta_{t,j}$ and vector $B_{t,j}$, which involves some simple algebraic manipulations:

$$A_{t,j} = \gamma_t + \sum_{i=1}^{j} \gamma_{t+i}\alpha_{t,i}\,, B_{t,j}^{\ell} = \sum_{i=\ell}^{j} \gamma_{t+i}\beta_{t+i}^{t+\ell}\,, \text{ and } \Theta_{t,j} = \sum_{i=1}^{j} \gamma_{t+i}\theta_{t,i}\,.$$

By contrast, for general lags, calculations are not so straightforward, see Lemmas 2.1 and 2.2 in Guigues and Sagastizábal (2010).

In future demand constraints, we can likewise replace ξ_{t+j} by its expression (9), and write all probabilistic constraints from (12) in the abstract form

$$\mathbb{P}(g(y) \geq X) \geq 1 - \varepsilon_{\mathtt{p}} \qquad \text{where } y := \prod_{j=0}^{T-t} \Big(u_{t+j}, df_{t+j} \Big)\,,$$

for different affine scalar functions $g(\cdot)$ and for different continuous random variables X, depending on the noises η_{t+j}.

To show that the feasible set (12) is a polyhedron, and (13) a linear program, we will use the equivalence

$$\begin{array}{lcl} \mathbb{P}(g(y) \geq X) \geq 1 - \varepsilon_{\mathtt{p}}, \quad \Longleftrightarrow \quad g(y) \geq F_X^{\leftarrow}(1 - \varepsilon_{\mathtt{p}}) & = & \mathbb{E}(X) + F^{-1}(1 - \varepsilon_{\mathtt{p}})\sigma(X) \\ & \uparrow & \text{only for Gaussian } X\,. \end{array} \tag{15}$$

In this relation, $\varepsilon_{\mathtt{p}} \in (0, 1)$, g is a deterministic function, and X is a random variable, keeping in mind that the rightmost identity holds only for Gaussian variables: $X \sim \mathcal{N}(\mathbb{E}(X), \sigma^2(X))$.

We now discuss the computation of the generalized inverse $F_X^{\leftarrow}(1 - \varepsilon_{\mathtt{p}})$ for the different random variables X involved in the chance constraints of (12).

4.3.1. *Critical minimal volume constraint.* Using the expression (14) for the last set of constraints in (12) (chance constraints on the reservoir levels), we obtain (15) for $g(y) := -\sum_{i=0}^{j} u_{t+i}$ and X the random variable

$$X_{t,j} = x^{crit}_{t+j+1} - x_t - A_{t,j}\tilde{\xi}_t - \sum_{\ell=1}^{j} B^{\ell}_{t,j}\eta_{t+\ell} - \Theta_{t,j},\ j = 1, \dots, T-t. \tag{16}$$

This is a Gaussian random variable, with mean and variance given by

$$\mathbb{E}(X_{t,j}) = x^{crit}_{t+j+1} - x_t - A_{t,j}\tilde{\xi}_t - \Theta_{t,j}, \tag{17}$$

$$\sigma^2(X_{t,j}) = \sum_{\ell=1}^{j} \left(B^{\ell}_{t,j}\sigma^{\eta}_{t+\ell}\right)^2. \tag{18}$$

As a result, the chance constraints on the reservoir levels have the form

$$x_t + A_{t,j}\tilde{\xi}_t + \Theta_{t,j} - \sum_{i=0}^{j} u_{t+i} \geq x^{crit}_{t+j+1} + F^{-1}(1-\varepsilon_{\mathtt{p}})\sqrt{\sum_{\ell=1}^{j} \left(B^{\ell}_{t,j}\sigma^{\eta}_{t+\ell}\right)^2} \tag{19}$$

for $j = 1, \dots, T-t$. These constraints are affine in the generation variables $(u_t, \dots, u_T)$.

An interesting feature of our approach is that chance-constraint (19) can be cast back into a mold akin to the original constraint, that is, akin to (5). The rewriting introduces new variables x^R_{t+j}, defined iteratively by transition equations, similar to (3):

$$\left[\begin{array}{lcl} x^R_t & = & x_t \\ x^R_{t+j+1} & = & x^R_{t+j} - u_{t+j} + \gamma_{t+j}\mathbb{E}(\xi_{t+j}|\tilde{\xi}_{[t]})\,, j = 0, \dots, T-t\,. \end{array}\right.$$

Then, it can be easily checked that

$$x^R_{t+j+1} = \mathbb{E}\left(x_{t+j+1}\middle|\tilde{\xi}_{[t]}\right) = x^{crit}_{t+j+1} - \sum_{i=0}^{j} u_{t+i} - \mathbb{E}(X_{t,j})\,,$$

so chance-constraint (19) becomes

$$x^R_{t+j+1} \geq x^{R\ crit}_{t+j+1} := x^{crit}_{t+j+1} + F^{-1}(1-\varepsilon_{\mathtt{p}})\sigma(X_{t,j}), \tag{20}$$

an inequality similar to (5), with an augmented righthand side term that can be explicitly computed using (18).

4.3.2. *Demand constraint.* Using relation (9), chance-constraints for the demand in (12) can be written as (15) with $g(y) := u_{t+j} + df_{t+j}$ and $X = \tilde{\mathcal{L}}_{t+j}(Y_{t,j})$ where, for $j = 1, \dots, T-t$, the random variable $Y_{t,j}$ is defined by

$$Y_{t,j} = (1-\gamma_{t+j})\left(\alpha_{t,j}\tilde{\xi}_t + \sum_{\ell=t+1}^{t+j} \beta^{\ell}_{t+j}\eta_{\ell} + \theta_{t,j}\right). \tag{21}$$

The random variable $Y_{t,j}$ is Gaussian with respective mean and variance

$$\mathbb{E}(Y_{t,j}) = (1-\gamma_{t+j})\left(\alpha_{t,j}\tilde{\xi}_t + \theta_{t,j}\right) \text{ and } \sigma^2(Y_{t,j}) = (1-\gamma_{t+j})^2 \sum_{\ell=t+1}^{t+j} (\beta^{\ell}_{t+j}\sigma^{\eta}_{\ell})^2.$$

In view of (15), we compute $F^{\leftarrow}_{\tilde{\mathcal{L}}_{t+j}(Y_{t,j})}(1-\varepsilon_{\mathrm{p}})$ directly from its definition. For this, we first note that $\tilde{\mathcal{L}}_{t+j}(\cdot)$ is a non increasing function. Indeed, for $x \in [a,b[$, we have $\tilde{\mathcal{L}}'_{t+j}(x) = 2Lx-(aL+1) < \tilde{\mathcal{L}}'_{t+j}(b) = 0$. It follows that $\tilde{\mathcal{L}}_{t+j}(\cdot)$ is strictly decreasing on $]-\infty, b]$ and is constant for $x \geq b$. In this context, the generalized inverse is defined for any $p \geq \tilde{\mathcal{L}}_{t+j}(b)$ by $\tilde{\mathcal{L}}^{\leftarrow}_{t+j}(p) = \inf\{x : \tilde{\mathcal{L}}_{t+j}(x) \leq p\}$. As a result, $\tilde{\mathcal{L}}^{\leftarrow}_{t+j}(p) = \mathrm{dem}_{t+j} - p$ for $p \geq \mathrm{dem}_{t+j} - a$. For $\mathrm{dem}_{t+j} - b + Lb(b-a) \leq p \leq \mathrm{dem}_{t+j} - a$, the generalized inverse is constant and equal to $\frac{1}{2L}\left(aL+1-\sqrt{(aL+1)^2 - 4L(\mathrm{dem}_{t+j}-p)}\right)$, noting that on the leftmost point, because $p = \mathrm{dem}_{t+j} - b + Lb(b-a) = \tilde{\mathcal{L}}_{t+j}(b)$, the identities $\tilde{\mathcal{L}}^{\leftarrow}_{t+j}(p) = b = \frac{aL+1}{2L}$ hold.

Letting $x_{\varepsilon_{\mathrm{p}}}$ denote $F^{\leftarrow}_{\tilde{\mathcal{L}}_{t+j}(Y_{t,j})}(1-\varepsilon_{\mathrm{p}})$ for short, either $\varepsilon_{\mathrm{p}} > \mathbb{P}(Y_{t,j} \leq b)$, and $x_{\varepsilon_{\mathrm{p}}} = \tilde{\mathcal{L}}_{t+j}(b)$; or $0 < \varepsilon_{\mathrm{p}} \leq \mathbb{P}(Y_{t,j} \leq b) = F\left(\frac{b-\mathbb{E}(Y_{t,j})}{\sigma(Y_{t,j})}\right)$. In the latter case,

$$1-\varepsilon_{\mathrm{p}} = \mathbb{P}\left(\tilde{\mathcal{L}}_{t+j}(Y_{t,j}) \leq x_{\varepsilon_{\mathrm{p}}}\right) = \mathbb{P}\left(Y_{t,j} \geq \tilde{\mathcal{L}}^{\leftarrow}_{t+j}(x_{\varepsilon_{\mathrm{p}}})\right) = 1 - F\left(\frac{\tilde{\mathcal{L}}^{\leftarrow}_{t+j}(x_{\varepsilon_{\mathrm{p}}}) - \mathbb{E}(Y_{t,j})}{\sigma(Y_{t,j})}\right),$$

and, hence,

$$x_{\varepsilon_{\mathrm{p}}} = F^{\leftarrow}_{\tilde{\mathcal{L}}_{t+j}(Y_{t,j})}(1-\varepsilon_{\mathrm{p}}) = \tilde{\mathcal{L}}_{t+j}\left(\mathbb{E}(Y_{t,j}) + F^{-1}(\varepsilon_{\mathrm{p}})\sigma(Y_{t,j})\right) \text{ for } 0 < \varepsilon_{\mathrm{p}} \leq F(\frac{b-\mathbb{E}(Y_{t,j})}{\sigma(Y_{t,j})}).$$

When the tailwater level becomes too high, the run-of-river energy is null and the loss function is no longer strictly monotone for $x \geq a$, as in (1). Nevertheless, it is still possible to derive explicit chance constraints proceeding as explained above, as long as functions $\tilde{\mathcal{L}}_t$ remain monotone.

4.3.3. *Explicit representation of* (12). Putting together the previous results, and letting

$$\mathrm{dem}^R_{t+j} := \mathrm{dem}_{t+j} + F^{-1}(1-\varepsilon_{\mathrm{p}})\sigma(Y_{t,j}),$$

by (1) and the fact $-F^{-1}(\varepsilon_{\mathrm{p}}) = F^{-1}(1-\varepsilon_{\mathrm{p}})$, we see that for $0 < \varepsilon_{\mathrm{p}} \leq \mathbb{P}(Y_{t,j} \leq b) = F\left(\frac{b-\mathbb{E}(Y_{t,j})}{\sigma(Y_{t,j})}\right)$, the feasible set (12) has the representation

$$\left\{\begin{array}{l}
(u_t, df_t, \ldots, u_T, df_T) \in \prod_{j=0}^{T-t} \mathcal{S}_{t+j} : \\
u_t + df_t + (1-\gamma_t)\tilde{\xi}_t - \mathcal{L}_t((1-\gamma_t)\tilde{\xi}_t) \geq \mathrm{dem}_t \\
x^R_{t+1} + z_t \geq x^{crit}_{t+1} \\
z_t \geq 0 \\
x^R_{t+1} = x^R_t - u_t + \gamma_t\tilde{\xi}_t \text{ and, for } j = 1, \ldots, T-t: \\
x^R_{t+j+1} = x^R_{t+j} - u_{t+j} + \gamma_{t+j}\mathbb{E}(\xi_{t+j}|\tilde{\xi}_{[t]}), \\
u_{t+j} + df_{t+j} + \mathbb{E}(Y_{t,j}) - \mathcal{L}_{t+j}\left(\mathbb{E}(Y_{t,j}) + F^{-1}(\varepsilon_{\mathrm{p}})\sigma(Y_{t,j})\right) \geq \mathrm{dem}^R_{t+j}, \text{ and} \\
x^R_{t+j+1} \geq x^{R\ crit}_{t+j+1},
\end{array}\right\} \tag{22}$$

where $x_t^R = x_t$, the known initial state. As mentioned, the slack variable z_t ensures satisfaction of (5) for any given realization of the streamflows, and it is penalized in the objective function.

4.4. **CVaR constraints and uncertain demand.** It is also possible to use Conditional Value-at-Risk constraints; for example, requiring that

$$-CVaR_{\varepsilon_{\mathtt{p}}}\left(x_{t+j+1}\middle|\tilde{\xi}_{[t]}\right) \geq x^{crit}_{t+j+1} - \varepsilon_{\mathtt{c}}(x^{crit}_{t+j+1}+1)\,, \tag{23}$$

for some confidence level $\varepsilon_{\mathtt{c}} > 0$. Since the CVaR is translation invariant, using once more the expression (14), inequality (23) can be rewritten as

$$x^{crit}_{t+j+1} - \sum_{i=0}^{j} u_{t+i} - CVaR_{\varepsilon_{\mathtt{p}}}(-X_{t,j}) = -CVaR_{\varepsilon_{\mathtt{p}}}(x^{crit}_{t+j+1} - \sum_{i=0}^{j} u_{t+i} - X_{t,j}) \geq \tilde{x}^{crit}_{t+j+1}$$

with $\tilde{x}^{crit}_{t+j+1} = x^{crit}_{t+j+1} - \varepsilon_{\mathtt{c}}(x^{crit}_{t+j+1}+1)$ and where $X_{t,j}$ is defined in (16).

For a Gaussian random variable X for which higher values are preferred, and for the (bijective) function $\varphi : [0,1] \to [0,\frac{1}{2}]$ defined by $\varphi(0) = 0$ and $\varphi(x) = 1 - F\left(\frac{\exp(-(F^{-1}(1-x))^2/2)}{\sqrt{2\pi}x}\right)$ for $x \neq 0$, the equivalences

$$\underline{\mathrm{X}} \leq -CVaR_{\varphi^{-1}(\varepsilon_{\mathtt{p}})}(X) \iff \underline{\mathrm{X}} \leq -VaR_{\varepsilon_{\mathtt{p}}}(X) = -F_X^{-1}(\varepsilon_{\mathtt{p}}) \iff \mathbb{P}(X \geq \underline{\mathrm{X}}) \geq 1 - \varepsilon_{\mathtt{p}}$$

hold. Together with the rightmost relation in (15), and the expressions (17) and (18) for the mean and standard deviation, an equivalent formulation of (23) is obtained by

$$\text{replacing in (19)} \quad \left[\begin{array}{lll} x^{crit}_{t+j+1} & \text{by} & \tilde{x}^{crit}_{t+j+1} = x^{crit}_{t+j+1} - \varepsilon_{\mathtt{c}}(x^{crit}_{t+j+1}+1) \\ F^{-1}(1-\varepsilon_{\mathtt{p}}) & \text{by} & F^{-1}(1-\varphi(\varepsilon_{\mathtt{p}}))\,. \end{array}\right.$$

If the demand is not deterministic but uncertain, in (12) the future demand constraints need to be conditioned to past realizations of both the streamflows and the demand: $(\tilde{\xi}_{[t]}, \widetilde{\mathrm{dem}}_{[t]})$. Suppose, for simplicity, that the demand process is periodic autoregressive with lag one. Then, for $t = 1,\ldots,T$, the normalized variable $\bar{Z}_t = \frac{\mathrm{dem}_t - \bar{\mu}_t}{\bar{\sigma}_t}$ has mean $\bar{\mu}_t = \mathbb{E}(\mathrm{dem}_t)$ and standard deviation $\bar{\sigma}_t = \sigma(\mathrm{dem}_t)$, and there exist non-null coefficients $\bar{\Phi}_{t+1}$ such that $\bar{Z}_{t+1} = \bar{\Phi}_{t+1}\bar{Z}_t + \bar{\eta}_t$ for $t = 1,\ldots,T-1$. The Gaussian random variables $\bar{\eta}_1,\ldots,\bar{\eta}_T$ are independent, have standard deviations $\sigma(\bar{\eta}_t) = \sigma_t^{\bar{\eta}} > 0$ and define a process $\bar{\eta}$ independent of η. Similarly to (9), $\mathrm{dem}_{t+j} = \bar{\alpha}_{t,j}\,\mathrm{dem}_t + \sum_{\ell=t+1}^{t+j} \bar{\beta}^{\ell}_{t+j}\bar{\eta}_\ell + \bar{\theta}_{t,j}$ for $t = 1,\ldots,T-1$, $j = 1,\ldots,T-t$, where the coefficients are computed as in (10).

In the format (15), demand chance-constraints are written with $g(y) := u_{t+j} + df_{t+j}$ and with the random variable X defined by

$$X := \mathrm{dem}_{t+j}|\widetilde{\mathrm{dem}}_t + X_L \text{ where } X_L = \mathcal{L}_{t+j}(Y_{t,j}) - Y_{t,j} \text{ for } Y_{t,j} \text{ from (21).} \tag{24}$$

Note that X_L is independent of the Gaussian variable $\text{dem}_{t+j}|\widetilde{\text{dem}}_t$, whose mean and variance are given by

$$\mathbb{E}(\text{dem}_{t+j}|\widetilde{\text{dem}}_t) = \bar{\alpha}_{t,j}\,\widetilde{\text{dem}}_t + \bar{\theta}_{t,j}, \quad \sigma^2(\text{dem}_{t+j}|\widetilde{\text{dem}}_t) = \sum_{\ell=t+1}^{t+j} (\bar{\beta}_{t+j}^{\ell}\sigma_{\ell}^{\bar{\eta}})^2.$$

By the rightmost relation in (15), the inverse of the distribution function of X needs to be computed. This computation is not explicit, but estimated by dichotomy, because the distribution function is nondecreasing. More precisely, for any $x \in \mathbb{R}$ we have

$$F_X(x) = \frac{1}{\sqrt{2\pi}\sigma(\text{dem}_{t+j}|\widetilde{\text{dem}}_t)} \int_{y=-\infty}^{y=+\infty} F_{X_L}(x-y) \exp\left(-\frac{(y - \mathbb{E}(\text{dem}_{t+j}|\widetilde{\text{dem}}_t))^2}{2\sigma^2(\text{dem}_{t+j}|\widetilde{\text{dem}}_t)}\right) dy \tag{25}$$

where the distribution function of X_L is given by

$$F_{X_L}(x) = \begin{cases} 0 & \text{if } x < \mathcal{L}_{t+j}(b) - b, \\ 1 - F\left(\frac{aL+1-\sqrt{(aL+1)^2+4Lx}-2L\mathbb{E}(Y_{t,j})}{2L\sigma(Y_{t,j})}\right) & \text{if } \mathcal{L}_{t+j}(b) - b \le x \le -a, \\ 1 - F\left(\frac{-x-\mathbb{E}(Y_{t,j})}{\sigma(Y_{t,j})}\right) & \text{if } x \ge -a\,. \end{cases}$$

The integral in (25) can be approximated by a numerical integration method, or by finding $F_X^{-1}(1-\varepsilon_{\mathtt{p}})$ as a root of the function $G_X(x) = F_X(x) - (1-\varepsilon_{\mathtt{p}})$, for example by a Newton-Raphson method.

5. Numerical assessment

Our numerical results are organized in three parts, succinctly described below.

(1) To emphasize the interest of a rolling horizon in a risk-averse setting, we show the superiority of our risk-averse rolling-horizon (`ra-RH`) policy over a robust nonrolling-horizon (`rob-NRH`) policy derived from the first risk-averse problem, that is from (13)-(22), written with $t = 1$. The progressive lack of precision of policy `rob-NRH` as T increases is made more clear by making the comparisons for two different time horizons, namely $T = 4$ and $T = 12$.

(2) For a short time horizon, we compare our `ra-RH` approach with two policies, obtained with a risk-neutral variant of SDDP, implemented both in nonrolling-horizon and rolling-horizon settings.

(3) For an OOP of real size and large T, we compare `ra-RH` with a risk-neutral nonrolling-horizon SDDP policy. The reason for this choice is that a rolling-horizon SDDP is computationally out of reach for large problems.

The implementation for all the testing was done in Matlab, using Mosek's optimization library to solve linear programming problems (`http://www.mathworks.com/products/matlab/` and `http://www.mosek.com`). The runs were done on a Dell PowerEdge 2900 server with 2 CPUs Intel Xeon

E5345 (2.33 GHz, 8M of cache memory, 1333 MHz FSB), running under CentOS release 5, with 48 GB of RAM.

5.1. **Problem data.** All problem instances correspond to an OOP akin (but not identical) to Brazil's power system, from Guigues and Sagastizábal (2009), and considered in different variants for this numerical study.

5.1.1. *Power system.* For the large-scale case, the hydro-thermal power system operates over a horizon of 10 years, discretized in $T = 120$ time steps, from January 2005 to December 2014. For real-life decisions, the choice of such an extended time horizon is explained by the final condition $\mathcal{Q}_{T+1} \equiv 0$, which "pollutes" the output. The desired time horizon is in fact 5 years: doubling these 60 time stages amounts to eliminating the impact of the boundary condition on the output of interest.

Most of the data was made available by CEPEL[1] and corresponds to part of Brazil's power system. With respect to the simplified OOP described in Section 3, hydro and thermal power plants are now spread over 4 different geographical subsystems, which can import/export energy. Each subsystem, South-East (SE), South (S), North-East (NE), and North (N), corresponds to a geographical region; some energy exchanges make use of a fifth, fictitious, node. 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.

The objective function is given by the total thermal operating cost (ranging between R\$ 6.27 per MWh and R\$ 1047.4 per MWh) plus load shedding, set at R\$ 4170.44 per MWh (hydro generation cost is positive but negligible). Unnecessary spillage and exchanges are avoided by introducing penalties and trading costs between subsystems. In order to prioritize the use of run-of-river energy over hydro or thermal power, the unit cost of any energy transfer should be strictly smaller than all unit hydro costs, which are in turn strictly smaller than all thermal unit costs.

5.1.2. *Loss function, streamflows, and demand representation.* Losses resulting from run-of-river plants efficiencies are modeled per reservoir, using functions (1) constant on t, and defined by the parameters $a := (1010, 366, 137, 0)$ and $b := (5760, 6190, 1560, 10^9)$ for $m = 1, \dots, 4$.

Following the lines of Maceira and Damázio (2004), the monthly streamflows in each reservoir are modeled by a periodic autoregressive process with lags between 1 and 5. The calibration uses historical data from 1931 to 2005, with one important modification, relative to reducing the estimated standard deviations to have nonnegative streamflows almost surely. Due to this modification, our results should be interpreted as an illustration of our methodology, rather than reflecting the real behavior of the Brazilian power system.

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

The demand in each subsystem is a Gaussian random variable with mean given by (31055, 8297, 7103, 3367) MWMonth (with the convention 1 MWMonth=$\frac{365.25\times 24}{12}$ MWh= 730.5 MWh). The respective standard deviation is (3105.5, 829.7, 710.3, 336.7) MWMonth.

For the real-life instance considered in Section 5.4, the demand is considered deterministic, and equal to its average value.

5.2. **Risk-averse policies in nonrolling- and rolling-horizon modes.** After solving (13) written at $t = 1$, the first time step, the corresponding optimal generations

$$u^* := (u_1^*, \ldots, u_T^*), \tag{26}$$

can be viewed as robust generations for the OOP over the entire horizon $[1, T]$. Indeed, our constraints can be seen as robust constraints, using uncertainty sets given in Guigues and Sagastizábal (2009) that are confidence areas for the uncertain parameters.

This robust policy is considered in Bertsimas and Thiele (2006) for inventory problems. For our application we also project the policy, to make it feasible. This is necessary because uncertainty sets are merely confidence areas and policy (u^*, df^*) may not satisfy constraints almost surely for time stages greater than one. Therefore, given the hydro-thermal complement (2), the state $x_t^{\texttt{rob-NRH}}$ and the trajectory of realizations $(\tilde{\xi}_t, \widetilde{\mathrm{dem}}_t)$, we compute $\widetilde{\texttt{HTcomp}}_t := \texttt{HTcomp}(\tilde{\xi}_t, \widetilde{\mathrm{dem}}_t)$ and define the projections

$$u_t^{\texttt{rob-NRH}} := \max(0, \min(u_t^*, x_t^{\texttt{rob-NRH}} + \gamma_t\tilde{\xi}_t - x_{t+1}^{crit}, \widetilde{\texttt{HTcomp}}_t)) \quad \text{and} \quad df_t^{\texttt{rob-NRH}} := \widetilde{\texttt{HTcomp}}_t - u_t^{\texttt{rob-NRH}}.$$

The corresponding states follow the reservoir dynamics: $x_{t+1}^{\texttt{rob-NRH}} = x_t^{\texttt{rob-NRH}} - u_t^{\texttt{rob-NRH}} + \gamma_t\tilde{\xi}_t$, conditioned to the trajectory of realizations.

For the comparison, we consider a reduced instance of the OOP problem, over two different time horizons, as described in Table 1 below. The column "$|\tilde{\xi}^{sim}|$" therein refers to the number of different trajectories considered as scenario realizations in the simulation phase.

T	$\lvert\tilde{\xi}^{sim}\rvert$	x_1	x_t^{crit}	x_{T+1}^{crit}	Demand	Loss function
4	40	$0.5x_{\max}$	$0.2x_{\max}$	x_1	uncertain	with
12	40	$0.5x_{\max}$	$0.2x_{\max}$	x_1	uncertain	with

TABLE 1. Data for reduced OOP.

Table 2 reports a comparison of the two policies, denoted by `ra-RH` and `rob-NRH`, respectively, implemented with $\varepsilon_{\texttt{p}} = 0.19$. In the tables, "s.d." stands for standard deviation. For both horizons, the more conservative policy `rob-NRH` has a higher mean cost. For the shorter time horizon `rob-NRH` has a cost similar to the RH policy, with reduced volatility. The assertion that as T increases, policy `rob-NRH` deteriorates is confirmed by the results obtained with $T = 12$: the robust policy becomes more than 3 times more expensive than `ra-RH` in average, with an even higher increase in volatility (the robust standard deviation is almost 20 times higher than `ra-RH`'s).

Output	Mean	s.d.	VaR 90%	VaR 5%	VaR 1%
`ra-RH`, $T = 4$	$2.378{\times}10^8$	$3.396{\times}10^7$	$2.035{\times}10^8$	$2.848{\times}10^8$	$2.848{\times}10^8$
`rob-NRH`, $T = 4$	$2.516{\times}10^8$	$7.357{\times}10^6$	$2.431{\times}10^8$	$2.617{\times}10^8$	$2.650{\times}10^8$
`ra-RH`, $T = 12$	$5.842{\times}10^8$	$3.476{\times}10^7$	$5.506{\times}10^8$	$6.323{\times}10^8$	$6.572{\times}10^8$
`rob-NRH`, $T = 12$	$1.885{\times}10^9$	$6.697{\times}10^8$	$9.390{\times}10^8$	$2.466{\times}10^9$	$3.105{\times}10^9$

TABLE 2. Results for reduced OOP.

Although not frequently employed for power planning, robust non-rolling horizon policies have sometimes been used in multistage stochastic inventory problems; for instance in Bertsimas and Thiele (2006) with $T = 20$. For large T, the future "seen" by `rob-NRH` becomes less and less accurate, and the policy becomes too conservative (cf. our final comments in Section 6).

5.3. **Risk-averse rolling-horizon policy versus risk-neutral SDDP policy in nonrolling- and rolling-horizon modes.** In order to determine the potential benefits in implementing our `ra-RH` policy, we compare it with two SDDP variants, over the small OOP instance in Table 3. The column "$|\tilde{\xi}^{sim}|$" refers both to the number of trajectories considered by `ra-RH` and to the number of scenarios employed by SDDP in the simulation phase. The columns "$|\tilde{\xi}^{fwd}|$" and "$|\tilde{\xi}^{back}|$" stand, respectively, for the number of SDDP forward scenarios and the number of noise realizations at each stage in SDDP backward pass.

T	$\vert\tilde{\xi}^{sim}\vert$	$\vert\tilde{\xi}^{fwd}\vert$	$\vert\tilde{\xi}^{back}\vert$	x_1	x_t^{crit}	x_{T+1}^{crit}	Demand	Loss function
12	20	10	9	$0.5x_{\max}$	$0.2x_{\max}$	x_1	uncertain	with

TABLE 3. Data for small OOP.

Table 4 reports the cost statistics obtained with the different policies. In the table, "`sddp-NRH`" and "`sddp-RH`" stand for the nonrolling- and rolling-horizon variants of SDDP. We observe that using a rolling-horizon approach allows SDDP to reduce its mean cost (in 7%), at the expense of a higher volatility (the s.d. increase is more than 40%). All policies have similar mean costs, with reduced indicators for volatility with `ra-RH` policy: both the cost standard deviation and the number of extreme cost scenarios (VaRs for low values of the confidence level) are reduced. In relative terms, `ra-RH` s.d. is 6% of the average cost, while `sddp-NRH` and `sddp-RH` standard deviations represent 51% and 76% of the respective average costs. As for the cost of the extreme scenarios, the VaR 5% represents 108%, 165% and 191% of the corresponding `ra-RH`, `sddp-NRH`, and `sddp-RH` average cost.

5.4. **Determining the impact of modeling losses for a large-scale problem.** In view of the numerical experience reported so far, for the large-scale OOP in Table 5 that follows, it is

Output	Mean	s.d.	VaR 90%	VaR 5%	VaR 1%
ra-RH, $\varepsilon_{\mathtt{p}} = 0.19$	5.842×10^8	3.476×10^7	5.506×10^8	6.323×10^8	6.572×10^8
sddp-NRH	5.620×10^8	2.841×10^8	1.342×10^7	9.315×10^8	1.011×10^9
sddp-RH	5.260×10^8	3.981×10^8	8.740×10^7	1.006×10^9	1.231×10^9

TABLE 4. Measures of central tendency and of dispersion of the generation cost (R$) with $T = 12$.

not worthy to implement the rob-NRH policy. Likewise, since when the demand is uncertain the cumulative distribution inverse function estimation increases the computational work per iteration, we only consider deterministic demand in this test-case. Finally, it is not possible to implement SDDP in a rolling-horizon mode: to obtain the reported results with sddp-NRH, it has already taken about 3 weeks in our computers.

T	$\lvert\tilde{\xi}^{sim}\rvert$	$\lvert\tilde{\xi}^{fwd}\rvert$	$\lvert\tilde{\xi}^{back}\rvert$	x_1	x_t^{crit}	x_{T+1}^{crit}	Demand	Loss function
120	500	200	20	$x_{\max}$	$0.2x_{\max}$	$0.2x_{\max}$	deterministic	with and without

TABLE 5. Data for large OOP.

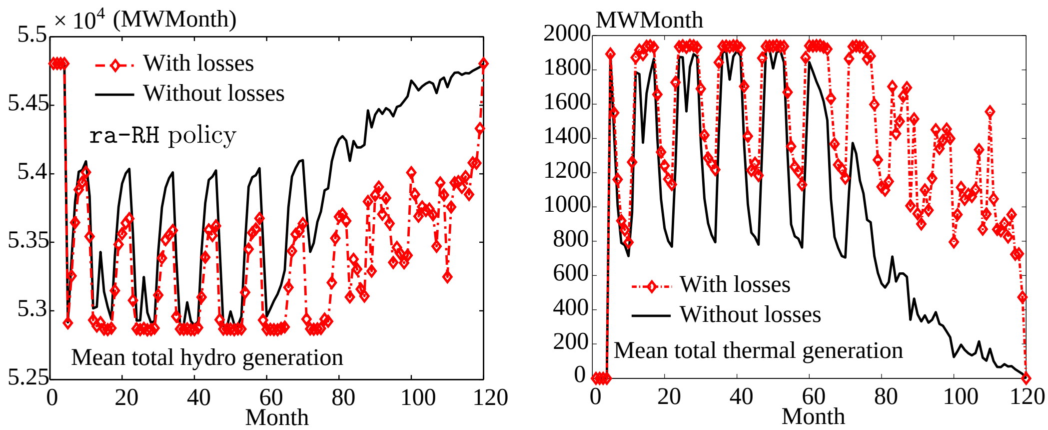

FIGURE 4. ra-RH policies with and without losses representation.

In order to evaluate the effect of modeling losses, we compared two different ra-RH policies, with and without modeling the run-of-river efficiencies. The first policy models the corresponding losses and takes $\varepsilon_{\mathtt{p}} = 0.19$. The second policy is obtained from a model without losses using $\varepsilon_{\mathtt{p}} = 0.15$ and $\varepsilon_{\mathtt{c}} = 0.01$ in (23).

For the two ra-RH policies, Figure 4 shows on the left the equivalent reservoir evolution, including the run-of-river generation, and on the right the average thermal generation. The model with losses keeps water and uses more thermal power. Incidentally, these graphs also put in evidence the "boundary effect" induced by the fact that the last recourse function is null: at the end of the optimization period, since water costs nothing, the optimal decision is to generate only hydro power. Note in addition that the boundary effect decreases as time stages get closer in the future. In fact, the perturbation becomes practically imperceptible for time steps smaller than $t = 70$, thus justifying the heuristic practice of doubling the time steps.

For the problem in Table 5, energy losses can be quite important: taken over all time steps, scenarios, and subsystems, the mean ratio $100\dfrac{\mathcal{L}_t((1-\gamma_t(m))\tilde{\xi}_t(m))}{(1-\gamma_t(m))\tilde{\xi}_t(m)}$ equals 4.8% while the maximal loss goes up to 36% of the streamflow. This shows the importance of modeling the run-of-river efficiencies. Notwithstanding, the loss modeling may significantly increase the generation cost. For all the policies considered, that is for `WS`, `ra-RH`, and `sddp-NRH`, modeling losses doubled the average cost. Also, as observed for the reduced OOP, `ra-RH` reduces volatility with respect to `sddp-NRH`.

A comparison of the different cost distributions can be found in Figure 5. Letting $C_{\mathtt{WS}}$ and $C_{\mathtt{ra-RH}}$, denote, respectively, the generation cost with policies `WS` and `ra-RH`, the figure reports the ratios $\frac{C_{\mathtt{ra-RH}} - C_{\mathtt{WS}}}{C_{\mathtt{WS}}}$ with losses on the left and without losses on the right.

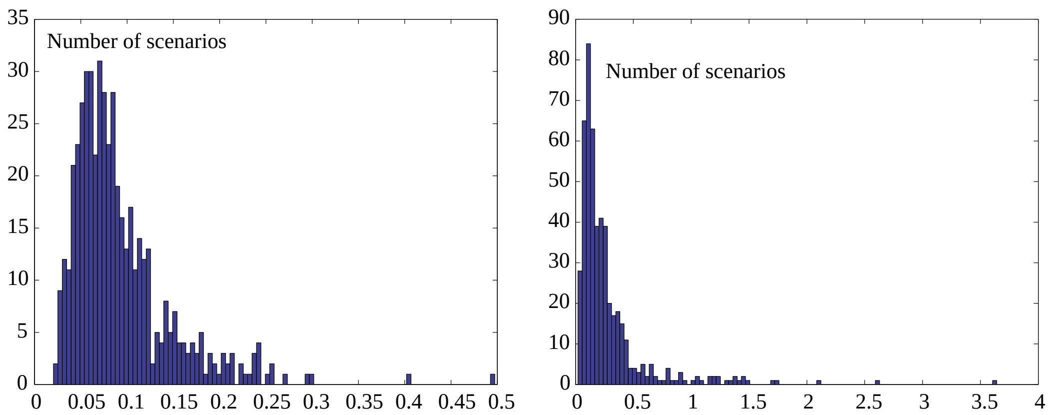

Figure 5. Ratio of empirical cost distributions with and without losses (left and right).

In relative terms, `ra-RH` mean cost is 8.6% or 43% higher than `WS` mean cost, depending on whether or not losses are incorporated in the model. As mentioned, taking into account losses results in a significant increase in the cost for all policies; but, as shown by the rightmost graph in Figure 5, empirical distributions of the cost of `ra-RH` and `WS` become “closer”, a phenomenon that can be explained by the fact that the run-of-river mean generation is less in this case.

Finally, Figure 6 shows the hydro and thermal generation for the three considered policies, taking into account the run-of-river efficiencies and, hence, modeling losses. We see that, for all policies, thermal generation is comparatively much smaller than hydro-generation, as expected in a hydro-dominated system (in our configuration, at each time step, thermal power can cover at most 12.8% of the average demand). When needed, thermal plants are committed in ascending order of their operational cost, to prevent load shedding. Globally, on all scenarios, $\mathtt{sddp} - \mathtt{NRH}$

uses slightly less water. On the first half of the optimization period, $\texttt{sddp} - \texttt{NRH}$ tends to use more water than the other two policies. Once more, the alluded "boundary effect" makes all policies use only hydro-power at the end of the optimization period.

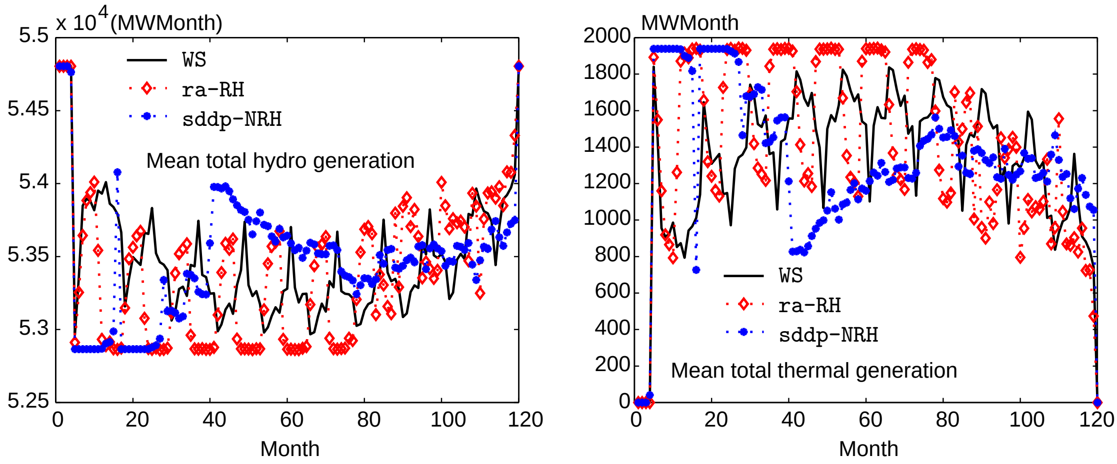

FIGURE 6. `WS`, `ra-RH`, and `sddp-NRH` policies with modeling of losses.

6. Concluding remarks

For large-scale problems, rolling-horizon models are often impractical, because, as shown by Figure 4, they need to generate scenario trees over each future step $(t+1, \dots, T)$, at each time stage t. The corresponding calculations are prohibitive, making non-rolling horizon approaches more popular for stochastic programming problems with many time stages. The situation is only worse when building risk-averse policies, since there are more variables and constraints. We propose an alternative risk-averse policy, implemented in a rolling horizon that solves a sequence of chance-constrained problems. The explicit formulation of chance constraints for each full subtree makes our `ra-RH` policy not only feasible, but also "close" to the optimal one.

In particular, the results presented for the optimal operation problem of a large hydro-thermal power system show that our risk-averse rolling-horizon model is both realistic and tractable. Indeed, the model can efficiently handle energy losses arising when the turbines of run-of-river plants attain their maximum capacity and some streamflow is spilled. It is possible for operators to declare *minzones*, keeping reservoirs above pre-defined limits, with high probability. Similarly, the model can easily incorporate additional probabilistic constraints for flood control purposes, maintaining the reservoirs below a given threshold, if desired. On the basis of our numerical experience and for the case-study analyzed, the consideration of energy losses appears as an important differential factor. We believe such is the case in general for hydro-dominated power systems, especially if operators aim at keeping reservoir levels above critical trajectories.

We also analyzed the numerical behaviour of two commonly used alternative policies, obtained by applying either robust optimization in a nonrolling-horizon mode or stochastic dual dynamic programming, both in nonrolling- and rolling-horizons settings. With respect to robust optimization, for some problems, the objective function value at a robust solution remains close to the optimal value for small data perturbations. Such is the case of a collection of "bad" `NETLIB` problems considered in Ben-Tal and Nemirovski (2000), for which a maximum 0.1% data perturbation yields a robust objective function value distant in at most 1% from the optimal value. However, for a stochastic linear program like ours, with many time stages, streamflow realizations can exhibit a large variation between scenarios. As a result, uncertainty sets, calibrated for $t = 1$ and used to define robust counterparts, become just too large for far ahead time stages, degrading substantially the quality of the robust solution for some scenarios. This unfortunate feature is related to the Robust Optimization premise establishing that all realizations of the uncertain parameters should belong to the uncertainty sets considered in the robust counterpart. But in practice this is not the case, because calibrated uncertainty sets are mere confidence regions. When some scenario not covered by calibration appears in the simulation, the system gets to an *unfeasible* state, highly penalized. In our application, such is the case when there is load shedding: the large deficit cost gives an extremely high cost, and, in particular, produces unacceptable marginal prices.

With respect to the risk-neutral `sddp-NRH` and `sddp-RH` policies, our `ra-RH` approach gives close results in much less computational time. Moreover, our policy can be defined *without resorting to scenario trees*, nor having to use loose termination criteria to stop the iterations.

We finish with an important remark. From the above, one could conclude that a rolling horizon approach should systematically be preferred to other solution methods. In this respect, it is important to keep in mind that our comparisons are both problem and parameter dependent. The rolling-horizon approach has the potential weakness of producing unfeasible risk-averse problems, if confidence levels are too small. Such is not the case in our application; however, for a different type of problem, finding an *a priori* sound choice of confidence levels may be a challenging question, and other methodologies may be more suitable.

References

[1] Andrieu, L., Henrion, R., and Römisch, W., *A model for dynamic chance constraints in hydro power reservoir management*, European Journal of Operational Research **207** (2010), 579–589.

[2] Artzner, P., Delbaen, F., Eber, J., and Heath, D. *Coherent measures of risk*, Mathematical Finance **9** (1999), 203–228.

[3] Beaulieu, N., Abu-Dayya, A., and MacLane, P., *Estimating the distribution of a sum of independent lognormal random variables*, IEEE Transactions on Communications **43** (1995), no. 12, 2869–2873.

[4] Ben-Tal, A., Goryashko, A., Guslitzer, E., and Nemirovski, A., *Adjustable robust counterpart of uncertain linear programs*, Mathematical Programming **99** (2003), 351–376.

[5] Ben-Tal, A., and Nemirovski, A., *Robust convex optimization*, Mathem. of Oper. Res. **23** (1998), no. 4, 769–805.

[6] Ben-Tal, A., and Nemirovski, A., *Robust solutions of linear programming problems contaminated with uncertain data*, Mathematical Programming **88** (2000), 411–424.

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

[8] Birge, J.R., and Louveaux, F., *Introduction to stochastic programming*, Springer Series in Operations Research and Financial Engineering, Springer, 1997, Corr. 2nd printing, 2000, XIX.

[9] Brignol, S., and Rippault, G.,*Risk management applied to weekly generation scheduling*, IEEE Winter Meeting Proceedings **1** (1999), 465–470.

[10] Cabero, J., and Baillo, A., and Cerisola, S., and Ventosa, M., and Garcia-Alcalde, A., and Peran, F., and Relano, G., *A medium-term integrated risk management model for a hydrothermal generation company*, IEEE Trans. on Power Syst. **20** (2005), no. 3, 1379–1388.

[11] P. Carpentier and J-P. Chancelier and G. Cohen and M. De Lara and P. Girardeau, *Dynamic consistency for stochastic optimal control problems*, Available at `http://arxiv.org/PS_cache/arxiv/pdf/1005/1005.3605v1.pdf`, (2010).

[12] Chen, Z.I., and Powell, W.B., *Convergent cutting-plane and partial-sampling algorithm for multistage stochastic linear programs with recourse*, Journal of Optimization Theory and Applications **102** (1999), 497–524.

[13] Conejo, A.J., Nogales, F.J., Arroyo, J.M., and Garcia-Bertrand, R., *Risk-constrained self-scheduling of a thermal power producer*, IEEE Trans. on Power Syst. **19** (2004), no. 3, 1569–1574.

[14] Eichhorn, A., and Römisch, W., *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 (2008), 1–8.

[15] Eichhorn, A., Römisch, W., and Wegner, I. *Polyhedral risk measures in electricity portfolio optimization*, PAMM Proc. Appl. Math. Mech 4 (2004), 7–10.

[16] Guigues, V., *SDDP for some interstage dependent risk averse problems and application to hydro-thermal planning*, Submitted to Computational Optimization and Applications, Available at `http://www.optimization-online.org/DB_HTML/2011/03/2970.html`, (2011).

[17] Guigues, V., Römisch, W., *Sampling-based decomposition methods for risk-averse multistage stochastic programs*, Submitted to SIAM Journal on Optimization, Available at `http://www.optimization-online.org/DB_HTML/2010/10/2763.html`, (2010).

[18] Guigues, V., *Robust production management*, Optimization and Engineering **10** (2009), 505–532.

[19] Guigues, V.,*A stabilized model and an efficient solution method for the yearly optimal power management*, Optimization Methods and Software, **26** (2011), no. 1, 67–88.

[20] Guigues, V., Ndiaye, P., Oustry, F., Romanet, F., and Aïd, R., *Robust Mid-Term power generation management*, Optimization **58** (2009), no. 3, 351–371.

[21] Guigues, V., and Sagastizábal, C., *Risk-averse feasible policies for large-scale multistage stochastic linear programs*, Available at `http://www.optimization-online.org/DB_FILE/2009/02/2215.pdf`, 2009.

[22] Guigues. V., and Sagastizábal, C., *Exploiting structure of autoregressive processes in risk-averse multistage stochastic linear programs*, Tech. report, IMPA, 2010.

[23] Maceira, M.E.P., and Damázio, J., *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 (2004), 397–402.

[24] Mo, B., Gjelsvik, A., and Grundt, A., *Integrated risk management of hydro power scheduling and contract management*, IEEE Trans. on Power Syst. **16** (2001), no. 2, 216–221.

[25] Pereira, M.V.F., and Pinto, L.M.V.G., *Multi-stage stochastic optimization applied to energy planning*, Mathematical Programming **52** (1991), 359–375.

[26] Philpott, A., and Guan., Z., *On the convergence of stochastic dual dynamic programming and related methods*, Oper. Res. Letters **36** (2008), no. 4, 450–455.

[27] Philpott, A., and de Matos, V., *Dynamic sampling algorithms for multi-stage stochastic programs with risk aversion*, submitted to European Journal of Operational Research, Available at `http://www.optimization-online.org/DB_HTML/2010/12/2861.html`, (2010).

[28] Rockafellar. R.T., and Uryasev, S., *Optimization of conditional value-at-risk*, The journal of Risk **2** (2000), no. 3, 21–41.

[29] Shapiro, A., *Analysis of stochastic dual dynamic programming method*, European Journal of Operational Research **209** (2011), 63–72.

[30] Shapiro, A., *On a time consistency concept in risk averse multi-stage stochastic programming*, Operations Research Letters, **37**, (2009), 143–147.

[31] Wallace, S.W., and Fleten, S.E., *Stochastic programming models in energy*, Handbooks in Operations Research and Management Science 10, North Holland, Amsterdam, 2003.

[32] Wets, R. J-B., *Stochastic programming models: wait-and-see versus here-and-now*, IMA Conference on Decision Making under Uncertainty, 2000.

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