

**Report for technical cooperation between Georgia
Institute of Technology and ONS - Operador Nacional
do Sistema Elétrico**

Risk Averse Approach

Alexander Shapiro and Wajdi Tekaya

School of Industrial and Systems Engineering,
Georgia Institute of Technology,
Atlanta, Georgia 30332-0205, USA

May 2011

This is the fourth report of the project.

Contents

1	Introduction	1
2	Case-study description	3
3	Regular Risk Averse Approach	5
3.1	Individual stage costs	6
3.2	Policy value	8
4	Adaptive Risk Averse Approach	12
4.1	Individual stage costs	13
4.2	Policy value	14
5	Comparison of the regular and adaptive approaches	19
6	Conclusions	20
7	Appendix	21
7.1	Regular Risk Averse SDDP	21
7.2	Adaptive Risk Averse SDDP	22

1 Introduction

We will continue to use the terminology of the first phase report. In this report we will deal with **risk averse** approaches to multistage stochastic programming. Let us look again at the formulation of (linear) multistage stochastic programming problems

$$\text{Min}_{\substack{A_1 x_1 = b_1 \\ x_1 \geq 0}} c_1^\top x_1 + \mathbb{E} \left[\min_{\substack{B_2 x_1 + A_2 x_2 = b_2 \\ x_2 \geq 0}} c_2^\top x_2 + \mathbb{E} \left[\cdots + \mathbb{E} \left[\min_{\substack{B_T x_{T-1} + A_T x_T = b_T \\ x_T \geq 0}} c_T^\top x_T \right] \right] \right]. \quad (1)$$

In that formulation the expected value $\mathbb{E} \left[\sum_{t=1}^T c_t^\top x_t \right]$ of the total cost is minimized subject to the feasibility constraints. That is, the total cost is optimized (minimized) *on average*. Since the costs are functions of the random data process, they are random and hence are subject to random perturbations. For a particular realization of the random process these costs could be much bigger than their average (i.e., expectation) values. The first histogram of Figure 8 (on page 9) shows the distribution of the optimal discounted policy value for the considered data set. We will refer to the formulation (1) as *risk neutral* as opposed to *risk averse* approaches which we will discuss below.

The goal of a risk averse approach is to avoid large values of the costs for some possible realizations of the data process. One such approach will be to maintain constraints $c_t^\top x_t \leq \theta_t$, $t = 1, \dots, T$, for chosen upper levels θ_t and *all* possible realizations of the data process. However, trying to enforce these upper limits under any circumstances could be unrealistic and infeasible. One may try to relax these constraints by enforcing them with a high (close to one) probability. However, introducing such so-called *chance constraints* can still result in infeasibility and moreover is very

difficult to handle numerically. So we consider here penalization approaches. That is, at every stage the cost is penalized while exceeding a specified upper limit.

In a simple form this leads to the following risk averse formulation

$$\text{Min}_{\substack{A_1 x_1 = b_1 \\ x_1 \geq 0}} c_1^\top x_1 + \mathbb{E} \left[\min_{\substack{B_2 x_1 + A_2 x_2 = b_2 \\ x_2 \geq 0}} f_2(x_2) + \mathbb{E} \left[\cdots + \mathbb{E} \left[\min_{\substack{B_T x_{T-1} + A_T x_T = b_T \\ x_T \geq 0}} f_T(x_T) \right] \right] \right], \quad (2)$$

where¹ $f_t(x_t) = c_t^\top x_t + \Phi_t [c_t^\top x_t - \theta_t]_+$ with θ_t and $\Phi_t \geq 0$, $t = 2, \dots, T$, being chosen constants. The additional terms $\Phi_t [c_t^\top x_t - \theta_t]_+$ represent the penalty for exceeding the upper limits θ_t . An immediate question is how to choose constants θ_t and Φ_t . In the experiments below we proceeded as follows. First, the risk neutral problem (1) was solved. Then at each stage the 95% quantile of the distribution of the cost $c_t^\top x_t$ of the corresponding optimal policy was estimated by randomly generating $M = 5000$ realizations of the random process and computing respective costs in the forward step procedure. These quantiles were used as the upper limits θ_t in the risk averse problem (2). For the constants Φ_t we used the same value Φ for all stages, this value was gradually increased in the experiments.

The SDDP algorithm with simple modifications can be applied to the problem (2) in a rather straightforward way (see the Appendix). It could be noted that in that approach the upper limits θ_t are fixed and their calculations are based on solving the risk neutral problem which involves all possible realizations of the data process. In other words in formulation (2) the upper limits are not adapted to a current realization of the random process. Let us observe that optimal solutions of problem (2) will be not changed if the penalty term at t -th stage is changed to $\theta_t + \Phi_t [c_t^\top x_t - \theta_t]_+$ by adding the constant θ_t . Now if we adapt the upper limits θ_t to a realization of the data process by taking these upper limits to be $(1 - \alpha_t)$ -quantiles of $c^\top x_t$ *conditional* on observed history $\xi_{[t-1]} = (\xi_1, \dots, \xi_{t-1})$ of the data process, we end up with penalty terms given by AV@R_{α_t} with $\alpha_t = 1/\Phi_t$. Recall that the *Average Value-at-Risk* of a random variable² Z is defined as

$$\text{AV@R}_\alpha[Z] = \text{V@R}_\alpha(Z) + \alpha^{-1} \mathbb{E} [Z - \text{V@R}_\alpha(Z)]_+, \quad (3)$$

with $\text{V@R}_\alpha(Z)$ being the $(1 - \alpha)$ -quantile of the distribution of Z , i.e., $\text{V@R}_\alpha(Z) = F^{-1}(1 - \alpha)$ where $F(\cdot)$ is the cumulative distribution function (cdf) of the random variable Z .

This leads to the following nested risk averse formulation of the corresponding multistage problem (cf., [2])

$$\text{Min}_{\substack{A_1 x_1 = b_1 \\ x_1 \geq 0}} c_1^\top x_1 + \rho_2 \mathbb{E}_{\xi_1} \left[\min_{\substack{B_2 x_1 + A_2 x_2 = b_2 \\ x_2 \geq 0}} c_2^\top x_2 + \cdots + \rho_T \mathbb{E}_{\xi_{[T-1]}} \left[\min_{\substack{B_T x_{T-1} + A_T x_T = b_T \\ x_T \geq 0}} c_T^\top x_T \right] \right]. \quad (4)$$

Here ξ_2, \dots, ξ_T is the random process (formed from the random elements of the data c_t, A_t, B_t, b_t), $\mathbb{E} [Z|\xi_{[t-1]}]$ denotes the conditional expectation of Z given $\xi_{[t-1]}$, $\text{AV@R}_{\alpha_t} [Z|\xi_{[t-1]}]$ is the conditional analogue of $\text{AV@R}_{\alpha_t} [Z]$ given $\xi_{[t-1]}$, and

$$\rho_{t|\xi_{[t-1]}} [Z] = (1 - \lambda_t) \mathbb{E} [Z|\xi_{[t-1]}] + \lambda_t \text{AV@R}_{\alpha_t} [Z|\xi_{[t-1]}], \quad (5)$$

with $\lambda_t \in [0, 1]$ and $\alpha_t \in (0, 1)$ being chosen parameters.

¹By $[a]_+$ we denote the positive part of number a , i.e., $[a]_+ = \max\{0, a\}$.

²In some publications the Average Value-at-Risk is called the Conditional Value-at-Risk and denoted CV@R_α . Since we deal here with conditional AV@R_α , it will be awkward to call it conditional CV@R_α .

In formulation (4) the penalty terms $\alpha^{-1} [c_t^\top x_t - \mathbf{V}@\mathbf{R}_\alpha(c_t^\top x_t)]_+$ are conditional, i.e., adapted to the random process by the optimization procedure. In the following experiments we fix the significance level $\alpha_t = 0.05$ and use the same constant $\lambda_t = \lambda$ for all stages. The constant λ controls a compromise between the average and risk averse components of the optimization procedure. Note that for $\lambda = 0$ problem (4) coincides with the risk neutral problem (1).

It is also possible to give the following interpretation of the risk averse formulation (4). It is clear from the definition (3) that $\mathbf{AV}@\mathbf{R}_\alpha[Z] \geq \mathbf{V}@\mathbf{R}_\alpha(Z)$. Therefore $\rho_{t|\xi_{[t-1]}}[Z] \geq \varrho_{t|\xi_{[t-1]}}[Z]$, where

$$\varrho_{t|\xi_{[t-1]}}[Z] = (1 - \lambda_t)\mathbb{E}[Z|\xi_{[t-1]}] + \lambda_t\mathbf{V}@\mathbf{R}_{\alpha_t}[Z|\xi_{[t-1]}]. \quad (6)$$

If we replace $\rho_{t|\xi_{[t-1]}}[Z]$ in the risk averse formulation (4) by $\varrho_{t|\xi_{[t-1]}}[Z]$, we will be minimizing the weighted average of means and $(1 - \alpha)$ -quantiles, which will be a natural way of dealing with the involved risk. Unfortunately such formulation will lead to a nonconvex and computationally intractable problem. This is one of the main reasons of using $\mathbf{AV}@\mathbf{R}_\alpha$ instead of $\mathbf{V}@\mathbf{R}_\alpha$ in the corresponding risk averse formulation. It is possible to show that in a certain sense $\mathbf{AV}@\mathbf{R}_\alpha(\cdot)$ gives a best possible upper convex bound for $\mathbf{V}@\mathbf{R}_\alpha(\cdot)$.

With a relatively simple additional effort the risk averse problem (4) can be solved by the SDDP algorithm (see the Appendix). We refer to the risk averse formulations (2) and (4) as *regular* and *adaptive*, respectively. It is interesting to note that the adaptive risk averse approach was applied, with a reasonable success, to a study of hydro-thermal scheduling in the New Zealand electricity system in the recent publication [1].

This report is organized as follows. In the next section we provide a description for the used case study. In section 3, we investigate the regular risk averse SDDP. In section 4, we examine the adaptive risk averse SDDP. Finally, in section 5 we compare the two approaches.

2 Case-study description

The numerical experiments described in this report were based on an aggregate representation of the Brazilian Interconnected Power System operation planning problem, with historical data available as of January 2011. The system can be represented by a graph with four generation nodes - comprising sub-systems Southeast (SE), South (S), Northeast (NE) and North (N) - and one (Imperatriz, IM) transshipment node (see Figure 1). The case's general data, such as hydro and thermal plants data and interconnections capacities were taken as static values throughout the planning horizon (120 months).




Figure 1: Case-study interconnected power system

Two different demand cases were considered using the same system configuration: one with a high demand (*high*) and the other with low increasing demand (*lowInc*). In the *high*-case the demand is seasonal, but was made invariant through years. In the *lowInc*-case the demand is also seasonal, but with lower values that increase throughout the study horizon. The two sets of demand values for each system are shown in Figure 2.

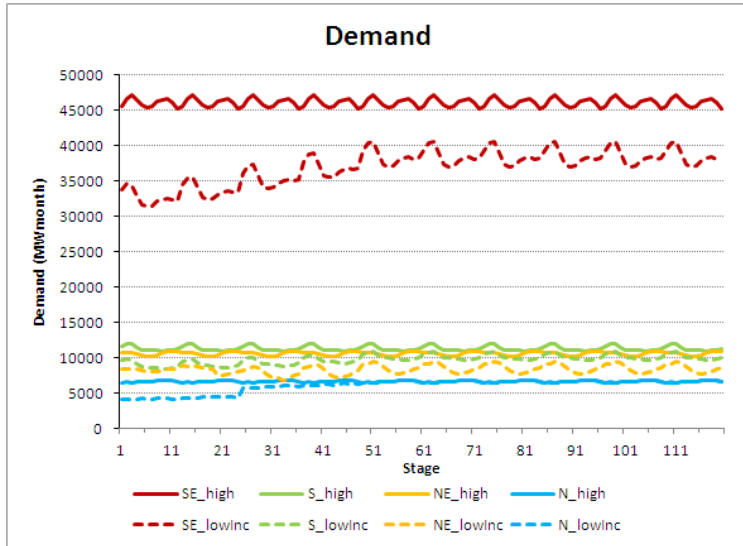


Figure 2: Demand values for each system and dataset

In each system the hydro generators are represented by one equivalent energy reservoir, and the thermal generators are considered individually. The number of thermal plants at each system is: 43 in the Southeast, 17 in the South, 33 in the Northeast and 2 in the North. The load of each area must be supplied by local hydro and thermal plants (see Figure 3) or by power flows among the interconnected areas, with transport capacity shown in Table 1, which may differ depending of the flow direction. The “ ∞ ” symbol in the table means that the energy exchange between the systems is considered unlimited.

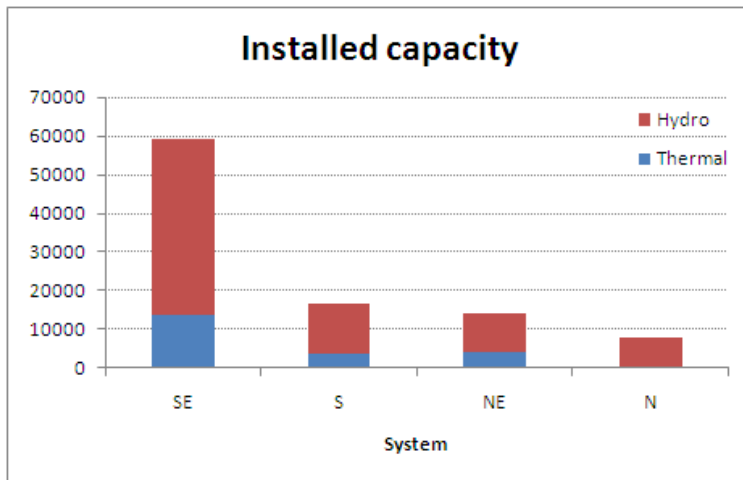


Figure 3: Demand values for each system and dataset

				To		
		SE	S	NE	N	IM
From	SE	–	7379	1000	0	4000
	S	5625	–	0	0	0
	NE	600	0	–	0	2236
	N	0	0	0	–	∞
	IM	3154	0	3951	3053	–

Table 1: Interconnection limits between systems

Four slack thermal generators with high cost accounts for load shortage at each system, with costs shown in Table 2. The capacity of each slack thermal plant is given as per unit value of the demand of the system, and corresponds to the increasing cost of load curtailment.

	Depth	Cost
1	0.05	1142.80
2	0.05	2465.40
3	0.10	5152.46
4	0.80	5845.54

Table 2: Deficit costs and depths

An annual discount rate of 12% was used in the current experiments. This specific value of the discount rate is used in the Brazilian system and it is approved by the national regulator.

A scenario tree consisting of $1 \times 200 \times 20 \times 20 \times \dots \times 20$ scenarios, for 120 stages, was sampled based on a simplified statistical model provided by ONS. In this (seasonal) model, a 3-parameter Log-normal distribution is fitted to each month and for every system. The scenario tree is generated by sampling from the obtained distributions using the Latin Hypercube Sampling scheme. The input data for the simplified statistical model is based on 79 historical observations of the natural monthly energy inflow (from year 1931 to 2009) for each of the four systems.

The sampling of the forward step in the SDDP is different (independently generated) for each experiment. Although the policy value was computed using the discount rate, in the graphs of individual stage cost we just plotted the stage cost without discounting (i.e., $c_t^\top x_t$). Finally, IBM ILOG CPLEX 12.2 was used as LP solver for all these experiments.

3 Regular Risk Averse Approach

We perform the following experiment:

1. Risk neutral SDDP

- run the SDDP for the risk neutral case for 2000 iterations (1 cut per iteration). We save the obtained cuts.
- evaluate the individual cost ($c_t^\top x_t$) at each stage over a sample of 5000 scenarios.

2. Regular Risk Averse SDDP

- run the regular risk averse SDDP (with θ_t being 95% quantile of risk neutral distribution of $c_t^\top x_t$ and $\Phi \in \{25, 50, 75, 100, 200, 300, \dots, 3000\}$) for 2000 iterations (1 cut per iteration). We save the obtained cuts.

- evaluate the individual cost $(c_t^\top x_t)$ at each stage over a sample of 5000 scenarios.

Throughout this section we consider reference penalty values $\Phi \in \{100, 2300\}$ for high-case and $\Phi = 900$ for lowInc-case. These choices will be justified at the end of section 3.2.

The means, quantiles and maximum values of constructed policies were estimated based on $M = 5000$ independently generated scenarios. It could be noted that as such these values are also subject to small variabilities, especially the estimated maximum values could be sensitive to the generated sample of scenarios.

3.1 Individual stage costs

Figure 4 illustrates the mean value obtained at each stage for the risk neutral and the regular risk averse case for some values of the penalty Φ .




Figure 4: Mean of the individual stage costs

Recall that in the high-case we have high demand throughout the stages. That is, the system is under higher load and expensive costs are more likely to occur. This can be seen in the risk neutral case where we have peaks of high costs in the dry season happening periodically. Notice that in the first stages the peaks are not as high as at later stages. This is somehow expected since we start with high stored volumes that allows covering the first shortages. For the final stages, the sudden decrease in cost happens because of the absence of future costs. This decrease is more evident in the risk averse case than in the risk neutral one. A possible explanation could be that, while approaching final stages, in the risk averse case we have higher reservoir levels than in the risk neutral case. This implies cheaper operation costs at the final stages.

It can be seen that in the risk averse case we have higher values of the mean in the first 100 stages. This observation is more noticeable in the first stages and can be justified by the used discounting. Indeed, in the first stages high values have more impact (lower discounting factor) on the present value of costs thus it is more important to be protected against them. However, the discount factor is present in both cases. A good question would be: why it has more impact on the risk averse approach? This can be justified by the fact that we have different objective functions. In the risk averse case, extreme values are penalized and, because of the discounting, they will have more impact if they occur in early stages. Also we can observe that higher penalty values Φ give higher average policy values. This is expected since a protection against high costs comes with an increase of policy value on average.

In the lowInc-case we have a lower demand that increases when progressing through the stages. In the risk neutral case, we can see that the cost for the first 50 stages is similar to the ones we

obtain for $\Phi = 900$. This is explained by the low demand and the capacity of the system to cover the requirement without using more expensive resources. When we progress further and the demand starts to increase, we observe peaks of high costs occurring in dry seasons when the system enters higher load regime. In the risk averse case a remarkable increase in the mean is observed when the system enters the higher load regime (i.e., when the demand starts to increase). An increase of the demand implies that shortage occurrence is more likely, and consequently, higher costs taking place and protection against them is assured by the risk averse approach.

In both cases we can see that there is a lower mean value of the *individual* costs in the final stages for the risk averse approach. It may happen that the *individual* stage costs in risk averse approaches can be lower, while the *sum* of all costs is always bigger. Saving energy at the first stages results in higher costs at the first stages and lower costs at the last stages.




Figure 5: 95% quantile of the individual stage costs

Figure 5 shows the 95% quantile for the risk neutral and regular risk averse case. The main difference between the risk averse and risk neutral cases with regard to the 95% quantile is in the last 10 stages. For almost all the previous stages, the quantile is almost the same (except for several stages) for both of the considered case studies. This behavior is somehow expected since in the regular risk averse approach for these experiments the 95% quantile of the risk neutral case defines a “static” threshold for the penalty to occur. This constitutes the main difference with the adaptive risk averse approach discussed in section 4 where this threshold is embedded in the optimization process and is “adaptive”.




Figure 6: 99% quantile of the individual stage costs

Figure 6 shows the 99% quantile for the risk neutral and regular risk averse cases. We can observe the significant impact of the risk averse approach on the 99% quantile by reducing its value whenever the system is under higher load (i.e., throughout all the stages in (high) and in the last stages in (lowInc)). However, interestingly in the high-case increasing the penalty Φ does not ensure lower 99% quantile value. This can happen since the reduction may occur in other higher quantiles (for instance the maximum see Figure 7).




Figure 7: Maximum of the individual stage costs

Figure 7 shows the maximum individual stage cost for the risk neutral and regular risk averse cases. In both cases we can see the contribution of the risk averse approach by reducing the maximum policy value compared to the risk neutral approach. In the high-case this reduction is more noticeable with higher penalty (i.e., $\Phi = 2300$) and it is approximately spread throughout most of the stages. In the lowInc-case this reduction is mainly perceivable in the final stages when the system enters the higher load regime.

3.2 Policy value

In this section, we compare the obtained policy values for each case. First, we start by examining the histograms of the discounted policy value of 120 stages for some penalty values for each of the case studies. Second, we plot the evolution of the discounted policy value for each of the cases as function of Φ .

The histograms for the risk neutral and regular risk averse approaches are shown in Figure 8 for the considered case studies. Figure 9 is just a zoom in of the histograms to show the extreme values for the distributions.

We can observe the effect of the risk averse approach on the distribution of the discounted policy value: the overall distribution is “pushed” to higher values compared to the risk neutral approach (see Figure 8) and the extreme values are reduced (see Figure 9).

Figure 10 shows evolution of the mean of the discounted policy value for the regular risk averse SDDP approach as function of Φ for the considered case studies. We can see that as the penalty Φ , for the values above the 95% risk neutral quantile, increases the mean of the policy value increases. This behavior is expected and it is the “price of risk aversion”. Note that in the high-case we have a relatively stable increase. However, we can see that in the lowInc-case some variability at higher values of penalty. It is not clear why this behavior is happening. A possible explanation might be related to the static threshold of the penalization defined by the risk neutral quantile.

Figure 11 shows the evolution of the 95% quantile of the discounted policy value for regular risk averse SDDP as function of Φ . First, notice the similar variability pattern of the 95% quantiles for




Figure 8: Histograms of discounted policy value for regular risk averse SDDP




Figure 9: Zoom on the histograms of discounted policy value for regular risk averse SDDP




Figure 10: Mean of discounted policy value for regular risk averse SDDP

both of the considered case studies. In the high-case only the attempted values of $\Phi \leq 200$ ensure lower quantile value than the risk neutral approach (i.e., when $\Phi = 0$). In the lowInc-case only $\Phi = 25$ resulted in lower quantile value for this quantile. Remember that the high-case has bigger demand throughout all the stages, which implies that the system is under higher load and shortage are most likely to occur. This observation justifies the fact that we had more penalty values in the high-case that achieve lower quantile. This observation will be clearer in higher quantile value (see




Figure 11: 95% quantile of discounted policy value for regular risk averse SDDP

Figure 12).




Figure 12: 99% quantile of discounted policy value for regular risk averse SDDP

Figures 12 and 13 show the evolution of the 99% quantile and maximum, respectively, of discounted policy value for regular risk averse SDDP as function of Φ . Figure 12 shows clearly the contribution of the risk averse approach. In the high-case the 99% quantile values are lower than the risk neutral quantile (i.e. $\Phi = 0$) for all the attempted penalty values. In the lowInc-case for almost all the attempted penalty values (except for $\Phi = 2200, 2500$) lower 99% quantiles value are obtained compared to the risk neutral case. Having higher demand in the high-case shows better the contribution of the risk averse approach since in that setting high costs are more likely to occur. A similar observation can be made for the behavior of maximum of the discounted policy value in terms of risk averse contribution: in the high-case we have more occurrences of lower maximum values than in the lowInc-case.

Figure 14 shows the evolution of relative loss in mean (in %), relative reduction of 95% quantile, 99% quantile and maximum (in %) of the discounted policy value with respect to the risk neutral case of (high) as function of Φ .

The idea behind using the risk averse approach is to avoid high costs to occur. This immunity is achieved at the price of losing in policy average value. Depending on what kind of protection we want to achieve, we decide the penalty parameter Φ . For instance, in the high-case if we want to achieve a protection at all costs against maximum policy values, a choice of $\Phi = 2300$ will ensure the highest reduction among the tried values (equal to -22.9%). In this case, there will be a loss




Figure 13: Maximum of discounted policy value for regular risk averse SDDP

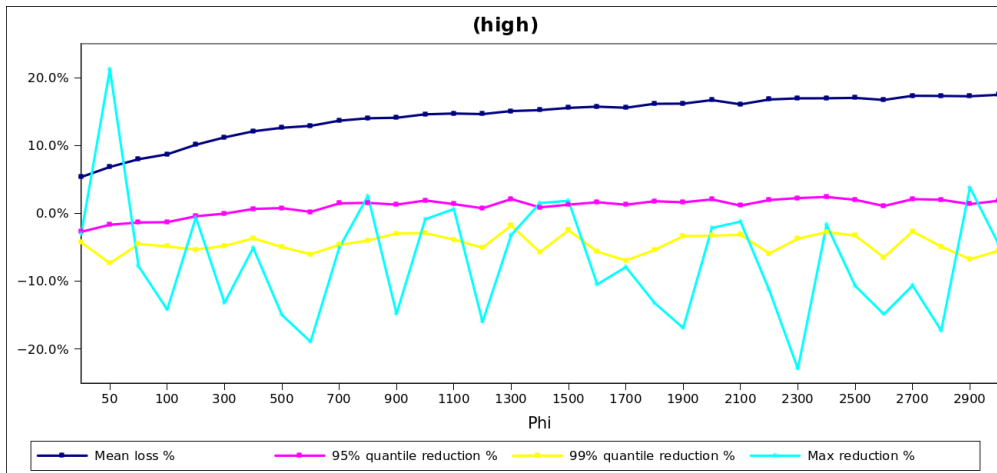


Figure 14: Mean loss and quantiles reduction in % of risk neutral approach of (high)

on average of 16.9%, a very modest reduction in the 99% quantile of -3.7% and an increase of the 95% quantile of 2.3% . However, if we seek more equilibrated protection a value of the penalty of $\Phi = 100$ ensures a reduction of -1.3% , -4.9% and -14.2% in the 95%, 99% quantiles and the maximum discounted policy value compared to the risk neutral case at the “reasonable” loss of 8.7% on average value.

Figure 15 shows the evolution of relative loss in mean (in %), relative reduction of 95% quantile, 99% quantile and maximum (in %) of the discounted policy value with respect to the risk neutral case of the lowInc-case as function of Φ .

First, note the significantly high maximum value for lower penalty ($\Phi = 25, 50$). Also, observe the high variability of the maximum value above the risk neutral case and the moderate improvement of the 99% quantile compared to the previous case as function of the penalty Φ . Recall that in the lowInc-case the shortage might occur only in the later stages with the increase of demand. In other words, the later stages is where we expect occurrence of high costs. This observation justifies the moderate contribution of the penalization process in this situation. A value of $\Phi = 900$ for the penalty ensures a reduction of -4.6% and -12.7% in the 99% quantile and maximum value. However, it leads to an increase of 2.2% and 13.4% in 95% quantile and mean policy value compared to the risk neutral case.

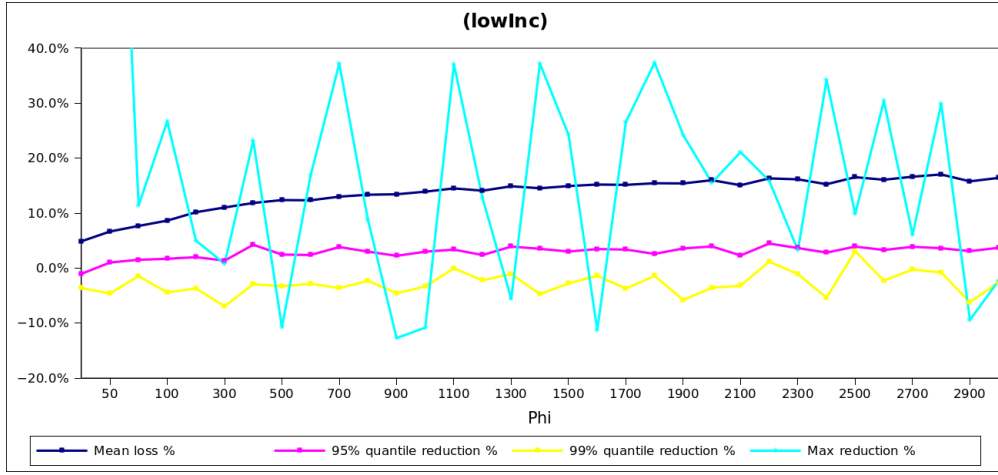


Figure 15: Mean loss and quantiles reduction in % of risk neutral approach of (lowInc)

It could be noted that the reduction of the higher quantiles (the 99% quantiles) and the maximum values doesn't have an increasing trend as function of Φ . At this point it is difficult to give a clear explanation of that phenomena and this behavior could be data dependent. Also recall that the calculated quantiles and maxima are *estimators* based on the generated random sample. As it was pointed earlier, for the higher quantiles and maxima these estimators could be unstable even with $M = 5000$ replications.




Figure 16: Mean, 95%, 99% quantile and maximum of policy value for regular risk averse SDDP

Figure 16 shows the evolution of the mean, 95% quantile, 99% quantile and maximum of the discounted policy value for regular SDDP as function of Φ for the considered case studies. We can observe the variability of the maximum discounted policy value compared to the other quantiles as function of Φ .

4 Adaptive Risk Averse Approach

In this section we present the results for the adaptive risk averse approach. We perform the following experiments:

- run the adaptive risk averse SDDP with $\alpha_t = 0.05$ and $\lambda_t = \lambda \in \{0, 0.01, 0.02, \dots, 0.27,$

0.3, 0.35, \dots, 0.5\} for 2000 iterations (1 cut per iteration). We save the obtained cuts.

- evaluate the individual cost $c_t^\top x_t$ at each stage over a sample of 5000 scenarios.

Throughout this section we consider reference penalty values $\lambda \in \{0.2, 0.3\}$ for high-case and $\lambda \in \{0.1, 0.45\}$ for lowInc-case. These choices will be justified at the end of section 4.2.

4.1 Individual stage costs

Figure 17 illustrates the mean value of individual stage costs for the risk neutral and the risk averse case for different values of λ for the considered case studies.




Figure 17: Mean of the individual stage costs

Recall that in the high-case we have high demand throughout the stages. This implies a higher load on the system and more likelihood for shortage to occur. This can be seen in the risk neutral case where we have peaks in the dry season happening periodically indicating a shortage taking place and using more expensive sources. Notice that at the first stages the peaks are not as high as at later stages. This is somehow expected since we start with high stored volumes that allows covering the first shortages. For the final stages, the significant drop in cost happens because of the absence of future costs. This drop is more important in the risk averse case than in the risk neutral one. A possible explanation could be that in the risk averse case we have higher reservoir levels than in the risk neutral case. This implies cheaper operation costs at the final stages. In the risk averse case we can see that we obtain higher values of the mean for most of the stages. Also we notice that for higher values of λ higher means occur. This is expected since a protection against high costs comes with an increase of policy value on average. Moreover, we observe a significant increase of the average policy value at the first stages. This can be justified by the employed discounting. Indeed, in the first stages high values have more impact (lower discounting factor) on the present value of costs thus it is more important to be protected against them. However, the discount factor is present in both cases. A good question would be: why it has more impact on the risk averse approach? This can be justified by the fact that we have different objective functions. In the risk averse case, extreme values are penalized and, because of the discounting, they will have more impact if they occur in early stages.

In the lowInc-case we have a lower demand which is increasing when progressing through the stages. In the risk neutral case we can see that the cost for the first 50 stages is similar. This is explained by the low demand and the capacity of the system to cover the requirement without using more expensive sources. When we progress further and the demand starts to increase, we

observe peaks of high costs occurring in dry seasons when the system enter higher load regime. In the risk averse case, a remarkable increase in the mean is observed when the system enters the higher load regime (i.e., when the demand starts to increase). An increase of the demand implies that shortage occurrence is more likely, and consequently, higher costs taking place and protection against them is assured by the risk averse approach.

Similar to the regular risk averse approach, in both cases we can see that there is a lower mean value of the *individual* costs in the final stages for the risk averse approach. This can happen since the individual stage costs in risk averse case can be lower. However, the *sum* of all costs is always bigger.




Figure 18: 95% quantile of the individual stage costs

Figure 18 illustrates the 95% quantile individual stage cost for the risk neutral and adaptive risk averse case for different values of λ for the considered case studies. Note that the 95% quantile is reduced most of the time for later stages. In the high-case we observe higher quantile values in the first stages and lower values starting around stage 30 than the risk neutral case. In the lowInc-case there is no significant reduction in the first stages (this is expected since the system is under low demand), an increase in the middle stages and significant reduction in the last stages.

Figure 19 shows the 99% quantile individual stage cost for the risk neutral and adaptive risk averse case for different values of λ for the considered case studies. We can observe the significant impact of the risk averse approach on the 99% quantile by reducing its value whenever the system is under higher load (i.e., throughout all the stages in (high) and in the last stages in (lowInc)). Also, we can see that for higher values of λ (i.e., more importance is given to the quantile minimization than the average) we obtain lower quantile values.

Figure 20 shows the maximum individual stage cost for the risk neutral and adaptive risk averse case for different values of λ for the considered case studies. Similarly for the maximum policy value, we can observe the reduction whenever the system is under high demand. Note also that in the high-case there are some stages (around 20s and 80s) where the maximum value does not change even with changing the values of λ . In the lowInc-case increasing the value of λ allows the reduction of the maximum value most of the time compared to the risk neutral case. This can be justified by noticing that low demand allows more “flexibility” in controlling quantiles by the risk averse approach.

4.2 Policy value

In this section we compare the obtained policy values for each case. First, we start by examining the histograms of the discounted policy value of 120 stages for some penalty values for each of the




Figure 19: 99% quantile of the individual stage costs




Figure 20: Maximum of the individual stage costs

case studies. Second, we plot the evolution of the discounted policy value for each of the cases as function of λ . We recall that this experiment was performed with $\alpha_t = 0.05$.

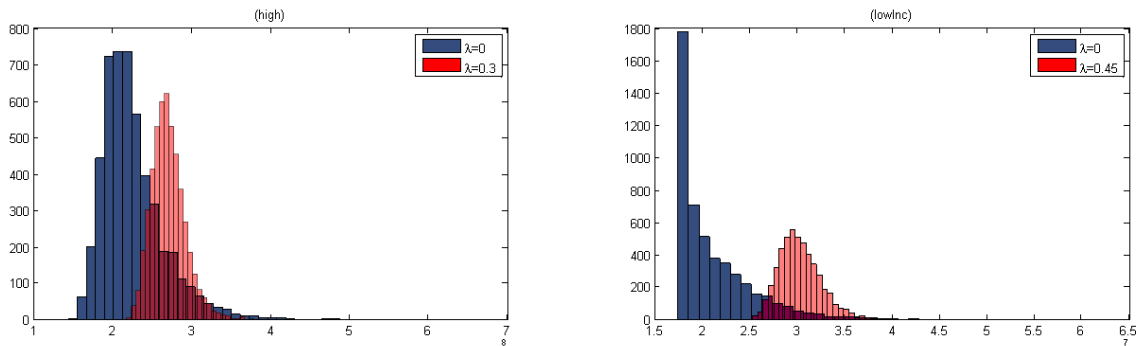


Figure 21: Histograms of discounted policy value for adaptive risk averse SDDP

The histograms for the risk neutral and adaptive risk averse approaches are shown in Figure 21 for the considered cases studies. Figure 22 is just a zoom in of the histograms to show the extreme values for the distributions. We can observe the effect of the risk averse approach on the distribution of the discounted policy value: the overall distribution is “pushed” to higher values

compared to the risk neutral approach (see Figure 21) and the extreme values are reduced (see Figure 22).




Figure 22: Zoom on the histograms of discounted policy value for adaptive risk averse SDDP




Figure 23: Mean of discounted policy value for adaptive SDDP

Figure 23 shows the evolution of mean of discounted policy value for adaptive SDDP as function of λ for the considered case studies. We can see similar behaviors of the mean for the two case studies: increasing values with the increase of λ . This behavior is expected since when increasing the value of λ we are more and more conservative. In other words, with higher values of λ , the algorithm attempts further to reduce the extreme high values to occur. This protection comes with an extra “price” which is an increase of the policy value (estimated by the mean).




Figure 24: 95% quantile of discounted policy value for adaptive SDDP

Figures 24, 25 and 26 show the evolution of 95% quantile, 99% quantile and maximum, respec-

tively, of policy value for adaptive risk averse SDDP as function of λ .




Figure 25: 99% quantile of discounted policy value for adaptive SDDP

In the high-case we obtain lower 95% quantile compared to the risk neutral case (i.e., $\lambda = 0$) for $\lambda \in \{0.04, \dots, 0.26\}$. In the lowInc-case we obtain lower 95% quantile value compared to the risk neutral case for $\lambda \in \{0.03, \dots, 0.2, 0.22, 0.25\}$. We can observe clearly the risk averse impact on especially the 99% quantile for both of the case studies. Indeed, we obtain lower values for this quantile compared to the risk neutral case for most of the values we tried for λ in the lowInc-case (except for $\lambda \leq 0.02$) and almost all the values of λ for (high) (except for $\lambda \leq 0.04$).




Figure 26: Maximum of discounted policy value for adaptive SDDP

In the high-case the maximum of the discounted policy value is reduced for all the tried values of λ except $\lambda \in \{0.02, 0.05, 0.09, 0.18, 0.25, 0.26, 0.5\}$. In the lowInc-case the maximum of the discounted policy value was reduced most for $\lambda = 0.45$. The main observation in this case is that the maximum value was varying around the risk neutral maximum value when the penalty parameter λ increases. Note also that the variation of the quantiles values is not a linear function of λ .

Table 3: Relative percentage of loss with respect to the risk neutral case

Case study/ λ	0.05	0.1	0.15	0.2	0.25	0.3	0.35	0.4	0.45	0.5
(high) in %	1.6	3.2	7.3	10.3	13.8	19.6	22.1	26.8	31.0	35.3
(lowInc) in %	1.9	3.3	7.3	11.5	16.9	23.7	29.0	35.8	42.2	46.7

Table 3 summarizes the relative percentage of loss in the policy mean value with respect to the risk neutral case. As we mentioned before, protection against high costs is assured with a certain loss in policy value on average. In other words, this can be seen as the “price of risk aversion”. A level of acceptable “loss” needs to be defined based on how much protection we need. We can see

that the loss on average is almost linear in the values of λ with a slight tendency to higher slope value with higher λ values.

Tables 4 and 5 summarize the quantile reduction in relative percentage with respect to the risk neutral values for the considered cases studies.

Table 4: (high): Relative percentage of reduction with respect to the risk neutral case in %

Quantile/ λ	0.05	0.1	0.15	0.2	0.25	0.3	0.35	0.4	0.45	0.5
95%	0.9	-4.4	-3.5	-2.4	-1.4	2.0	2.3	5.2	8.2	10.4
99%	0.7	-7.2	-6.5	-9.7	-8.4	-8.5	-9.2	-7.1	-4.0	-4.4
Maximum	-5.2	6.9	-15.1	-18.8	-10.8	-26.2	-7.7	-14.0	-25.8	7.0

Deciding what values of λ are adequate depends essentially on what kind of protection we seek. For example, in the high-case if we want to reduce as much as possible the maximum cost realization $\lambda = 0.3$ might be a reasonable choice incurring a loss on average of 19.6%. $\lambda = 0.15$ and $\lambda = 0.2$ allow more “uniform” protection at all quantiles with moderate loss on average (7.3% and 10.3% respectively). Notice that the contribution of risk aversion is more perceptible in the case study (high) because of the high demand level throughout all the stages. This configuration puts the system under higher load and causes expensive costs to happen.

Table 5: (lowInc): Relative percentage of reduction with respect to the risk neutral case in %

Quantile/ λ	0.05	0.1	0.15	0.2	0.25	0.3	0.35	0.4	0.45	0.5
95%	0.3	-2.5	-1.8	-0.6	1.4	5.1	7.4	10.6	14.7	17.5
99%	-3.9	-8.2	-7.2	-7.3	-5.6	-7.4	-8.2	-3.6	-2.0	-1.0
Maximum	12.7	-4.3	6.8	3.0	70.1	-14.1	70.4	15.0	-23.1	134.6

Similar observations can be said for the lowInc-case: $\lambda = 0.1$ provides an equilibrated reduction of the quantiles with reasonable loss on average policy value of 3.3%. In the case study (lowInc) the demand is lower than in the high-case and it does not put the system under high load. This observation means that there will not be high costs occurring with the risk neutral case and the contribution of risk aversion is not observed as much as in the high-case.




Figure 27: Mean, 95% quantile, 99% quantile and maximum of policy value for adaptive SDDP

Figure 27 shows the combined evolution of mean, 95% quantile, 99% quantile and maximum of the discounted policy value for adaptive SDDP as function of λ for the considered case studies. One observation that can be made at this point is the relative high variability of the maximum policy value as function of λ compared to the other quantiles and average.

5 Comparison of the regular and adaptive approaches

A natural question is which approach performs better - the adaptive or regular? In this section we discuss this point.

Figure 28 shows the average of the discounted policy value as function of the penalty parameter for the adaptive and regular risk averse approaches.




Figure 28: Mean of policy value for adaptive and regular approaches

The key observation at this point is the shape of the nondecreasing average policy value as function of the penalty in both of the considered case studies. In the regular approach we notice a significant increase for the low penalty values and a lower increase for higher penalty values. In the adaptive approach we observe slow increments for small penalty parameter values and higher increments for higher values. This indicates a fair advantage for the adaptive risk averse approach in the sense that the “price of risk aversion” is lower in the considered range of values (i.e., for $0 < \Phi \leq 3000$ and $0 < \lambda \leq .25$).




Figure 29: 95% quantile of policy value for adaptive and regular approaches

Figure 29 shows the 95% quantile of the discounted policy value as function of the penalty parameter for the adaptive and regular risk averse approaches. Adaptive risk averse approach performs better than the regular risk averse approach with respect to the 95% quantile for almost all penalty value parameter. This result is expected since in the regular method penalty starts to occur when the cost exceeds the “static” threshold defined by the 95% quantile of the risk neutral approach. In other words, as long as the cost does not exceed this level there is no penalty. However, in the adaptive approach this threshold is continuously changing in the optimization process and the penalization is defined by the penalty parameter λ .

Figure 30 shows the 99% quantile of the discounted policy value as function of the penalty parameter for the adaptive and regular risk averse approaches. For $\lambda = 0.17$ and $\lambda = 0.1$ better




Figure 30: 99% quantile of policy value for adaptive and regular approaches

reduction is achieved by the adaptive approach for the high-case and the lowInc-case, respectively, within the considered penalty values between the two approaches. Remember (see Figure 28) that the price that we pay for this protection is lower for the adaptive approach for both of these cases.




Figure 31: Maximum of policy value for adaptive and regular approaches

Figure 31 shows the maximum of the discounted policy value as function of the penalty parameter for the adaptive and regular risk averse approaches. The improvement in the maximum policy value is mostly similar between the two approaches. However, the adaptive approach achieves the better reductions for $\lambda > 0.07$ for both of the considered case studies.

6 Conclusions

In this report we investigated the regular and adaptive risk averse approaches and compared their performance. In the regular risk averse approach, discussed in section 3, the contribution of the method was mainly observed in the 99% quantiles with minor effect on the “static” threshold of 95% quantiles and moderate impact on the maximum values. The adaptive risk averse approach, discussed in section 4, showed better impact on the 95% and 99% quantiles and moderate reduction in the maximum value.

The comparison between the two approaches, discussed in section 5, suggests a fair advantage of the adaptive method by ensuring lower “price of risk aversion” (i.e., loss on policy average value) with better protection against high costs as compared to the regular risk averse approach. An intuitive explanation of this is that the adaptive method employs a dynamical embedding of the quantile minimization in the optimization process, while the regular method relies on a “static” predefined threshold of penalization.

7 Appendix

Let $\theta > 0$ denote the monthly discounting factor.

7.1 Regular Risk Averse SDDP

<p>Initialization: $\Omega_t = \{0\}$ for $t = 2, \dots, T + 1$ and choose $\Phi, \theta_t \forall t$</p> <p>Step 1: Forward and Backward recursion Sample M random scenarios from scenario tree For $k = 1, \dots, M$ /*Forward step*/ $[\bar{x}_1^k, \bar{u}_1^k, \gamma_2^k] = \arg \min \quad c_1 x_1 + \Phi u_1 + \theta \gamma_2$ s.t. $A_1 x_1 = b_1$ $u_1 \geq c_1 x_1 - \theta_1$ $[x_1, u_1, \gamma_2] \in \Omega_2, x_1 \geq 0, u_1 \geq 0$</p> <p>For $t=2, \dots, T-1, T$ $[\bar{x}_t^k, \bar{u}_t^k, \gamma_{(t+1)}^k] = \arg \min \quad \tilde{c}_{tk} x_t + \Phi u_t + \theta \gamma_{t+1}$ s.t. $\tilde{A}_{tk} x_t = \tilde{b}_{tk} - \tilde{B}_{tk} \bar{x}_{t-1}^k$ $u_t \geq c_t x_t - \theta_t$ $[x_t, u_t, \gamma_{t+1}] \in \Omega_{t+1}, x_t \geq 0, u_t \geq 0$</p> <p>End For /*Backward step*/ For $t=T, T-1, \dots, 2$ For $j = 1, \dots, N_t$ $\tilde{Q}_{tj}(\bar{x}_{t-1}^k) = \min \quad \tilde{c}_{tj} x_t + \Phi u_t + \theta \gamma_{t+1}$ s.t. $\tilde{A}_{tj} x_t = \tilde{b}_{tj} - \tilde{B}_{tj} \bar{x}_{t-1}^k$ ($\tilde{\pi}_{tj}^k$: dual variable) $u_t \geq c_t x_t - \theta_t$ $[x_t, u_t, \gamma_{t+1}] \in \Omega_{t+1}, x_t \geq 0, u_t \geq 0$</p> <p>End For $\tilde{Q}_t(\bar{x}_{t-1}^k) \leftarrow \frac{1}{N_t} \sum_{j=1}^{N_t} \tilde{Q}_{tj}(\bar{x}_{t-1}^k), \tilde{g}_t^k \leftarrow \frac{1}{N_t} \sum_{j=1}^{N_t} -\tilde{B}_{t,j} \tilde{\pi}_{tj}^k$ $\Omega_t \leftarrow \{[x_{t-1}, u_{t-1}, \gamma_t] \in \Omega_t : [-\tilde{g}_t^k \quad 0 \quad 1] \begin{bmatrix} x_{t-1} \\ u_{t-1} \\ \gamma_t \end{bmatrix} \geq \tilde{Q}_t(\bar{x}_{t-1}^k) - \tilde{g}_t^k \bar{x}_{t-1}^k\}$</p> <p>End For</p>
<p>Step 2: lower bound update $\underline{z} = \min \quad c_1 x_1 + \Phi u_2 + \theta \gamma_2$ s.t. $A_1 x_1 = b_1, u_1 \geq c_1 x_1 - \theta_1$ $[x_1, u_1, \gamma_2] \in \Omega_2, x_1 \geq 0, u_1 \geq 0$</p>
<p>Step 3: Stopping criterion If (Total number of iterations > itermax); STOP!; Otherwise go to Step 1.</p>

Table 6: Regular risk averse SDDP algorithm

7.2 Adaptive Risk Averse SDDP

<p>Initialization: $\Omega_t = \{0\}$ for $t = 2, \dots, T+1$ $\lambda_{T+1} = 0$ and choose $\lambda_t, \alpha_t, \forall t$</p>
<p>Step 1: Forward and Backward recursion Sample M random scenarios from scenario tree For $k = 1, \dots, M$ /*Forward step*/ $[\bar{x}_1^k, \bar{u}_1^k, \gamma_2^k] = \arg \min \quad c_1 x_1 + \theta [\lambda_2 u_1 + \gamma_2]$ s.t. $A_1 x_1 = b_1$ $[x_1, u_1, \gamma_2] \in \Omega_2, x_1 \geq 0$</p> <p>For $t=2, \dots, T-1, T$ $[\bar{x}_t^k, \bar{u}_t^k, \gamma_{(t+1)}^k] = \arg \min \quad \tilde{c}_{tk} x_t + \theta [\lambda_{t+1} u_t + \gamma_{t+1}]$ s.t. $\tilde{A}_{tk} x_t = \tilde{b}_{tk} - \tilde{B}_{tk} \bar{x}_{t-1}^k$ $[x_t, u_t, \gamma_{t+1}] \in \Omega_{t+1}, x_t \geq 0$</p> <p>End For /*Backward step*/ For $t=T, T-1, \dots, 2$ For $j = 1, \dots, N_t$ $\tilde{Q}_{tj}(\bar{x}_{t-1}^k) = \min \quad \tilde{c}_{tj} x_t + \theta [\lambda_{t+1} u_t + \gamma_{t+1}]$ s.t. $\tilde{A}_{tj} x_t = \tilde{b}_{tj} - \tilde{B}_{tj} \bar{x}_{t-1}^k \quad (\tilde{\pi}_{tj}^k : \text{dual variable})$ $[x_t, u_t, \gamma_{t+1}] \in \Omega_{t+1}, x_t \geq 0$</p> <p>End For $\tilde{Q}_t(\bar{x}_{t-1}^k, \bar{u}_{t-1}^k) \leftarrow \frac{1}{N_t} \sum_{j=1}^{N_t} \{(1 - \lambda_t) \tilde{Q}_{tj}(\bar{x}_{t-1}^k) + \lambda_t \alpha_t^{-1} [\tilde{Q}_{tj}(\bar{x}_{t-1}^k) - \bar{u}_{t-1}^k]_+\}$ $\tilde{g}_{tj}^k \leftarrow -\tilde{B}_{t,j} \tilde{\pi}_{tj}^k, S_t^k = \sum_{j=1}^{N_t} (1_{\{\tilde{Q}_{tj}(\bar{x}_{t-1}^k) > \bar{u}_{t-1}^k\}})$ $\tilde{g}_t^k \leftarrow \frac{1}{N_t} [(1 - \lambda_t) \sum_{j=1}^{N_t} \tilde{g}_{tj}^k + \lambda_t \alpha_t^{-1} \sum_{j=1}^{N_t} \tilde{g}_{tj}^k 1_{\{\tilde{Q}_{tj}(\bar{x}_{t-1}^k) > \bar{u}_{t-1}^k\}}, -\lambda_t \alpha_t^{-1} S_t^k]$</p> $\Omega_t \leftarrow \{[x_{t-1}, u_{t-1}, \gamma_t] \in \Omega_t : \begin{bmatrix} -\tilde{g}_t^k & 1 \end{bmatrix} \begin{bmatrix} x_{t-1} \\ u_{t-1} \\ \gamma_t \end{bmatrix} \geq \tilde{Q}_t(\bar{x}_{t-1}^k, \bar{u}_{t-1}^k) - \tilde{g}_t^k \begin{bmatrix} \bar{x}_{t-1}^k \\ \bar{u}_{t-1}^k \end{bmatrix} \}$ <p>End For</p>
<p>Step 2: lower bound update $\underline{z} = \min \quad c_1 x_1 + \theta [\lambda_2 u_1 + \gamma_2]$ s.t. $A_1 x_1 = b_1$ $[x_1, u_1, \gamma_2] \in \Omega_2, x_1 \geq 0$</p>
<p>Step 3: Stopping criterion If (Total number of iterations > itermax); STOP!; Otherwise go to Step 1.</p>

Table 7: Adaptive risk averse SDDP algorithm

Note that $1_{\{A\}} = \begin{cases} 1 & \text{if statement } A \text{ holds} \\ 0 & \text{otherwise} \end{cases}$.

References

- [1] Philpott, A.B. and de Matos, V.L., Dynamic sampling algorithms for multi-stage stochastic programs with risk aversion,
http://www.optimization-online.org/DB_FILE/2010/12/2861.pdf.
- [2] Shapiro, A., Analysis of Stochastic Dual Dynamic Programming Method, *European Journal of Operational Research*, vol. 209, pp. 63-72, 2011.