Optimal Dam Management (Deterministic and Under Uncertainty)

Michel De Lara et Vincent Leclère

October 10, 2017

Contents

1 Problem statement

We consider a dam manager intenting to maximize the intertemporal payoff obtained by selling power produced by water releases.

1.1 Dam dynamics

Consider the *dam dynamics* $s_{t+1} = \texttt{dyn}\,(s_t, u_t, a_t)$, where

$$\texttt{dyn}\,(s, u, a) = \max\{\underline{s}, \min\{\bar{s}, s - u + a\}\} \tag{1}$$

where

- time $t \in [\![t_0, T]\!]$ is discrete (such as days, weeks or months),
- s_t is the *stock level* (in water volume) at the beginning of period $[t, t+1[$, belonging to $\mathbb{S} = [\![\underline{s}, \bar{s}]\!]$, where $\underline{s}$ and $\bar{s}$ are the minimum and maximum volume of water in the dam,
- the variable u_t is the control decided at the beginning of period $[t, t+1[$ belonging to $\mathbb{U} = [\![0, \bar{u}]\!]$ (it can be seen as the period during which the turbine is open, and the effective *water release* is $\min\{s_t + a_t, u_t\}$, necessarily less than the water $s_t + a_t$ in the dam at the moment of turbinating),
- a_t is the *water inflow* (rain, hydrology, etc.) during the period $[t, t+1[$.

1.2 Intertemporal payoff criterion

We consider a problem of *payoff maximization* where turbining one unit of water has unitary *price* p_t. On the period from t_0 to T, the payoffs sum up to

$$\sum_{t=t_0}^{T-1} p_t \min\{s_t + a_t, u_t\} + K\big(s_T\big) \ , \tag{2}$$

where K is the final valorization of the water in the dam.

1.3 Uncertainties and scenarios

Both the inflows a_t and the prices p_t are uncertain variables. We denote by $w_t := (a_t, p_t)$ the couple of *uncertainties* at time t. A *scenario*

$$w(\cdot) := (w_{t_0}, \cdots, w_T) \tag{3}$$

is a sequence of uncertainties, inflows and prices, from initial time t_0 up to the horizon T.

1.4 Generation of trajectories and payoff by a policy

A *policy* $\phi : [\![t_0, T-1]\!] \times \mathbb{S} \to \mathbb{U}$ assigns a control $u = \phi(t, s)$ to any state s of dam stock volume and to any time $t \in [\![t_0, T-1]\!]$.

Given a policy ϕ and a scenario $w(\cdot)$ as in (3), we obtain a volume trajectory $s(\cdot) := \big(s_{t_0}, \dots, s_{T-1}, s_T\big)$ and a control trajectory $u(\cdot) := \big(u_{t_0}, \dots, u_{T-1}\big)$ produced by the "closed-loop" dynamics

$$
\begin{array}{rcl}
s_{t_0} & = & s_0 \\
u_t & = & \phi\big(t, s_t\big) \\
s_{t+1} & = & \mathtt{dyn}\big(s_t, u_t, a_t\big) \ .
\end{array}
\tag{4}
$$

Pluging the trajectories $s(\cdot)$ and $u(\cdot)$ given by (4) in the criterion (2), we obtain the evaluation

$$
\mathsf{Crit}^{\phi}\big(t_0, s_0\big) := \sum_{t=t_0}^{T-1} p_t \min\{s_t + a_t, u_t\} + K\big(s_T\big) \ . \tag{5}
$$

1.5 Numerical data

We consider a weekly management over a year, that is $t_0 = 0$ and $T = 52$, with

$$
s_0 = 30 \text{ hm}^3 \ , \quad \underline{s} = 1 \text{ hm}^3 \ , \quad \bar{s} = 60 \text{ hm}^3 \ , \quad \bar{u} = 10 \text{ hm}^3 \ . \tag{6}
$$

2 Evaluation of a strategy

We begin by generating a scenario of inflows and prices as follows.

Question 1 *Download and execute in Scilab the file* `dam_management2.sce`*, so that a certain number of macros are now available. Generate one scenario by calling the macro* ***PricesInflows*** *with the argument set to* `1` *(corresponding to a single scenario generation).*

Then, we are going to implement the code corresponding to §1.4 under the form of a macro `simulation_det` whose input arguments are a scenario and a policy (itself given under the form of a macro with two inputs and one output).

Question 2 *Complete the Scilab macro* ***simulation_det*** *by implementing a time loop from initial time t_0 up to the horizon T. Within this loop, follow the dynamics* (4) *and use formula* (5) *to compute the payoff. The outputs of this macro will be the gain* (5) *(scalar), and the state and control trajectories (vectors of sizes $T - t_0 + 1$ and $T - t_0$) given by* (4)*.*

This done, we are going to test the above macro `simulation_det` with simple strategies.

Question 3 *Using the macro* ***strat_constant****, compute the payoff, the state and control trajectories attached to the scenario generated in Question 1. Plot the evolution of the stocks levels as a function of time. Design other policies, like, for instance, the constant strategies ($u_t = k$ for $k \in [\![Dmin, Dmax]\!]$) and the myopic strategy consisting in maximizing the instantaneous profit $p_t \min\{s_t + a_t, u_t\}$. Compare the payoffs given by those different strategies.*

3 Optimization in a deterministic setting

In a deterministic setting, we consider that the sequences of prices $p(\cdot)$ and inflows $w(\cdot)$ are known, and we optimize accordingly. The optimization problem we consider is

$$\max_{(s_t,u_t)_{t\in\{t_0,\dots,T\}}} \quad \sum_{t=t_0}^{T-1} p_t \min\{s_t + a_t, u_t\} + K(s_T) \tag{7}$$

$$s_{t_0} = \quad s_0 \tag{8}$$

$$s_{t+1} = \quad \mathtt{dyn}\,(s_t, u_t, a_t)\ . \tag{9}$$

The theoretical Dynamic Programming equation is

$$\begin{cases} V(T,s) = \underbrace{K(s)}_{final\ profit} \quad , \\ V(t,s) = \max_{0\le u\le \bar{u}} \left[\underbrace{p_t \min\{s_t + a_t, u_t\}}_{instantaneous\ profit} + V\big(t+1, \underbrace{\mathtt{dyn}\,(s_t,u_t,a_t)}_{future\ stock\ level}\big) \right] \end{cases} \tag{10}$$

3.1 Zero final value of water

Here, we fix the final value K of water in problem (7) to 0. This means that the water remaining in the dam at time T presents no economic interest for the manager.

Question 4 *Write the theoretical Dynamic Programming equation attached to Problem* (7)*. Then*

- *complete the Scilab macro* ***`optim_det`*** *that compute the Bellman value of this problem,*
- *complete the Scilab macro* ***`simulation_det_Bellman`*** *which constructs the optimal strategy given a Bellman's Value function,*
- *simulate the stock trajectory using the macro* ***`simulation_det`****,*
- *plot the evolution of the water levels, of the prices and of the controls.*

What can you say about the level of water at the end of the period ? Can you explain ?

Question 5 *Theoretically, what other mathematical methods could have been used to solve the dynamic optimization problem* (7)*?*

3.2 Determining the final value of water

We are optimizing the dam on one year. However at the end of this year the dam manager will still have to manage the dam, thus the more water in the dam at the end, the better for the next year. The question is how to determine this value.

The main idea is the following : we want to optimize the management of our dam on a very long time, however we would like to actually solve the problem only on the first year, representing the remainings years by the final value function K in (7). Thus $K(s)$ should represent how much we are going to earn during the remaining time, starting at state s.

Question 6 *Consider the optimal strategy s_N^* obtained when we solve problem* (7) *on N years, with zero final value ($K = 0$). Using the Dynamic Programming Principle find the theoretical function K_N such that the restriction of the strategy s_N^* on the first year is optimal for the one year problem (Problem 7) with final value $K = K_N$.*

Thus, choosing the final value $K = K_N$ means that we take in consideration the gains on N years. We would like to have N going to infinity, however $K_{N+1} - K_N$ is more or less the gain during one year, thus K_N will not converge. In the following question we will find a way of determining a final value converging with N that represents the problem on a long time.

Question 7 *Consider the optimal control problem* (7) *with final value K, and the same problem with final value $K + c$, where c is a constant. What can you say about their optimal strategies ? their optimal values ?*

If K is the value of remaining water, what should be the value of $K(\underline{s})$ (in the sense that how much the future manager of the dam is ready to pay for you to keep the minimum water level in the dam) ?

How do you understand the macro `final_value_det` ? Test it and comment it. Plot the final value obtained as a function of the stock level.

3.3 Introducing a constraint on the water level during summer months

For environemental and touristic reasons the level of water in the dam is constrained. We expect that, during the summer months (week 25 to 40), the level of water in the dam must be above a minimal level $\underline{s}'$.

Question 8 *Recall that a constraint can be integrated in the cost function : whenever the constraint is violated the cost function should be infinite.*

Create a Scilab macro `optim_det_constrained` to integrate this constraint.

Compare the evolution (for different minimal levels $\underline{s}'$) of stock trajectories and optimal values. What can you say about it ? What should you do in order to compute a final value of water adapted to the problem with constraints ?

3.4 Closed-loop vs open-loop control

A closed-loop strategy is a policy given by $\phi : [\![t_0, T-1]\!] \times \mathbb{S} \to \mathbb{U}$, which assigns a water turbined $u = \phi(t, s)$ to any state s of dam stock volume and to any decision period $t \in [\![t_0, T-1]\!]$, whereas an open-loop strategy is a predetermined planning of control, that is a function $\phi : [\![t_0, T-1]\!] \to \mathbb{U}$.

Let us note that, formally, an open-loop strategy is a closed-loop strategy.

Question 9 *In a deterministic setting show that a closed-loop strategy is equivalent to an open-loop strategy in the sense that, for a given initial stock s_0, the stock and control trajectories of* (4) *will be the same.*

Write a Scilab macro that constructs an optimal open-loop strategy from the optimal closed-loop solution.

However, one can make an error in his prediction on inflows or prices and open-loop control may suffer from this. In order to represent this, we will proceed in the following way.

1. We simulate a scenario of prices and inflows.
2. We determine the optimal closed-loop strategy via Dynamic Programming.
3. We determine the associated optimal open-loop strategy.
4. We test both strategies on the original scenario.
5. We modify slightly the original scenario (keep in mind that all inflows must be integers)
6. We test both strategies on the modified scenario.

The "slight" modification of the original scenario must be simple and well understood. Thus we should change either the price or the inflow, at a few times only. However the size of the modification can be substantial.

Question 10 *Write a Scilab macro* `comparison_openVSclosed_loop` *that will implement this procedure and test it. Are there any differencies of value and stock trajectories for the original scenario ? Are there any differencies of value and stock trajectories for the modified scenario ? Why ?*

In the same macro, compute the optimal strategy for the modified scenario and compare the results of the open-loop and closed-loop strategies derived from the original scenario to the optimal result of the modified scenario.

Comment on the pro and cons of closed-loop strategies against open-loop strategies (in a deterministic setting).

4 Optimization in a stochastic setting

In section 3 we have made optimization and simulation on a single scenario. However water inflows and prices are uncertain, and we will now take that into account.

4.1 Probabilistic model on water inputs and expected criterion

We suppose that sequences of uncertainties $\big(a_{t_0}, \dots, a_{T-1}\big)$, $\big(p_{t_0}, \dots, p_{T-1}\big)$ are discrete random variables with known probability distribution. Moreover we will assume that $a(\cdot)$ and $p(\cdot)$ are independent, and that each of them is a sequence of independent random variables.

Notice that the random variables $\big(a_{t_0}, \dots, a_{T-1}\big)$ are independent, but that they are not necessarily identically distributed. This allows us to account for seasonal effects (more rain in autumn and winter).

To each strategy ϕ, we associate the *expected payoff*

$$\mathbb{E}\Big[\mathsf{Crit}^{\phi}\big(t_0, s_0\big)\Big] = \mathbb{E}\Big[\sum_{t=t_0}^{T-1} p_t \min\{s_t + a_t, u_t\} + K\big(s_T\big)\Big] . \tag{11}$$

This expected payoff will be estimated by a Monte-Carlo approach. In order to do that we will use the macros `Price` and `Inflows` that generate a table of random trajectories of the noise, each line being one scenario. The expected payoff of one strategy will be estimated as the empirical mean of the payoff on these scenarios. In order to compare two strategies we have to use the same scenarios for the Monte-Carlo estimation. Thus, we fix a set of simulation scenarios $(\omega_i)_{i\in[\![1,n]\!]}$, where $\omega_i = \{p^i_1, a^i_1, \cdots, p^i_T, a^i_T\}$. and we will always evaluate the criterion $\mathbb{E}\mathsf{Crit}^{\phi}$ as $\frac{1}{N}\sum_{i=1}^{N} \mathsf{Crit}^{\phi}(\omega_i)$.

Consequently the problem is now written as

$$\begin{aligned}
\max_{u_t, s_t} \quad & \mathbb{E}\sum_{t=t_0}^{T-1} p_t \min\{s_t + a_t, u_t\} + K(s_T) && (12)\\
s_{t_0} = \quad & s_0 && (13)\\
s_{t+1} = \quad & \mathtt{dyn}\,(s_t, u_t, a_t) && (14)\\
u_t = \quad & \phi(t, s_t) \ . && (15)
\end{aligned}$$

4.2 Simulation of strategies in a stochastic setting

Here, we will use the macros `simulation` and `simulation_Bellman` that simulate a strategy on each scenario giving a vector of gains, as well as a matrix of stock and control trajectories.

Question 11 *As in Question 1, test the constant strategies and compare the results.*

4.3 Open-loop control of the dam in a probabilistic setting

We have seen that, in the deterministic case (without any errors of prevision), an open-loop strategy is equivalent to a closed-loop strategy. Thus, in a probabilistic setting, one can be tempted to determine an optimal open-loop strategy.

In a first part, we will work on a mean scenario to derive an open-loop strategy.

Question 12 *Complete the macro* `simu_mean_scenario`*, using the macros from the deterministic study, to compute the optimal strategy for the mean scenario.*

In a second part we compute the best open-loop strategy using the function `optim` built-in in Scilab. We choose a set of optimization scenarios $(\omega'_i)_{i\in[\![1,N_{mc}]\!]}$, where $\omega'_i = \{p'^i_1, a'^i_1, \cdots, p'^i_T, a'^i_T\}$. (let us note that this set of scenarios is fixed and that it is different from the set of simulation scenarios). Then we construct a cost function $J(u)$ as

$$J(u) := \frac{1}{N_{mc}} \sum_{i=1}^{N_{mc}} \mathsf{Crit}(u)(\omega'^i)$$

where u is a vector of $T-1$ variables reprensenting the planning of control. Thus we have

$$\mathsf{Crit}(u)(\omega'^i) = \sum_{t=t_0}^{T-1} p^i_t \min\{s^i_t + a'^i_t, u_t\} + K\big(s^i_T\big)$$

with

$$s^i_{t+1} = \mathtt{dyn}\{s^i_t, u_t, a'^i_t\}$$

Question 13 *Use the macro* `best_open_loop` *to obtain the best possible open-loop strategy. Test it and compare to the strategy obtained for the mean scenario. You will consider the simulation of both strategies on the optimization scenarios and on the simulation scenarios.*[1]

4.4 Stochastic Dynamic Programming Equation

4.4.1 Decision-Hazard framework

We will now focus on finding an optimal closed loop solution for problem (12) The dynamic programming equation associated to the problem of *maximizing the expected profits* is

$$\left\{\begin{array}{lcl} V(T,s) & = & \underbrace{K(s)}_{final\ profit} \quad , \\ V(t,s) & = & \max\limits_{0\le u\le \bar{u}} \mathbb{E}\big[\underbrace{p_t \min\{s_t + a_t, u_t\}}_{instantaneous\ profit} + V\big(t+1, \underbrace{\mathtt{dyn}\,(s_t, u_t, a_t)}_{future\ stock\ level}\big)\big] \ , \end{array}\right. \tag{16}$$

Question 14 *Complete the function* `DP`*, that solves the dynamic programming equation (We consider that $K = 0$).*

Then write a macro `simulation_Bellman_DH` *that will simulate the optimal strategy on a set of simulation scenario.*

Plot an histogram of the payoffs and plot an evolution of the stocks level. Compare the gains obtained with this strategy to the open-loop strategy derived from the mean-scenario. You can also compare this strategy to the optimal open-loop strategy.

[1] As the computation can take some time you might want to go to the next section before the computation is done.

4.4.2 Hazard-Decision framework

One may note that, in practice, the dam manager often assume that the weekly inflows and prices are perfectly known. Indeed at the beginning of the week meteorologists and economists can give some predictions. Moreover this problem is only an approximation of the real one, as a dam is managed per hour and not per week, thus the manager has more information than what we assume in a Decision-Hazard setting. Consequently we will now change slightly problem (12) by assuming that, at each time step t we know the price p_t and inflow a_t.

Problem (12) is turned into

$$\max_{u_t, stock_t} \quad \mathbb{E}\left(\sum_{t=t_0}^{T-1} p_t \min\{s_t + a_t, u_t\} + K(s_T)\right) \tag{17}$$

$$s_{t_0} = \quad s_0 \tag{18}$$

$$s_{t+1} = \quad \mathtt{dyn}\,(s_t, u_t, a_t) \tag{19}$$

$$u_t = \phi(s_t, p_t, a_t) \tag{20}$$

Question 15 *Write a macro `DP_HD` that will solve problem* (17) *in a hazard-decision setting. Test it and compare to the solution from the decision-hazard setting (question 14).*

4.5 (Anticipative) upper bound for the payoff

The choice of the probabilistic model of noises (prices and inflows) is quite important. Until now, we have represented the noises as independent variables, and this is not the more precise probabilistic model we could have used. Consequently we might want to estimate the potential gain in using a more precise (but numerically less tractable) probabilistic model. Thus we would like to have an upper bound on our problem. Such an upper bound can be found by doing an anticipative study : for each scenario we compute the best possible gains on this scenario.

Let us stress out that this *will not give a strategy* that can be used. It only gives an upper bound on the possible gain for a set of simulation scenario, a-posteriori.

Question 16 *Write a macro `Simu_anticipative`, that computes for each scenario the upper bound given by the deterministic optimisation. Compare the results obtained by the differents strategies with this upper bound.*