



[Water Resources Research]

Supporting Information for

[Designing with Information Feedbacks: Forecast Informed Reservoir Sizing and Operation]

[F. Bertoni¹, A. Castelletti¹, M. Giuliani¹, P. M. Reed²]

[¹Department of Electronics, Information, and Bioengineering, Politecnico di Milano, Milano, Italy]

[²Department of Civil and Environmental Engineering, Cornell University, Ithaca, NY, USA]

Contents of this file

1. Text S1 to S7
2. Figure S1 to Figure S8
3. Table S1

Introduction

In this study, we generate three sets of infrastructure designs, namely alternative dam sizes with associated candidate operating policies, for the Kariba dam system by solving the following three distinct multi-objective optimization problems: (i) Basic Infrastructure Design (BID), where the operating policies associated to different dam sizes are informed using a basic set of policy inputs, consisting of the reservoir storage and the month of the year; (ii) Perfect Operating Policy (POP), identified with respect to the full, deterministic known trajectory of external drivers (i.e., streamflows) over the entire evaluation horizon for each of the three dam sizes selected within the set of BID, obtaining a sequence of optimal release decisions that an ideal system operator would follow under perfect knowledge on the future; (iii) Informed Infrastructure Design (IID), which differs from the Basic one only in the formulation of the operating policies associated to different dam sizes, which are now dependent upon the selected set of informative forecast lead times determining the reservoir releases. This supplement contains seven sections. The first provides a mathematical formulation of the parameterized reservoir operating policy, whereas the second of the pure management optimization problem to be solved for each of the three dam sizes selected within the set of BID for identifying the corresponding Perfect Operating Policy. The third displays the runtime evolution of the Borg MOEA search to ensure that the algorithm search is at convergence and that the ten random seeds optimization covers the entire diversity and convergence space, finding high-quality solutions. The fourth analyzes the results of the Iterative Input Selection phase performed over both seasonal and inter-annual streamflow forecasts in order to identify the most informative lead times to be included in the optimal infrastructure design phase. The fifth discusses the Kariba system dynamics achieved under different dam sizes, comparing Basic and Informed Infrastructure Designs in order to thoroughly understand the effects of informative forecast lead times on enhancing the infrastructure design. Similarly, the sixth presents the system dynamics achieved under over-estimated forecasts for different dam sizes. The seventh performs a sensitivity analysis of the Informed Infrastructure Designs with respect to different accuracies in all the biased forecasts generated, evaluated in terms of percentage bias (Pbias) performance metric.

Section S1: Gaussian radial basis functions

In this study, the water reservoir operating policy is parameterized according to non-linear approximating networks, and in particular Gaussian radial basis functions (RBFs), and the reservoir release decision u_t is therefore calculated as follows:

$$u_t = \beta + \sum_{i=1}^N w_i \varphi_i(I_t)$$

where N is the number of RBFs $\varphi(\cdot)$, β is the linear parameter associated to the decision variable u_t , and w_i is the non-negative weight of the i -th RBF ($w_i \geq 0, \forall i$). As for the single RBF, it is defined as follows:

$$\varphi_i(I_t) = \exp \left[- \sum_{j=1}^M \frac{[(I_t)_j - c_{j,i}]^2}{b_{j,i}^2} \right]$$

where M is the number of policy inputs I_t , c_i and b_i are the M -dimensional center and radius vectors of the i -th RBF. In particular, the centers must lie within the input bounded space and the radii must be strictly positive (Busoniu et al, 2011). As a result, the parameters vector employed for the parametrization of the operating policy is defined as $\theta = [c_{j,i}, b_{j,i}, w_i, \beta] \in \mathbb{R}^{n_\theta}$ where $i = 1, \dots, N$, $j = 1, \dots, M$, $n_\theta = nu + N(2M+nu)$, and nu = number of policy outputs (i.e., $nu = 1$ as we consider a single reservoir release decision).

Section S2: Pure management optimization problem

As discussed in section 3.2 of the manuscript, the Perfect Operating Policy is identified under perfect foresight assumption by solving the management side of the joint optimization problem 3 in the manuscript with respect to the vector of $n = 2$ management objectives, for a fixed dam size $\bar{\alpha}$ and associated irrigation diversion parameters $\bar{\theta}_{irr}$ selected within the set of Basic Infrastructure Designs. This pure management optimization problem can be formulated as follows:

$$p^* = \arg \min_p J_p^{POP}$$

$$\text{where } J_p^{POP} = |J_p^{hyd}, J_p^{irr}|$$

$$\bar{\alpha}, \bar{\theta}_{irr} \text{ given}$$

where the optimal operating policy p^* is identified with respect to $n = 2$ management objectives (i.e., J_p^{hyd}, J_p^{irr}). This optimization problem can be solved by either a local optimization method (e.g., gradient-based) or a global optimization method (e.g., direct search). Conversely, if the objective function is time-separable, Deterministic Dynamic Programming (DDP) can be used (Bellman, 1957), which is able to provide an almost exact solution much more efficiently than other nonlinear optimization methods. We employ DDP to solve the pure management optimization problem with respect to the full, deterministically known trajectory of streamflows over the entire evaluation horizon H . For

each dam size, we therefore obtain an optimal operating policy p^* , from which we derive a target sequence of optimal release decisions (i.e., target output of the corresponding iterative input selection procedure) that an ideal system operator would follow under deterministic knowledge on the future.

Section S3: Runtime evolution of the Borg MOEA search

Our choice of running only 10 random seed trials was based on a preliminary diagnostic analysis of the Borg MOEA search in terms of hypervolume runtime dynamics (Figure S1). As can be observed, at the beginning of the search process the 10 random seed trials present an hypervolume that varies between 0.4 and 0.5. Such seeds variability quickly decreases as the number of function evaluations increases up to 1 million, when the hypervolume metric negligibly ranges from 0.75 to 0.77 (shaded area extension). In addition, from 750,000 to 1 million function evaluations the hypervolume value increases from 0.754 to 0.758 on average (bold line), meaning that no additional function evaluation is needed and that the best possible approximation of the Pareto front has been identified. This evolution of the search progress suggests that the algorithm eventually reaches convergence with little variability across the seeds, covering the entire diversity and convergence space and finding high-quality solutions.

Section S4: Iterative Input Selection for seasonal and inter-annual forecasts

We identify via Iterative Input Selection (IIS) algorithm the most informative seasonal forecast lead times that explain the sequence of optimal release decisions associated to the three target trade-offs highlighted in Figure 3 of the manuscript for the large L, medium M, and small S dam sizes selected.

First, we perform a regression on a sample dataset consisting of the Kariba storage s_t and month of the year t in order to prevent this basic set of information, highly correlated with the release trajectory to be explained, to overshadow the real contribution of forecast lead times if jointly considered in the information selection phase. Results are summarized in Table S1, where each row corresponds to a different dam size, operated under three alternative target trade-offs, namely an hydropower-prone (H), a compromise (C) and an irrigation-prone (I) operating policy. The last column reports the share of the variance in the reservoir release trajectories obtained under the target trade-offs that is not explained ($1-R^2$) by s_t and t respectively. The unexplained variance $1-R^2$ systematically increases from I to H operational trade-off, regardless of dam size. This has already been noticed in section 4.1 of the manuscript when discussing the maximum space for improvement in Figure 3 that decreases from H to I. Since the maximum space for improvement measures the distance between Basic Infrastructure Design and Perfect Operating Policy, and the former is informed with the basic set of information consisting of s_t and t , basic irrigation-prone policies are already close to the target trade-off and would not benefit from their conditioning upon other informative variables besides storage and time. This is particularly evident for large L dam sizes, where 13% and 37% of the variance in the releases associated to I and H respectively is not explained by reservoir storage and time. The added benefit of including informative forecast lead times in addition to storage and time to explain the sequence of optimal release decisions will be thus more significant for large dam sizes operated under a hydropower-prone rather than an irrigation-prone policy. As for small S dam sizes, $1-R^2$ assumes almost constant values and equal to 0.27 across the three operational trade-offs. However, such values are pretty low, and the advantages of adding information are limited across all the operational trade-offs. This is reflected by the

value of the maximum space for improvement presented in Figure 3 of the manuscript, which assumes its lowest value for S dam sizes whose Basic and Perfect solutions are very close throughout the entire objective space. As expected, medium M dam sizes behave in between small and large reservoirs.

After calculating the model residuals of s_t and t_t , they are employed as dependent variable when running the Iterative Input Selection algorithm with the set of perfect seasonal streamflow forecasts as independent variable. The aim of the IIS algorithm is to select the most informative seasonal forecast lead times and temporal aggregations (i.e., maximum and minimum over 7 months) describing the target optimal release sequence not explained by reservoir storage and time.

Figure S2 shows the results of this information selection phase, reporting the variables that have been selected more frequently by IIS throughout the 50 runs of the algorithm and the associated average cumulated performance in terms of R^2 . This procedure was applied repeatedly to filter the randomness associated to the construction of the extra-tress models employed by the input selection algorithm (Galelli and Castelletti, 2013). Each row corresponds to a specific dam size, each column to an alternative target trade-off. Regardless of dam size and operational trade-off, the most informative variables selected always provide information on the future streamflow extremes (i.e., maximum for flood peaks or minimum for drought periods) rather than on the cumulated water volume entering the reservoir over different future lead times. In particular, hydropower-prone policies are best informed by the maximum future streamflow over 7 months $qm7_t$ for all three dam sizes. This variable allows the dam operator to acquire perfect knowledge on the maximum flood peak that will enter the reservoir in the next 7 months, and act accordingly by lowering the levels in order not to spill and thus maximize hydropower production (please refer to section 4.2 of the manuscript for further insights into system dynamics). The second most informative variable that is selected by the IIS algorithm is the minimum future streamflow over 7 months $qm7_t$. However, this latter contributes to about 3% of the final cumulated performance averaged across the dam sizes and can be thus considered almost negligible. Dually, irrigation-prone operating policies are associated to the minimum future streamflow over 7 months $qm7_t$ as the only most informative variable for all three dam sizes. $qm7_t$ allows the reservoir operator to acquire perfect knowledge on the most severe drought that the system will experience in the next 7 months, and act accordingly by storing sufficient water to satisfy the irrigation demands and thus minimize the irrigation deficit. Not surprisingly, the compromise policy presents a less clear pattern than the other two extreme solutions, since it has to balance two competing objectives into a single operating strategy ensuring both satisfactory hydropower productions and irrigation deficits. Under a compromise trade-off, each dam size may achieve this balance by operating the reservoir differently based on its own active storage capacity. As a consequence, the most informative lead times explaining the different release trajectories obtained under each dam size change accordingly.

Then, we employ the Iterative Input Selection algorithm to identify the most informative lead times out of an additional set of inter-annual streamflow forecast medians computed over 12, 24, 36, 48, and 60 months ahead (i.e., $qmed12_t$, $qmed24_t$, $qmed36_t$, $qmed48_t$, $qmed60_t$). The selected lead times are then used in the Informed Infrastructure Design phase to assess whether inter-annual streamflow forecasts can further benefit particularly large dam sizes. Figure S3 shows the results of this second information selection phase, reporting the lead times that have been selected more frequently by IIS throughout the 50 runs of the algorithm and the associated average cumulated performance in terms of R^2 .

This procedure was applied repeatedly to filter the randomness associated to the construction of the extra-trees models employed by the input selection algorithm (Galelli and Castelletti, 2013). Each row corresponds to a specific dam size (i.e., small S, medium M, large L), each column to an alternative target trade-off (i.e., hydropower-prone H, compromise C, irrigation-prone I).

Section S5: System dynamics under Informed Infrastructure Design

Figure S4 displays the system dynamics for large L and small S dam sizes in terms of level trajectories (panel b and d) associated to the basic (orange), informed (cyan) and perfect (grey) solutions LH and SH highlighted in panels a and c, respectively. As observed in section 4.2, since the maximum future streamflow over 7 months allows the system operator to acquire perfect knowledge on the flood events that will occur in the near future, he/she is able to keep the large and small reservoir levels respectively 2.5 and 0.5 meters higher than the Basic and closer to the Perfect trajectories without spilling. The advantages of adding informative forecast lead times during the infrastructure design phase are less evident for small dam sizes as they have a restricted space of operation discretion due to a rather small active storage capacity compared to both medium and large dam sizes.

Section S6: System dynamics under over-estimated forecasts

Figure S5a displays the performance of the small S dam size in terms of J^{hyd} and J^{irr} achieved under the set of Basic Infrastructures Designs (orange) and Perfect Operating Policies (grey), along with the set of Informed Infrastructure Designs under over-estimated (yellow), under-estimated (green), under-dispersed (purple) and perfect (cyan) seasonal forecasts. In addition, Figure S5b shows the system dynamics in terms of level and release trajectories for the small dam size, when operated under a hydropower-prone operating policy (i.e., SH solutions highlighted in Figure S5a). When comparing Informed (Over-est) and Informed (Perfect) solutions, over-estimated forecasts, which over-estimate the wet season flood peaks, force the reservoir to release more water from January to May in order not to spill, while lowering releases during the dry season (September-December) in order not to excessively lower the reservoir levels and still maintain a satisfactory level of hydropower production. On the contrary, perfect forecasts enable the reservoir to release less during the wet season, saving water for the dry season when irrigation demand is higher. These dynamics allow both Informed Infrastructure Designs to attain the same hydropower production, yet over-estimated forecasts attain a 12% higher irrigation deficit.

Section S7: Sensitivity analysis with respect to forecasts Pbias

We perform a sensitivity analysis of the Informed Infrastructure Designs with respect to different accuracies in all the biased forecasts generated, evaluated in terms of percentage bias (Pbias) performance metric. We test a low, moderate and high Pbias for each of the forecast biases considered, namely over-estimation, under-estimation, and under-dispersion. In particular, we use a +20% (low), +30% (moderate), and +40% (high) Pbias for over-estimated, -20% (low), -30% (moderate), and -40% (high) for under-estimated, and in the end -6% (low), -10% (moderate), and -15% (high) for under-dispersed seasonal forecasts. Figure S6-Figure S8 display the performance of the hypervolume metric associated to a large L (panel a), medium M (panel b), and small S

(panel c) dam size and identified under a basic set of information (orange), together with perfect (cyan), over-estimated (Figure S6), under-estimated (Figure S7), and under-dispersed (Figure S8) seasonal forecasts characterized by different percentage biases (low - green, moderate - yellow, high - red). In all three figures, the grey bar with hypervolume equal to one corresponds to the Perfect Operating Policies, designed under a full, deterministic knowledge of the future. As already discussed in section 4.4, regardless of dam size, the sets of Informed Infrastructure Designs are less sensitive to both under-estimated and under-dispersed forecasts with both a low and moderate Pbias, attaining almost the same hypervolume as for perfect forecasts. However, when associated to a high Pbias, their performance in terms of hypervolume worsens, moving closer to the Basic Infrastructure Design value, especially for large and medium dam sizes. On the contrary, over-estimated forecasts lead to the lowest values of the hypervolume metric for all three dam sizes, whose under-performance becomes even more pronounced under a +40% over-estimation.




Figure S1. Hypervolume runtime dynamics for the 10 random seeds optimization of the Informed Infrastructure Design. The shaded area is bounded by 5th and 95th percentiles of the hypervolume performance value across the multiple random seeds at each 250,000 runtime function evaluation (NFE), whereas the bold line identifies the average performance.

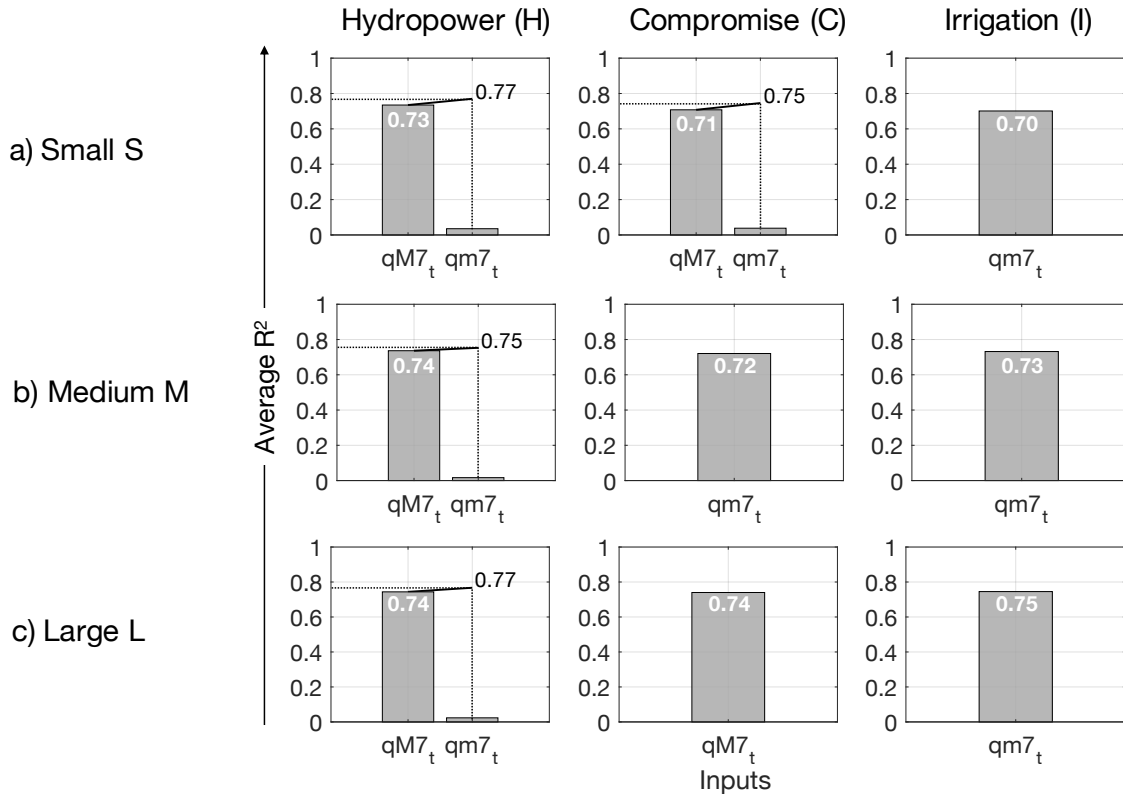


Figure S2. Information selection results obtained by performing 50 runs of the IIS algorithm in terms of average cumulated performance. Each row refers to a single dam

size, each column corresponds to a specific target trade-off of the Perfect operating policies to be explained by the most informative seasonal forecast lead times selected.

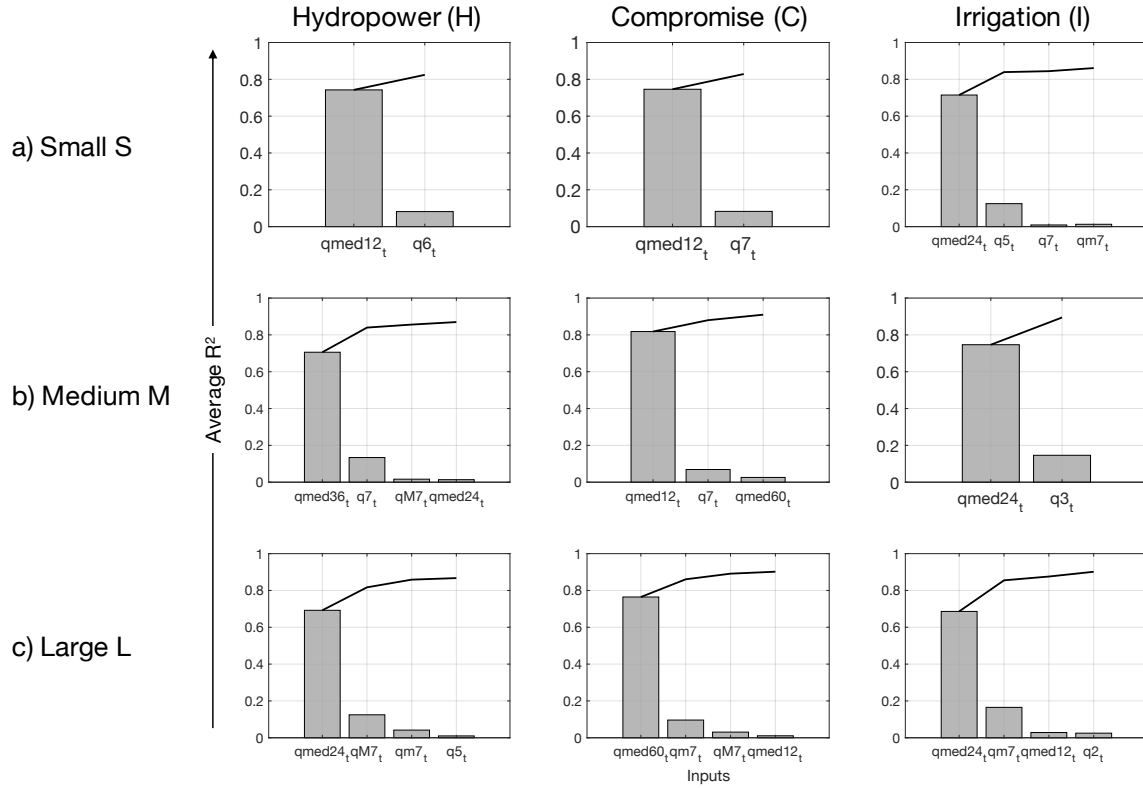


Figure S3. Information selection results obtained by performing 50 runs of the IIS algorithm in terms of average cumulated performance. Each row refers to a single dam size, each column corresponds to a specific target trade-off of the Perfect operating policies to be explained by the most informative inter-annual forecast lead times selected.

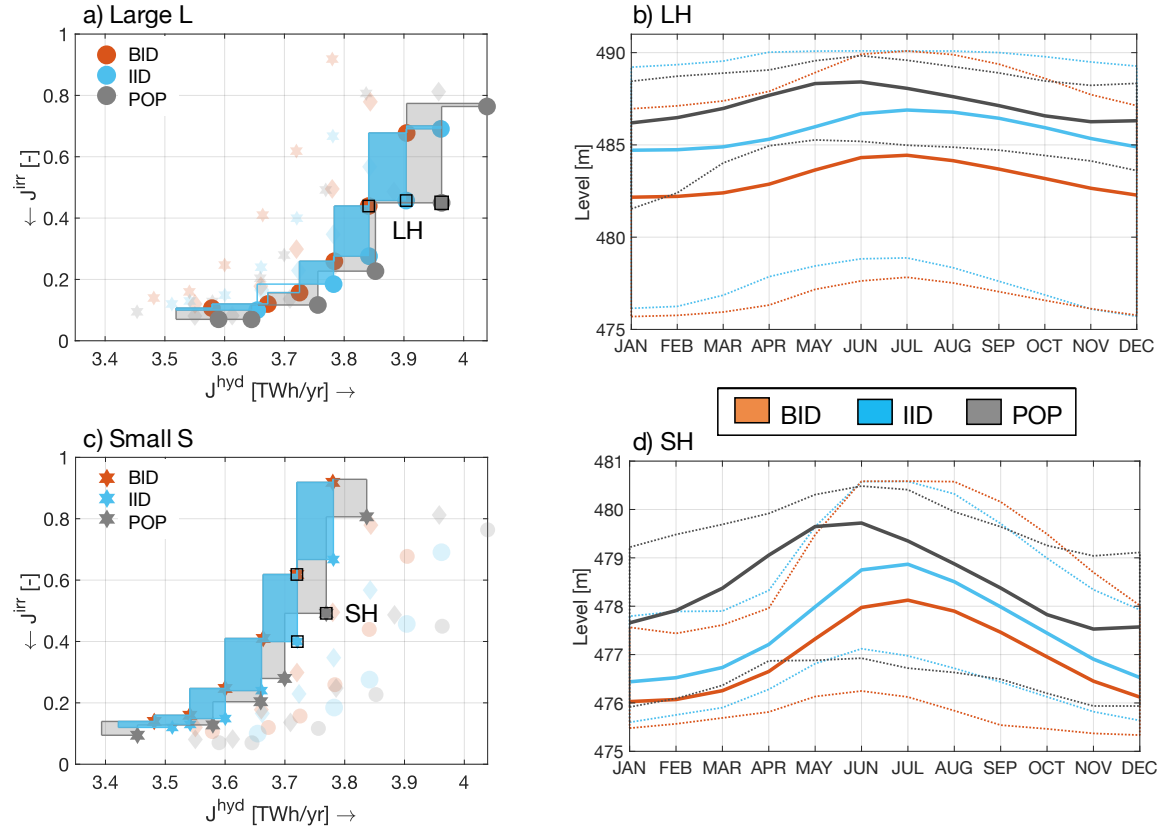


Figure S4. Panel a/c: 2-D management objective space for the large L and small S dam sizes respectively, where the BID (orange), IID (cyan) and POP (grey) solutions associated to a hydropower-prone operating policy H are squared in black. Panel b/d: monthly cyclo-stationary level trajectories for the three solutions highlighted in panels a and c, respectively. Dotted lines bound the 5-th and 95-th percentiles of the monthly levels, whereas bold lines identify the monthly cyclo-stationary average.




Figure S5. Monthly cyclo-stationary level and release trajectories (panel b) for a small S dam size associated to a hydropower-prone operating policy H, corresponding to the solutions highlighted in panel a. Dotted lines bound the 5-th and 95-th percentiles of the monthly levels and releases, whereas bold lines identify the monthly cyclo-stationary average.




Figure S6. For each dam size, namely large L (panel a), medium M (panel b), small S (panel c), the forecast value associated to the Informed Infrastructure Designs identified under basic information (orange), together with perfect (cyan), +20% over-estimated (green), +30% over-estimated (yellow), and +40% over-estimated (red) seasonal forecasts is quantified in terms of hypervolume. Arrows indicate the direction of preference in the hypervolume metric.




Figure S7. For each dam size, namely large L (panel a), medium M (panel b), small S (panel c), the forecast value associated to the Informed Infrastructure Designs identified under basic information (orange), together with perfect (cyan), -20% under-estimated (green), -30% under-estimated (yellow), and -40% under-estimated (red) seasonal forecasts is quantified in terms of hypervolume. Arrows indicate the direction of preference in the hypervolume metric.




Figure S8. For each dam size, namely large L (panel a), medium M (panel b), small S (panel c), the forecast value associated to the Informed Infrastructure Designs identified under basic information (orange), together with perfect (cyan), -6% under-dispersed (green), -10% under-dispersed (yellow), and -15% under-dispersed (red) seasonal forecasts is quantified in terms of hypervolume. Arrows indicate the direction of preference in the hypervolume metric.

Table S1. Share of the variance in the reservoir release trajectories (dependent variable) that is not explained by reservoir storage and month of the year (independent variables) in terms of coefficient of non-determination $1-R^2$. Such trajectories are obtained under a hydropower-prone (H), compromise (C) and irrigation-prone (I) perfect operating policies associated to small S, medium M and large L dam sizes.

Dam size	Operational trade-off	$1-R^2$
Small (S)	Hydropower (H)	0.29
	Compromise (C)	0.28
	Irrigation (I)	0.24
Medium (M)	Hydropower (H)	0.34
	Compromise (C)	0.27
	Irrigation (I)	0.20
Large (L)	Hydropower (H)	0.37
	Compromise (C)	0.23
	Irrigation (I)	0.13

References

Bellman, R., 1957. *Dynamic Programming*. Princeton University Press.

Busoniu, L., Ernst, D., De Schutter, B., & Babuska, R. (2011). Cross-entropy optimization of control policies with adaptive basis functions. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, 41(1), 196–209.

Galelli, S., Castelletti, A., 2013. Tree-based iterative input variable selection for hydrological modeling. *Water Resources Research*, 49, 4295–4310.