Designing with Information Feedbacks: Forecast Informed Reservoir Sizing and Operation

F. Bertoni[1], **A. Castelletti**[1], **M. Giuliani**[1], **P. M. Reed**[2]

[1]Department of Electronics, Information, and Bioengineering, Politecnico di Milano, Milano, Italy
[2]School of Civil and Environmental Engineering, Cornell University, Ithaca, NY, USA

Key Points:

- The benefit from using streamflow forecasts in informing joint dam design and operation is explored for the first time
- The use of seasonal streamflow forecasts could maximally enable a 20% reduction in capital costs for large reservoirs
- The value of forecast information changes remarkably with dam size and how operational trade-offs are balanced

Corresponding author: A. Castelletti, `andrea.castelletti@polimi.it`

Abstract

The value of streamflow forecasts to inform water infrastructure operations has been extensively studied. Yet, their value in informing infrastructure design is still unexplored. In this work, we investigate how dam design is shaped by information feedbacks. We demonstrate how flexible operating policies informed by streamflow forecasts enable the design of less costly reservoir relative to alternatives that do not rely on forecast information. Our approach initially establishes information bounds by selecting the most informative lead times of perfect streamflow forecasts to be included in the infrastructure design. We then analyze the design and operational sensitivities relative to realistic imperfect streamflow forecasts synthetically modeled to explicitly represent different biases. We demonstrate our approach through an ex-post analysis of the Kariba dam in the Zambezi river basin. Results show that informing dam design with perfect forecasts enable attaining the same hydropower production of the existing dam, while reducing infrastructure size and associated capital costs by 20%. The use of forecasts with lower skill reduces this gain to approximately 15%. Finally, the adoption of forecast information in the operation of the existing system facilitate an annual average increase of 60 GWh in hydropower production. This finding, extrapolated to the new planned dams in the basin, suggests that consideration of forecast informed policies could yield power production benefits equal to 75% of the current annual electricity consumption of the Zambian agricultural sector. Forecast information feedbacks have a strong potential to become a valuable asset for the ongoing hydropower expansion in the basin.

1 Introduction

Dam design and operation are classically treated as two independent problems. Optimal reservoir capacity sizing has typically been addressed using a least-cost problem framing, aimed at minimizing total costs (e.g., Perelman et al., (2013) in the water sector; Rodriguez et al., (2015) in the energy sector). In particular, when dealing with water infrastructures such as dams, sizing has usually been performed via the Rippl method, based on a sequence of pre-specified, desired releases and observed inflows (e.g., U.S. Army Corps of Engineers, 1975, 1977; Stephenson & Petersen, 1991). In these design frameworks, operations are represented using pre-defined operating policies (e.g., Thomas Jr & Burden, 1963; Hall et al., 1969; Montaseri & Adeloye, 1999; Soils Incorporated Ltd, 2000; Jeuland & Whittington, 2014; Matrosov et al., 2015), without directly coupling reservoir planning and management as interdependent design choices. The studies that have focused on solving planning and management jointly have typically employed Linear Programming (LP) (e.g., Eastman & ReVelle, 1973; Houck et al., 1980; Lall & Miller, 1988; Mousavi & Ramamurthy, 2000; Afzali et al., 2008; Satishkumar et al., 2010; Afshar et al., 2015; Cervigni et al., 2015) or heuristic optimization techniques (Shourian et al., 2008; Paseka et al., 2018) in open-loop formulations (i.e., they do not account adaptive, information feedbacks). More recently, there are examples emerging in the literature that show the potential value of better abstracting adaptive information feedbacks in joint reservoir design and operation (e.g., Geressu & Harou, 2015; Bertoni et al., 2019, and references therein). The joint optimization of dam size and operations in these studies is accomplished by coupling heuristic multi-objective search with stochastic simulation, where operations are abstracted using storage-dependent reservoir release policies. In all of these studies, the operating rules considered are only conditioned upon the reservoir level/storage, and do not account for other potentially valuable sources of information that are now widely acknowledged for forecast informed reservoir operations (Hejazi et al., 2008; Lu et al., 2017; Turner et al., 2017; Baker et al., 2019).

The value of employing forecasts to enhance the operations of existing infrastructures has long been acknowledged (e.g., Kelman et al., 1990; Kim & Palmer, 1997; Faber

& Stedinger, 2001). The theoretically attainable improvement in performance across operating objectives by the forecast informed system is referred to as forecast value (Murphy, 1993). Forecast value may change according to the temporal dynamics of the operating objectives as well as dam size (Anghileri et al., 2016). For example, in a water reservoir system primarily operated to satisfy short-term operating objectives (e.g., flood control), short-term forecasts might be most informative, allowing the system operator to create a buffer volume in advance for mitigating the upcoming flood peak and thus minimize flood damages (e.g., Saavedra Valeriano et al., 2010; Wang et al., 2012; Raso et al., 2014; Zhao et al., 2014). Alternatively, reservoirs operated with respect to long-term objectives (e.g., irrigation water supply) might benefit more from seasonal forecasts (e.g., Hamlet et al., 2002; Maurer & Lettenmaier, 2004; K. Georgakakos et al., 2005; Voisin et al., 2006; Block, 2011; Steinschneider & Brown, 2012; Anghileri et al., 2016). When dealing with a multi-purpose water reservoir system, estimating the associated forecast value becomes more challenging since both short-term and long-term operating objectives must be balanced (e.g., A. Georgakakos et al., 2012; Sreekanth et al., 2012; Xu et al., 2015; Denaro et al., 2017; Lu et al., 2017; Fuchs et al., 2018; Nayak et al., 2018). Prior studies have shown forecast value may change significantly also with dam size (e.g., You & Cai, 2008; Sankarasubramanian et al., 2009; Graham & Georgakakos, 2010; Anghileri et al., 2016; Turner et al., 2017). For example, when operated to guarantee downstream water supply, the operations of both under-sized and over-sized dams become trivial, since the former case always leads to a structural deficit in the system, whereas in over-sized context managers are always able to satisfy demands. Forecasts might therefore have no value in improving their operations. However, it is important to note that the aforementioned studies assess the sensitivity of forecast value with respect to multiple discrete sampled system configurations characterized by different structural features (e.g., different dam sizes, storage capacity-inflow ratios, storage capacity-demand ratios) chosen in absence of a broader search of the candidate design space. Given the complex interdependencies for how dynamics, information choices, and reservoir design decisions interact, the discrete problem decompositions tacit to these sensitivity analyses have significant limitations.

In this study, we investigate the value of streamflow forecasts for informing the coupled design of a water reservoir size and its operations. Building on the robust dam design framework proposed in Bertoni et al. (2019), we assess whether more flexible operating policies informed by streamflow forecasts enable the design of less costly but operationally effective reservoir systems. We achieve this by first identifying the most informative lead times for perfect streamflow forecasts across different dam sizes and operational trade-offs. Then, the forecasts over the selected lead times are included within the coupled dam design and operation problem to quantify an upper bound estimate for the associated forecast value. Lastly, we explore the sensitivity of the forecast value to realistic streamflow forecast biases. We demonstrate the value of our methodological contribution through an ex-post design analysis of the Kariba dam in the Zambezi river basin, a region where there are a large number of dams planned in the near future (World Bank, 2010), motivating the need for innovations in dam design. Kariba is the largest man-made reservoir in Africa, and the dam's reservoir has the potential for large inter-annual carry over water volumes that could benefit from streamflow forecasts to mitigate seasonal and inter-annual drought anomalies.

In summary, the main contributions of the paper include a framework for: (i) theoretically bounding cases for how perfect knowledge of the future inflows shapes changes in dam sizing and operational trade-offs; (ii) selecting the most informative lead times of perfect streamflow forecasts for different dam sizes and operational performance trade-offs; (iii) analyzing the potential benefits associated with including the selected perfect information forecasts into the coupled reservoir design and operation problem; and (iv) assessing the sensitivity of the forecast informed reservoir design and operation alternatives to more realistic imperfect forecast information.

2 Case Study Description

2.1 Kariba Dam

The Kariba reservoir is a regulated lake in the transboundary Zambezi River Basin (ZRB) in southeastern Africa (Figure 1a). Draining 1.37 million km^2, the ZRB is the 4th largest basin in Africa. The river is shared among eight countries, with Zambia, Zimbabwe and Mozambique encompassing nearly 70% of the entire basin (SADC, 2012). Hydropower is a main source of electricity within the basin, generated by four major regulated reservoirs with a total hydropower capacity of 5,145 MW. About 35% of this overall capacity is installed at the Kariba dam, built in 1960 and impounding the largest man-made reservoir in Africa with a surface area of about 5,600 km^2 and a total storage capacity of about 180 km^3 (65 km^3 of which are active storage). Kariba dam feeds two hydropower plants, the North Bank Station in Zambia and the South Bank Station in Zimbabwe, for a total nameplate capacity of about 2,000 MW. The two plants are jointly operated by the Zambezi River Authority. The Kariba dam system is complemented by two irrigation districts, located respectively upstream and downstream the reservoir.

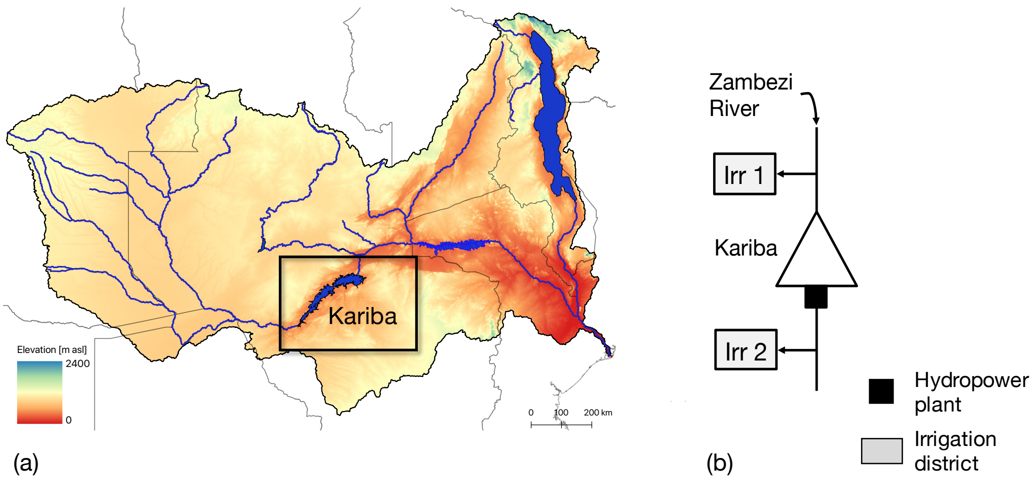

Figure 1: Panel a: Map of the Zambezi River Basin with Kariba dam squared in black. Panel b: Schematic representation of the Kariba multi-purpose reservoir system.

2.2 Model Description

As shown in Figure 1b, the model of the Kariba reservoir system consists of three main components, the reservoir with its hydropower plant and two irrigation districts upstream and downstream. A monthly modeling time-step is employed to capture the Kariba reservoir's dynamics through the following water mass balance equation:

$$s_{t+1} = s_t + i_{t+1} - r_{t+1} - e_t \cdot S_t \tag{1}$$

where s_t is the storage at the beginning of month t, i_{t+1} is the inflow to the reservoir, r_{t+1} is the volume of water released and $e_t \cdot S_t$ is the water evaporated in the time interval $[t, t+1)$. In particular, e_t is the mean monthly evaporation rate, while S_t is the reservoir surface uniquely defined by a non-linear relation given s_t. The actual release $r_{t+1} = f(s_t, u_t, i_{t+1}, e_t, \alpha)$ is formulated according to the non-linear, stochastic relation $f(\cdot)$ between r_{t+1} and the release decision u_t (Soncini-Sessa et al., 2007), which is constrained within a certain zone of operational discretion by the maximum $\bar{u}_t(\alpha)$ and minimum $\underline{u}_t(\alpha)$ feasible release functions, due to the presence of physical (i.e., spillway activation) constraints. Such release functions directly depend upon the dam size $\alpha \in A$, and the extension of the dam operation discretion space enlarges/shrinks proportionally to the dam size considered. As for the reservoir release decision u_t, at each time step it is uniquely defined from an operating policy $u_t = p_{\theta_{res}}(\cdot)$, based on a certain set of inputs (e.g., reservoir storage, time, streamflow forecasts). The policy $p_{\theta_{res}}$ belongs to a pre-defined class of functions, according to which it is parameterized within the space of the parameters $\theta_{res} \in \Theta_{res}(\alpha)$. Note that the interdependency between dam size and operation is expressed in terms of the direct dependence of the feasibility set $\Theta_{res}(\alpha)$ of the policy parameters θ_{res} upon the dam size α. The physical dam size, therefore, constrains the space of operational discretion to reside within a limited range.

Kariba provides storage for two hydropower plants managed by the same operator but with different features: the South Bank is equipped with deeper, more efficient yet smaller turbines than the North Bank (Gandolfi & Togni, 1997). To account for that, in our model we assume the total release r_{t+1} to be split into $r^N_{t+1} = r_{t+1} \cdot \Delta$ for the North and $r^S_{t+1} = r_{t+1} \cdot (1 - \Delta)$ for the South Bank. Δ and $1 - \Delta$ are the normalized turbines efficiencies of the North and South Bank respectively, given the sum of their actual efficiencies (i.e., η^N and η^S, respectively) equal to one.

As for the two irrigation districts (id=1,2), they can abstract water a^{id}_{t+1} from the river through a regulated water diversion channel. The volume of water they can abstract is calculated according to a non-linear hedging rule (Celeste & Billib, 2009):

$$a^{id}_{t+1} = \begin{cases} \min(q^{id}_{t+1}, w^{id}_t \cdot [\frac{q^{id}_{t+1}}{h^{id}}]^{m^{id}}) & \text{if } q^{id}_{t+1} \leq h^{id} \\ \min(q^{id}_{t+1}, w^{id}_t) & \text{else} \end{cases} \tag{2}$$

where q^{id}_{t+1} is the volume of water available in the river upstream of the id-th irrigation district, and w^{id}_t is the monthly water demand. As for the time invariant parameters h^{id} and m^{id} regulating the two diversion channels (id=1,2), they can be grouped into the vector $\theta_{irr} = [h^1, m^1, h^2, m^2] \in \Theta_{irr}$. The diversion rules allow hedging the water abstractions to account for downstream users.

3 Methods and Tools

This study builds on Bertoni et al. (2019) to include streamflow forecasts within the coupled reservoir design and operation problem in order to assess the potential benefits of conditioning design alternatives on their use of key information feedbacks. Our methodology, illustrated in Figure 2, is composed of the following three methodological steps.

The first step is the generation of streamflow forecasts and the selection of the most informative lead times to be included within the dam design phase. Three different sets of streamflow forecasts are considered, perfect seasonal (blue arrows), perfect inter-annual (black arrows), and realistic seasonal (red arrows) forecasts. The perfect seasonal and inter-annual forecasts are broadly used to evaluate the upper bound theoretical value of information at these timescales. The realistic seasonal forecast is intended to reproduce a more realistic decision making environment, where forecasts are affected by systematic biases (i.e., over-estimation, under-estimation, and under-dispersion). Then, we select the most informative lead times of both seasonal and inter-annual forecasts for different dam sizes, based on their ability to best explain the target sequence of reservoir releases derived from a theoretical operating policy (i.e., *perfect operating policy, POP*) (Giuliani et al., 2015), informed by perfect knowledge of the future.

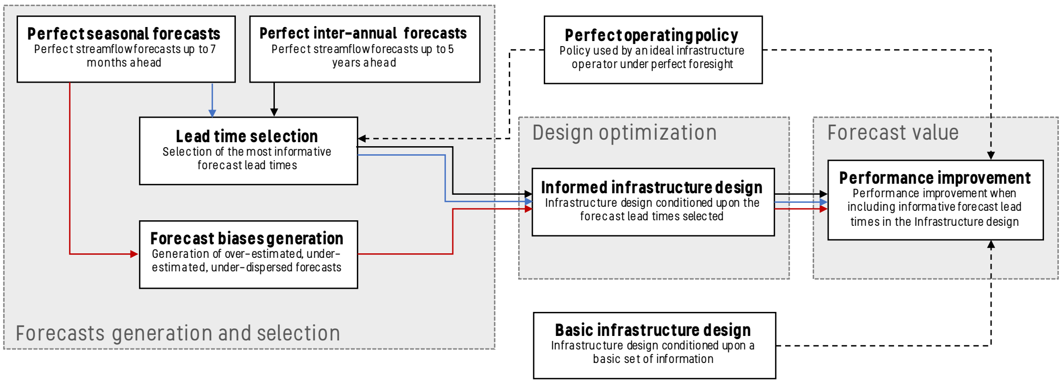

Figure 2: Flowchart of the methodology employed in this study. Each line is colored differently based on the set of streamflow forecasts it refers to, namely perfect seasonal (blue), perfect inter-annual (black), and realistic seasonal (red) forecasts.

The second step of our methodology consists of a joint optimization of reservoir size and operations. In particular, we identify the *informed infrastructure design (IID)*, where operations are informed by the most effective forecast lead times selected in the previous step. This set of optimal infrastructure designs is compared with the *basic infrastructure design (BID)*, which represents the lower bound system performance as the system operations depend upon a basic set of policy inputs traditionally employed in the literature (e.g., reservoir storage).

The last step of our methodological flowchart displayed in Figure 2 is the estimation of the forecast value, namely the performance improvement that could be attained when including the most informative forecast lead times within the infrastructure design phase. Given the upper bound of the forecast value as the difference between the *BID* and *POP* performance, the *IID* is expected to approach the *POP* by partially filling this performance gap. The identified performance improvement represents the forecast value.

3.1 Generation and selection of forecasts

The generation and selection forecasts step consists of first identifying two arrays of perfect streamflow forecasts at both seasonal and inter-annual time scale, from which the most informative lead times are selected to inform the search for design and operation alternatives. Then, a set of realistic seasonal streamflow forecasts characterized by different biases is generated to assess the sensitivity of the resulting forecast informed reservoir design and operation alternatives to the forecast biases.

3.1.1 Perfect forecasts and lead time selection

Initially, we identify an array Ξ_t^s of perfect seasonal streamflow forecasts for different monthly lead times up to a maximum of 7 months ahead. The 7-month lead time is selected to reflect the maximum lead time for seasonal forecasts available currently. The European Centre for Medium-Range Weather Forecasts (ECMWF) (Owens & Hewson, 2018), as well as the most skilful lead time of the long-range Ensemble Streamflow Prediction (ESP) system, developed by the NOAA's National Weather Service (Franz et al., 2003), were taken as references. The selection of the most informative forecast lead times is not a straightforward process. Since the forecast skill usually degrades with a longer lead time (Doblas-Reyes et al., 2011), shorter, more accurate forecast lead times would ideally be selected and relied on. However, since the operational value of forecasts is strictly related to the structural characteristics of the system considered (e.g., reservoir size), bigger dam sizes with an annual carry over capacity might favor longer, less precise forecast lead times to further enhance their operations. In order to select the most informative lead times for different dam sizes, we use an input variable selection technique, following the Information Selection and Assessment (ISA) procedure proposed by Giuliani et al. (2015). Specifically, according to the guidelines in Galelli et al. (2014) and Giuliani et al. (2015), we employ the Iterative Input Selection (IIS) algorithm (Galelli & Castelletti, 2013b) coupled with Extremely Randomized Trees (Geurts et al., 2006; Galelli & Castelletti, 2013a). The IIS algorithm is a hybrid model-based/model-free approach characterized by modelling flexibility (i.e., ability to approximate strongly non-linear functions), computational efficiency, and scalability with respect to the number of candidate inputs. For each dam size, the IIS algorithm follows an iterative procedure to select the most informative seasonal forecast lead times $\mathbf{I}_t^s \in \Xi_t^s$ that best model a target sequence of reservoir releases representing the optimal operation of the system. This sequence is derived for a specific dam size from a theoretical operating policy used by an ideal system operator under perfect knowledge on the future (i.e., perfect operating policy) (for further details about the identification of the target output, refer to both sections 3.4 and S2 in the Supplementary Material).

The role of longer-term forecasts, such as decadal forecasts, is being recently investigated in the literature (e.g., Ham et al., 2014; Schuster et al., 2019; Smith et al., 2019). Decadal forecasts aim at modeling future climatic conditions over a longer time horizon (i.e., the next 10-30 years). Inter-annual and decadal time frames are particularly relevant for infrastructure planners, water resources managers, and others (Meehl et al., 2009), potentially bringing added socioeconomic benefits to infrastructure operations (Choudhury et al., 2019). As a means of theoretically bounding the value of inter-annual forecasts to inform the coupled infrastructure design, we generate a second array Ξ_t^i of perfect streamflow forecasts up to a maximum lead time of 5 years ahead. This longer forecast horizon might benefit particularly large dam sizes, as they have enough storage capacity to carry over large water volumes one or more years. These long carry over periods hold the potential to help mitigate the impacts of inter-annual anomalies related to global climate oscillations (e.g., El Niño Southern Oscillation). When informing their operations with inter-annual forecasts, the active storage of large reservoirs may be managed accordingly to decrease (increase) and make room (compensate) for the large (small) water volumes that will enter the reservoir in the upcoming years, achieving high system performance across the entire evaluation horizon. The same input variable selection technique described for seasonal forecasts is applied to inter-annual forecasts, in order to identify their most informative lead times $\mathbf{I}_t^i \in \Xi_t^i$ for different dam sizes.

3.1.2 Generation of biases in forecasts

In the previous section, we explored perfect forecasts as a prescriptive upper bound for evaluating forecast value. In order to reproduce a more realistic decision making environment and assess the sensitivity of the coupled reservoir design and operation alternatives with respect to typical forecast errors, we generate a set of realistic seasonal streamflow forecasts by incorporating different systematic biases reflecting common errors (i.e., over-estimation, under-estimation, under-dispersion; Cassagnole et al. (2017)). Over-estimated and under-estimated forecasts are symmetrical, while under-dispersed forecasts under-estimate high flows and over-estimated low ones. The accuracy of each biased forecast is evaluated in terms of the percent bias (Pbias) performance metric, which calculates the average tendency of the estimated forecasts with respect to the perfect ones (Moriasi et al., 2007).

3.2 Design optimization

The infrastructure design consists of a coupled multi-objective optimization of both reservoir size and operations, which can be formulated as follows:

$$\begin{aligned} & \pi^* = \arg\min_{\pi} \mathbf{J}_\pi \\ \text{where} \quad & \mathbf{J}_\pi = \left| J^{cost}_\pi, J^{hyd}_\pi, J^{irr}_\pi \right| \\ \text{subject to} \quad & \text{equations 1, 2} \end{aligned} \tag{3}$$

where $\pi = |\alpha, p_{\theta_{res}}, \theta_{irr}|$ is the decision vector, including the dam size $\alpha \in A$, the parametric operating policy $p_{\theta_{res}}$, and the time invariant parameters regulating the two irrigation diversion channels θ_{irr}. Such decision variables are optimized with respect to one planning J^{cost}_π and two management J^{hyd}_π, J^{irr}_π objectives, formulated as follows:

- Minimization of dam construction costs J^{cost}_π [$] discounted over the lifespan of the project:
$$J^{cost}_\pi = c(\alpha) \cdot \frac{r}{1-(1+r)^{-L}} \tag{4}$$
where $c(\alpha)$ [$] is the reservoir construction costs that is proportional to the dam size α, r [yr-1] is the interest rate set at 0.05 (IRENA, 2012) and L [yr] is the lifespan of the project set at 100 years (ibid.), over which construction costs are discounted.

- Maximization of hydropower production J^{hyd}_π [TWh/yr]:
$$J^{hyd}_\pi = \frac{1}{H} \sum_{t=0}^{H-1} \delta g \gamma (\eta^N \bar{h}^N_t q^N_{t+1} + \eta^S \bar{h}^S_t q^S_{t+1}) \tag{5}$$
where δ is a conversion factor to turn hydropower production into [TWh/yr], η^N and η^S are the turbines efficiencies of the North and South Bank, $g = 9.81$ [m/s^2] is the gravitational acceleration, $\gamma = 1000$ [kg/m^3] is the water density, $\bar{h}^N_t$ and $\bar{h}^S_t$ [m] are the net hydraulic heads of the North and South Bank (i.e., reservoir level minus tailwater level), while q^N_{t+1} and q^S_{t+1} [m^3/s] are the turbinated flows at the North and South Bank. If we focus, for example, on the North Bank, the turbinated flow is calculated as follows: $q^N_{t+1} = \min(r^N_{t+1}, \bar{q}^N)$, where $\bar{q}^N$ is the maximum capacity of the turbines at the North Bank and $r^N_{t+1} = \Delta \cdot r_{t+1}$ is the water flowing through them (for further discussion, refer to section 2.2). The same relation holds for the South Bank. In the end, at each time step the total hydropower production is given by the sum of the productions at the two power plants.

- Minimization of total squared irrigation deficit J^{irr}_π [-] normalized with respect to the squared irrigation demand of each district:
$$J^{irr}_\pi = \frac{1}{H} \sum_{t=0}^{H-1} \sum_{id=1}^{2} \left(\frac{\max(w^{id}_t - a^{id}_{t+1}, 0)}{w^{id}_t} \right)^2 \tag{6}$$
where w^{id}_t and a^{id}_{t+1} are the monthly irrigation water demand and abstraction for the id-th irrigation district respectively.

At first, we identify the basic infrastructure design (*BID*) by solving problem 3 with basically informed operating policies associated to alternative dam sizes. Such policies are informed with a basic set of policy inputs (s_t, t), consisting of the reservoir storage s_t (state of the system) and the month of the year t. This basic information is traditionally employed in the literature as the minimum set of information required when designing reservoir operations (Bertsekas, 1976), providing information feedbacks for the reservoir release decision u_t, namely $u_t = p_{\theta_{res}}(s_t, t)$. By conditioning their operations on a minimum number of informative variables, the resulting basic infrastructure designs represent our lower bound system performance.

Then, we solve problem 3 to identify the informed infrastructure design (*IID*), conditioning the operating policies $p_{\theta_{res}}$ associated to alternative dam sizes upon an enlarged set of policy inputs, consisting of storage s_t, month t, as well as the forecast over the identified lead times $\mathbf{I}_t^f \in \Xi_t^f$, with $f = [s, i]$. The resulting informed infrastructure designs differ from the basic only in the formulation of the operating policies associated to different dam sizes, which are now dependent upon the extended set of informative forecast lead times $\mathbf{I}_t^f$ determining the release decision u_t, namely $u_t = p(s_t, t, \mathbf{I}_t^f)$.

We solve both the basic and informed joint planning and management formulations by following the approach proposed in Bertoni et al. (2019), where Evolutionary Multi-Objective Direct Policy Search (EMODPS; Giuliani et al. (2016)) is expanded to include reservoir sizing in addition to operations of both the reservoir itself and the two irrigation diversion channels. EMODPS is a parameterization-simulation-optimization approach (Guariso et al., 1986; Oliveira & Loucks, 1997; Koutsoyiannis & Economou, 2003) that searches candidate parameterized operating policies $p_{\theta_{res}}$ in the space of the parameters $\theta_{res} \in \Theta_{res}$ via Multi-Objective Evolutionary Algorithms (MOEAs). MOEA-based search identifies the set of Pareto approximate policies (i.e., trade-off solutions) whose performance in any single objective can only be improved at the cost of one or more other objectives (Coello Coello et al., 2007). Finding the optimal reservoir operating policy $p^*_{\theta_{res}}$ is therefore equivalent to finding the associated optimal policy parameters θ^*_{res}. In this study, a single optimization problem must be solved over a complex search space to identify the optimal decision vector π, formed by both continuous (i.e., policy parameters θ_{res} and irrigation diversion parameters θ_{irr}) and discrete (i.e., dam size α) decision variables. The water reservoir operating policy is parameterized according to non-linear approximating networks, and in particular Gaussian radial basis functions (RBFs) (for further details about the mathematical formulation of the parameterized operating policy, refer to Section S1 in the Supplementary Material).

The search tasks were implemented using the self-adaptive Borg MOEA algorithm (Hadka & Reed, 2013), since it has been proven to be highly robust across a wide number of challenging multi-objective problems by meeting or exceeding the performance of other state-of-the-art MOEAs (Reed et al., 2013). The Borg MOEA employs multiple global probabilistic search operators for mating, selection and mutation, whose probability of being selected during the optimization phase is linked to their demonstrated ability of generating quality solutions. In particular, we use the hierarchically parallelized version of the Borg MOEA, termed the Multi-Master Borg MOEA (Hadka & Reed, 2015), which has proven to be successful in complex reservoir control problems (Salazar et al., 2017; Giuliani et al., 2018; Quinn et al., 2018). The Multi-Master Borg MOEA exploits communication across multiple master-worker parallel implementations of the Borg algorithm, improving both the algorithm's reliability across random seed trials and the performance of the worst seeds, without degrading the best seeds (Hadka & Reed, 2015; Salazar et al., 2017; Giuliani et al., 2018).

In addition to the *BID* and *IID* formulations, we have to determine the target output of the iterative input selection procedure described in section 3.1.1. This is achieved by solving the management side of the joint optimization problem 3 with respect to the vector of $n = 2$ management objectives (i.e., J_π^{hyd} in equation 5, J_π^{irr} in equation 6), for a fixed dam size $\bar{\alpha}$ and irrigation diversion parameters $\bar{\theta}_{irr}$. For further details about the mathematical formulation of the optimization problem, refer to Section S2 in the Supplementary Material. This pure management problem is solved via Deterministic Dynamic Programming (DDP) (Bellman, 1957) for different dam sizes identified under basic infrastructure design and with respect to the full, deterministically known trajectory of external drivers (i.e., streamflows) over the entire evaluation horizon H. For each dam size, we therefore obtain an optimal operating policy, 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 (refer to section 3.4 for further details). Given that this policy is identified under perfect knowledge of the future (i.e., POP), it represents the upper bound system performance.

3.3 Forecast value

Prior literature has defined forecast value as the operational value of employing forecasts to enhance system operations, namely their effectiveness in supporting decisions (Anghileri et al., 2016; Turner et al., 2017), and quantified in terms of performance improvement in the system operation objectives (Murphy, 1993). In this study, we first assess the theoretical upper bound of attainable forecast value by estimating the maximum space for improvement - also known as Expected Value of Perfect Information (EVPI) - that is in principle attainable under the assumption of full and perfect (deterministic) information on the future when operational decisions must be made (Giuliani et al., 2015; Denaro et al., 2017, for further details, refer to Section S2 in the Supplementary Material). We define a multi-objective performance envelope bounded by the *POP* and *BID* results. When dealing with single-objective problems, the single management objective considered assumes one scalar value for the *POP* and the *BID*, making it trivial to quantify the difference between these performance bounds. For multi-objective problems, both *BID* and *POP* performance is a vector solution set of the optimization problem 3 (i.e., Pareto optimal or approximate set). Following Zitzler et al. (2003), we use finite set theory to quantify the candidate space for multi-objective performance improvement, using the hypervolume (HV) indicator that measures the volume of objective space dominated by a Pareto set. The HV measure has the benefit of capturing both convergence ("proximity to the best known solutions") and diversity ("representation of the full extent tradeoffs"). The upper bound of the forecast value is calculated as the difference in hypervolume between the ideal optimal Pareto front (i.e., the *POP*'s performance) and the approximation set attained using only information on the state of storage (i.e., the no forecast constrained *BID*), where the Pareto front associated to the higher hypervolume is the better. As a rule, the larger the delta is between these two sets of solutions, the more the system can benefit from including more forecast information during its joint optimization of candidate designs and operations.
This Expected Value of Perfect Information is expected to be partially covered by the informed infrastructure design, filling the performance gap between the perfect (upper bound) and basic (lower bound) solutions and drawing its system performance as close as possible to the *POP*. Such performance improvement corresponds to the actual forecast value and is calculated as the difference in hypervolume between the two approximation sets (i.e., informed minus basic infrastructure designs).

3.4 Computational experiment

Our computational experiment has the following structure:

- *Perfect seasonal streamflow forecasts*: forecasts of Kariba inflows from 1974 to 2005 computed over seven different lead times, ranging from 1 month to 7 months ahead. In addition to the cumulative future streamflow, both the minimum and maximum over 7 months are also included (see Table 1). The minimum future streamflow allows to acquire perfect knowledge on the most severe drought that will affect the system in the next 7 months. The maximum future streamflow allows to acquire perfect knowledge on the maximum flood peak that will enter the Kariba dam in the next 7 months.

Table 1: Set of perfect seasonal streamflow forecasts calculated over different lead times.

Name	Description	Period
$q1_t,\ldots,q7_t$	Cumulative future streamflow over 1,...,7 months	1974-2005
$qm7_t$	Minimum future streamflow over 7 months	1974-2005
$qM7_t$	Maximum future streamflow over 7 months	1974-2005

- *Perfect inter-annual streamflow forecasts*: in addition to the set of seasonal forecasts in Table 1, we also consider inter-annual perfect forecasts of Kariba inflows from 1974 to 2005. In particular, we add the median of future streamflows over 12, 24, 36, 48, and 60 months ahead. Other temporal aggregation metrics used to characterize streamflow forecasts, such as the cumulative future streamflows, as well as the maximum and minimum over subsequent months, usually provide exact information on the amount of water that will enter the system in the near future. Therefore, their skill rapidly degrades with longer lead times (Doblas-Reyes et al., 2011). Due to the multi-year time resolution associated to inter-annual forecasts and the difficulties related to their exact estimate, we use the median to characterize them because it provides a rough estimate of the water volume entering the system in the next years, suggesting whether the upcoming years will be rather wet/dry in median.
- *Realistic streamflow forecasts*: since the accuracy of each realistic forecast is evaluated in terms of the Pbias performance metric, we use a +30% Pbias for over-estimated, -20% for under-estimated and -10% for under-dispersed seasonal forecasts according to Cassagnole et al. (2017). A sensitivity analysis of the *IID* alternatives with respect to different realistic forecast Pbiases is reported in Section S7 in the Supplementary Material.
- *Basic infrastructure design*: *BID* solutions are designed via EMODPS over the 1974-2005 evaluation horizon. Based on Bertoni et al. (2019), three dam sizes are selected such that they uniformly cover the entire set of optimal system configurations identified under basic information, namely a small S = 128 km^3, a medium M = 148 km^3 and a large L = 188 km^3 dam size. Note that this includes an alternative that is very similar to the existing Kariba dam's size (188 km^3 vs 180 km^3 respectively).
- *Perfect operating policy*: *POP* solutions are designed via DDP over the 1974-2005 evaluation horizon for each of the three dam sizes selected (for further details, refer to Section S2 in the Supplementary Material). Since DDP requires to solve a single-objective problem, we use the weighting method (Saaty & Gass, 1954) to convert the 2-objective problem discussed in section 3.2 into a single-objective one via convex combinations. The operational trade-offs between the hydropower production and irrigation deficit objectives are explored by varying the weights used for aggregating the objectives. For each dam size, three target *POP*s associated to three different target trade-offs between the two management objectives are used as target outputs of the iterative input selection procedure to identify the most informative forecast lead times.
- *Information selection*: for each target *POP* trade-off to be explained and each dam size, the Iterative Input Selection algorithm is used to select the most in-

formative forecast lead times that mostly explain the target sequence of optimal releases. At first, we perform a regression on a sample dataset consisting of the Kariba storage s_t and month of the year t. Being these two variables highly correlated with the target output to be explained, they would overshadow the real contribution of other potentially informative variables if jointly considered in the information selection phase. Then, the IIS algorithm is run on the set Ξ_t^s of perfect seasonal streamflow forecasts presented in Table 1 to select the most informative lead times and temporal aggregation metrics (i.e., maximum and minimum over 7 months) $\mathbf{I}_t^s \in \Xi_t^s$ explaining the model residuals of s_t and t. The same procedure is then repeated for the set of perfect inter-annual streamflow forecasts Ξ_t^i, where the IIS algorithm must select the most informative lead times only $\mathbf{I}_t^i \in \Xi_t^i$ since inter-annual forecasts are characterized by the median of future streamflows over multiple years.

- *Informed infrastructure design*: *IID* solutions are designed via EMODPS over the 1974-2005 evaluation horizon. Five sets of informed infrastructure designs are optimized, conditioning their operations upon: (i) most informative lead times and temporal aggregation metrics selected for perfect seasonal forecasts; (ii) most informative lead times selected for perfect inter-annual forecasts; (iii) +30% over-estimated seasonal forecasts; (iv) -20% under-estimated seasonal forecasts; (v) -10% under-dispersed seasonal forecasts. In order to estimate the forecast value, for each set of solutions we analyze the same three dam sizes selected under basic information from the set of optimal, informed system configurations.

Both the *BID* and *IID* formulations have been solved using the Multi-Master Borg MOEA algorithm (see section 3.2), which is based on an epsilon dominance archiving, requiring the users to specify a numerical precision for each optimization objective below which they are insensitive to changes in performance. We use epsilon dominance values equal to 0.06 for J_π^{hyd}, 0.01 for J_π^{irr}, and $4.8 \cdot 10^9$ for J_π^{cost}, representing the significance of precision that is considered consequential in evaluating decision trade-offs. Each optimization problem was run for 10 random seeds in order to improve solution diversity and avoid randomness dependence, using a 4-master implementation. Each seed was run up to 1 million function evaluations, proved to be sufficient by visual inspection of search progress, with little variability across seeds (refer to Section S3 in the Supplementary Material). The remainder of the Multi-Master Borg MOEA algorithm's parameters were set using the defaults recommended in prior studies (Hadka & Reed, 2015; Salazar et al., 2017). For each computational experiment, the final set of Pareto approximate system configurations was computed as the reference set of non-dominated solutions obtained across the 10 optimization trials. The experiments were run on the Cube cluster at the Cornell Center for Advanced Computing, running CentOS 7.6 across 32 compute nodes with Dual 8-core E5-2680 CPUs at 2.7 GHz, 128 GB of RAM, using 192 cores per island for a total of 3840 computational hours.

4 Results and Discussion

4.1 Influence of Dam Size on the Expected Value of Perfect Information

In order to quantify the actual value of forecasts and their effectiveness in enhancing dam design, we must first find the maximum space for improvement that is theoretically attainable with full and perfect (i.e., deterministic) foresight, namely the Expected Value of Perfect Information. This represents the upper bound system performance. A small space for improvement means that the potential gain in system performance achievable by using streamflow forecasts is negligible (and the converse for a large space for improvement).

Figure 3 shows the performance for three different dam sizes, namely small S (stars, panel a), medium M (diamonds, panel b), and large L (circles, panel c), in terms of hydropower production (J^{hyd}) and irrigation deficit (J^{irr}), and the associated EVPI obtained by comparing *BID* results exploiting only standard storage information (orange) and the *POP* baseline with perfect foresight (grey). For each dam size, the corresponding EVPI is calculated as the difference in hypervolume between basic and Perfect solutions (grey shaded area) and reported in the bottom right corner of each panel (the absolute values of the hypervolume metric for both sets of solutions are reported in Figure 7). The Expected Value of Perfect Information increases from 0.54 to 0.62 as we move from small to large dam sizes. Larger reservoirs provide an increased active storage capacity and operational flexibility to carry over significant water volumes across different temporal scales (i.e., from intra-monthly to inter-annually). The operations of large dams might thus benefit more from additional information on future hydro-climatic conditions of the system (e.g., future streamflows accumulated over several months). Moreover, regardless of dam size, the grey shaded area (i.e., EVPI) shrinks as we move from a hydropower preference (top right) to an irrigation focus (bottom left) in the objective space. When the reservoir is operated to maximize hydropower, information on future hydrologic conditions is needed every month of the evaluation horizon to always keep the reservoir full to sustain constant releases while minimizing spillages. An irrigation focused operating policy is dominated by the need to meet the target irrigation demands with a maximum peak in August/September. Since there is a structural deficit in the system, namely a small deficit always occurs regardless of water availability and reservoir operations, the information on the amount of water entering the reservoir in subsequent months is less valuable. Therefore, better informing reservoir operations can never achieve zero deficit.

In each of the three panels in Figure 3, we highlighted three different operational preferences: increased hydropower (H), compromise (C) and an emphasis on irrigation (I) (black squares). Our evaluations of the value of streamflow forecasts in the subsequent results are based on these 9 nine solutions.

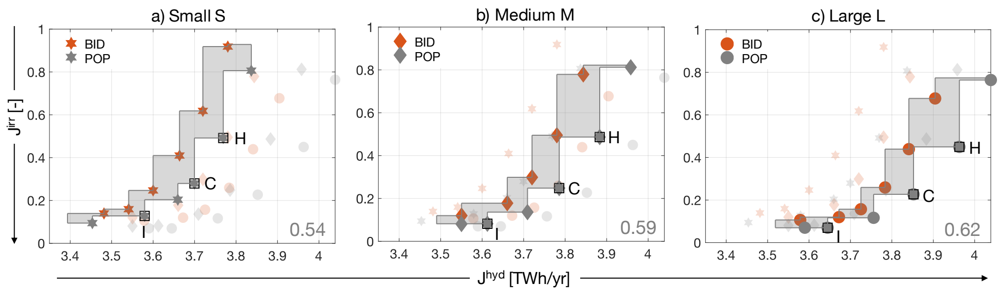

Figure 3: The objective space performance comparisons of the basic infrastructure designs (orange) and the corresponding perfect operating policies (grey) for each of the three dam sizes selected, namely small S (stars, panel a), medium M (diamonds, panel b), and large L (circles, panel c). The grey shaded area represents the Expected Value of Perfect Information, also reported in the bottom right corner of each panel. Arrows indicate the direction of preference in the objectives.

4.2 Forecast Informed Infrastructure Design Using Perfect Seasonal Forecasts

After assessing the Expected Value of Perfect Information as well as the potential benefits that are achievable when hydropower strongly shapes solution preferences, we now identify via Iterative Input Selection the most informative seasonal forecast lead times $\mathbf{I}_t^s \in \Xi_t^s$ for the three target trade-offs for each of the three dam sizes (Figure

3). More details on the results of this information selection phase can be found in Section S4 in the Supplementary Material. Independent of dam size and operational performance trade-offs, the most informative variables selected always provide information on 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 monthly lead times. In particular, hydropower focused policies are best informed by the maximum future streamflow over 7 months $qM7_t$ for all three dam sizes. This input variable for the control policies 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. We employ this variable as the first, most informative input to be included in the informed infrastructure design phase in addition to reservoir storage and time. The second most informative input, the minimum future streamflow over 7 months $qm7_t$, had a marginal impact control policies' actions and was therefore not considered any further.

Figure 4 compares the associated forecast value obtained by *BID* (orange) and *IID* (cyan) solutions for the small S (stars, panel a), medium M (diamonds, panel b), and large L (circles, panel c) dam sizes in terms of J^{hyd} and J^{irr}. For each dam size, the forecast value is computed as the difference in hypervolume between basic and informed solutions (cyan shaded area), whereas the grey shaded area represents the residual space for improvement that is not yet explained (the absolute values of the hypervolume metric for all sets of solutions are reported in Figure 7). Regardless of dam size, the informed system performs better than the basic, moving closer to the set of *POP*s.
For large dam sizes, the forecast value is characterized by a +32% hypervolume increase from 0.38 (basic) to 0.50 (informed), covering about 20% of the corresponding space for improvement. This corresponds to a +80 GWh/yr hydropower production increase and no changes in terms of irrigation deficit on average across all the basic and informed solutions. As for medium dam sizes, the forecast value is the smallest, with a +22% hypervolume increase from 0.41 (basic) to 0.50 (informed), covering about 15% of the corresponding space for improvement. This corresponds to a +42 GWh/yr hydropower production increase and a 0.05 normalized irrigation deficit decrease on average across all the basic and informed solutions. In the end, small dam sizes are associated to the highest forecast value, with a +33% hypervolume increase from 0.46 (basic) to 0.61 (informed), covering about 28% of the corresponding space for improvement. This corresponds to a 0.13 normalized irrigation deficit decrease and no changes in terms of hydropower production on average across all the basic and informed solutions.
Such improvements are particularly evident in the hydropower focused region of the objective space, which was already expected to attain the highest enhancement when including informative variables in infrastructure design. Here, the cyan shaded area is larger and the grey shaded area shrinks accordingly. In particular, if we fix a specific high level of hydropower production (e.g., $J^{hyd} \simeq 3.84$ TWh/yr), a 20% reduction in capital costs could be attained by designing a medium dam operated with forecast information (i.e., cyan diamond in Figure 4b), that produces the same hydropower as that of a larger reservoir informed with the basic set of policy inputs (i.e., orange circle in Figure 4c). Given this fixed level of hydropower production and a large dam size (Figure 4c), the informed infrastructure design (cyan) is able to produce 0.06 TWh/yr (60 GWh/yr) more hydropower than the corresponding basic solution (orange), moving from a 3.84 TWh/yr to a 3.90 TWh/yr absolute value in the hydropower objective performance. This improvement is particularly significant as it corresponds to more than 25% of the yearly average electricity consumption by the agriculture sector in Zambia, where the Kariba dam is located, recorded over the 2014-2017 period (IEA, 2019). Attained under perfect streamflow forecasts, 60 GWh/yr of additional hydropower production might not be fully achievable in the real world where forecasts may be affected by systematic errors. However, achieving 50% of such increase (i.e.,

30 GWh/yr) might be realistic and still significant, as it corresponds to the yearly average electricity consumption by the transport sector in Zambia, recorded over the 2014-2017 period (IEA, 2019). It is also important to notice that this hydropower improvement is attained at no additional cost for irrigation, as the irrigation deficit remains unchanged.

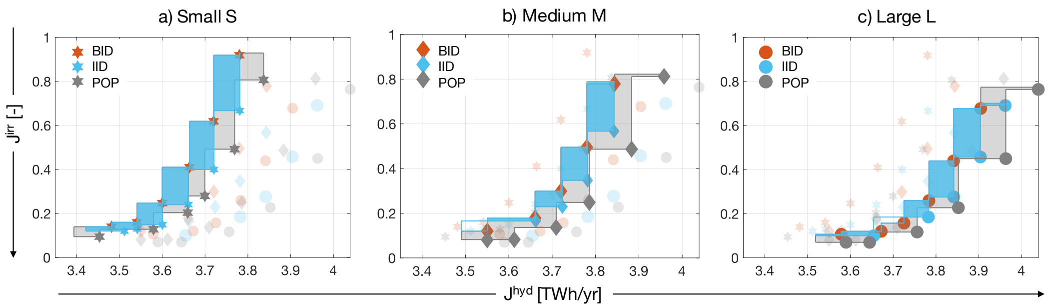

Figure 4: Comparison of the two objective tradeoffs that result from the basic infrastructure designs (orange), the informed infrastructure designs (cyan) and perfect operating policies (grey) for each of the three dam sizes selected, namely small S (stars, panel a), medium M (diamonds, panel b), and large L (circles, panel c). The cyan shaded area represents the forecast value, whereas the grey shaded area corresponds to the residual space for improvement to still be filled. Arrows indicate the direction of preference in the objectives.

To further understand the effects of streamflow forecasts on enhancing the reservoir system design, we analyze the system dynamics achieved under a hydropower focused policy, where we have observed the greatest improvement. Figure 5 displays the system dynamics for medium dam sizes in terms of levels (panel b), inflows (panel c) and dry year levels (panel d) trajectories associated to the basic (orange), informed (cyan) and perfect (grey) solutions MH highlighted in panel a (for other dam sizes, refer to Section S5 in the Supplementary Material). Since the maximum future streamflow over 7 months $qM7_t$ 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 reservoir levels about 2 meters higher than the basic and closer to the perfect trajectories without spilling (Figure 5b). This leads the informed solution to a +2% further increase in hydropower production with respect to the basic one, approaching the performance achieved under *POP* (Figure 5a). On the contrary, since *BID* relies on the reservoir storage and time only, the system operator does not have any information on the future streamflows entering the reservoir. Being afraid of spilling and consequently wasting possible production, the operator keeps the reservoir levels very low, without exploiting the full hydropower potential of the dam. This is particularly evident during dry years (e.g., 1994), when low reservoir levels contribute to further decreasing hydropower production under basic infrastructure design (Figure 5d). Since less water enters the reservoir (red dotted line in Figure 5c), releases must be reduced in order not to further lower the levels and thus the hydropower potential, causing production to be decreased even further.

4.3 Forecast Informed Infrastructure Design Using Perfect Inter-Annual Forecasts

To this point, our analysis has considered the value of an array of perfect seasonal streamflow forecasts over different monthly lead times up to the maximum lead time

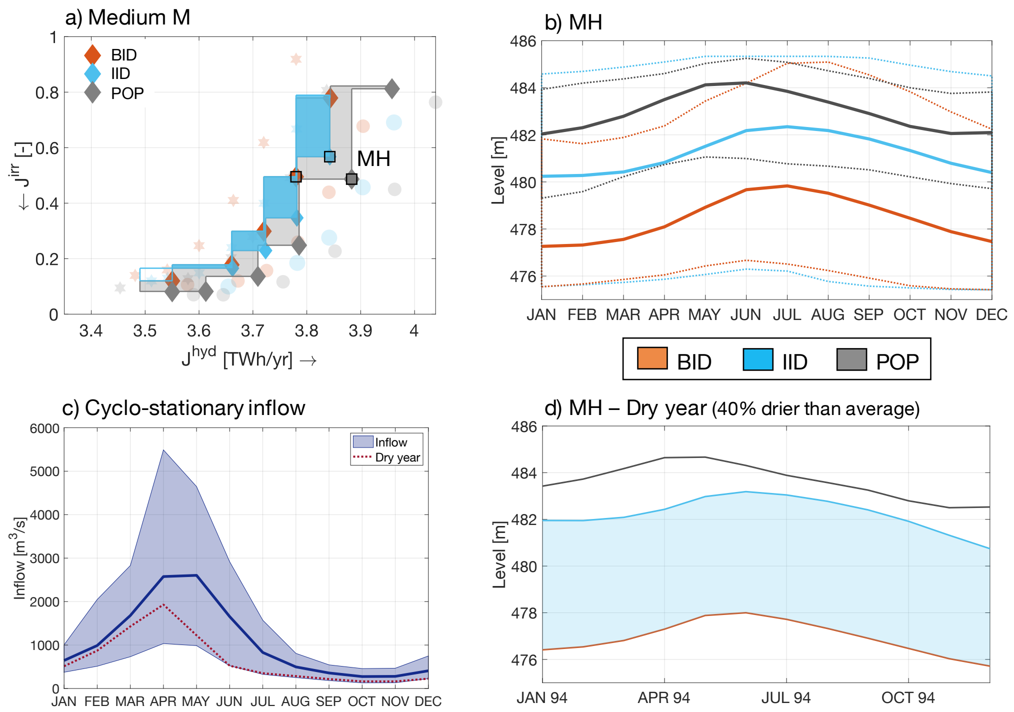

Figure 5: Panel a: Two objective performance trade-offs for the medium M dam size, where the *BID* (orange), IID (cyan) and *POP* (grey) solutions associated to a hydropower-prone operating policy H are squared in black. Panel b: monthly cyclo-stationary level trajectories for the three solutions highlighted in panel a. Dotted lines bound the 5-th and 95-th percentiles of the monthly levels, whereas bold lines identify the monthly cyclo-stationary average. Panel c: monthly cyclo-stationary inflow trajectory of the Kariba dam. The shaded area is bounded by the 5-th and 95-th percentiles of the monthly inflows, whereas the bold line identifies the monthly cyclo-stationary average. The red dotted line corresponds to the inflow trajectory of a dry year (i.e., 1994). Panel d: monthly level trajectories for the three solutions highlighted in panel a during a dry year (i.e., 1994). The cyan shaded area covers the distance between the *BID* and *IID* level trajectories.

of seasonal forecasts provided by weather forecast centers. However, it is also interesting to assess whether inter-annual streamflow forecasts contribute any further benefit particularly for large dam sizes, which can carry over large water volumes year-to-year. Even if inter-annual forecasts are not yet very accurate, their skill is expected to considerably increase in the near future. To this end, we employ the Iterative Input Selection algorithm to identify the most informative lead times $\mathbf{I}_t^i$ out of an additional set of perfect inter-annual streamflow forecasts Ξ_t^i. The full details for our analysis of the inter-annual perfect information selection phase can be found in Section S4 in the Supplementary Material. In addition to the maximum future streamflow over 7 months qM7_t, the IIS algorithm selected the median of future streamflows over the next 12 months qmed12_t, as an additional, not negligible informative variable to be included in the informed infrastructure design phase for all dam sizes.

Figure 6a displays the performance of large L dam sizes in terms of J^{hyd} and J^{irr} achieved under basic (orange), informed under seasonal forecasts (cyan), informed under inter-annual forecasts *IID* - IA (blue) infrastructure designs and perfect operating policies (grey). Inter-annual forecasts coupled with qM7_t bring particular advantages in the hydropower focused region of the objective space, allowing the *IID* - IA alternatives to approach the *POP* set. This information allows the system operator to acquire perfect knowledge not only on the magnitude of the upcoming flood peak, but also whether the next year will be wet or dry relative to the median. The operator is therefore confident in storing more water and keeping the levels higher without spilling, consequently increasing hydropower production and reducing irrigation deficit. However, this operating strategy can be applied to large dam sizes only, for which inter-annual forecasts are valuable as they can store significant water volumes and carry them over inter-annually. Adding inter-annual forecasts does not bring any additional benefits in the irrigation-prone area, where the potential improvement is extremely limited.

Figure 6 displays the levels (panel b), inflows (panel c) and wet year levels (panel d) trajectories associated to the basic (orange), informed under seasonal forecasts (cyan), informed under inter-annual forecasts (blue), and perfect (grey) solutions LH highlighted in panel a. As expected, under *IID* - IA the reservoir levels are about 2.5 and 0.5 meters higher than the basic and seasonally informed respectively on average, moving closer to the Perfect trajectories (Figure 6b). Such difference is not big enough for allowing *IID* - IA to further increase hydropower production with respect to the seasonally informed system, yet it is sufficient for storing more water needed to satisfy the irrigation demand, attaining a 20% reduction in the irrigation deficit. When compared with the basic system, however, the difference in the reservoir levels is significant, allowing *IID* - IA to achieve a 2% higher hydropower production and a 15% lower irrigation deficit (Figure 6a). These system dynamics are particularly evident during wet years (e.g., 1976, red dotted line in Figure 6c), where the contribution of qmed12_t is evident in terms of keeping the levels 3 meters higher than the informed and closer to the Perfect trajectories (Figure 6d). In this case, the system operator is aware not only that the flood peak entering the reservoir will be one of the highest recorded across the entire time-period, but also how wet the next 12 months will be, increasing the levels accordingly without spilling.
It must be noted that such performance improvements have been attained under perfect inter-annual streamflow forecasts. However, in a realistic decision making environment inter-annual forecasts are still not reliable, since their longer-term resolution makes their correct estimation challenging.

4.4 Forecast Informed Infrastructure Design Using Biased Forecasts

This section explores how much performance might degrade with imperfect seasonal streamflow forecasts characterized by different biases, namely over-estimation, under-estimation, under-dispersion. We employ the same settings of the *IID* analysis presented in section 4.2, where the number of policy inputs to be included in the optimal infrastructure design is limited to s_t, t, and qM7_t.

Figure 7 displays the performance of the small S (stars, panel a), medium M (diamonds, panel b) and large L (circles, panel c) dam sizes in terms of J^{hyd} and J^{irr} and the associated hypervolume metrics, associated to the set of basic infrastructure designs (orange) and perfect operating policies (grey), along with the suite of over-estimated (yellow), under-estimated (green), under-dispersed (purple) and perfect (cyan) seasonal forecasts. Regardless of dam size, the sets of *IID*s are less sensitive to both under-estimated and under-dispersed forecasts, attaining almost the same hypervolume as for perfect forecasts. On the contrary, over-estimated forecasts lead to the lowest value of the HV metric for all three dam sizes. Since informed solutions are conditioned upon the maximum future streamflow over 7 months qM7_t, over-estimated forecasts lead the system to release significant water volumes in order not to waste excess flood water. However, when less water actually flows into the reservoir, the levels (and thus hydropower production) decrease while irrigation deficit increases, since less water is then available to be released for meeting the demand. Both L and S dam sizes achieve a 4% lower HV under over-estimated with respect to perfect seasonal forecasts,

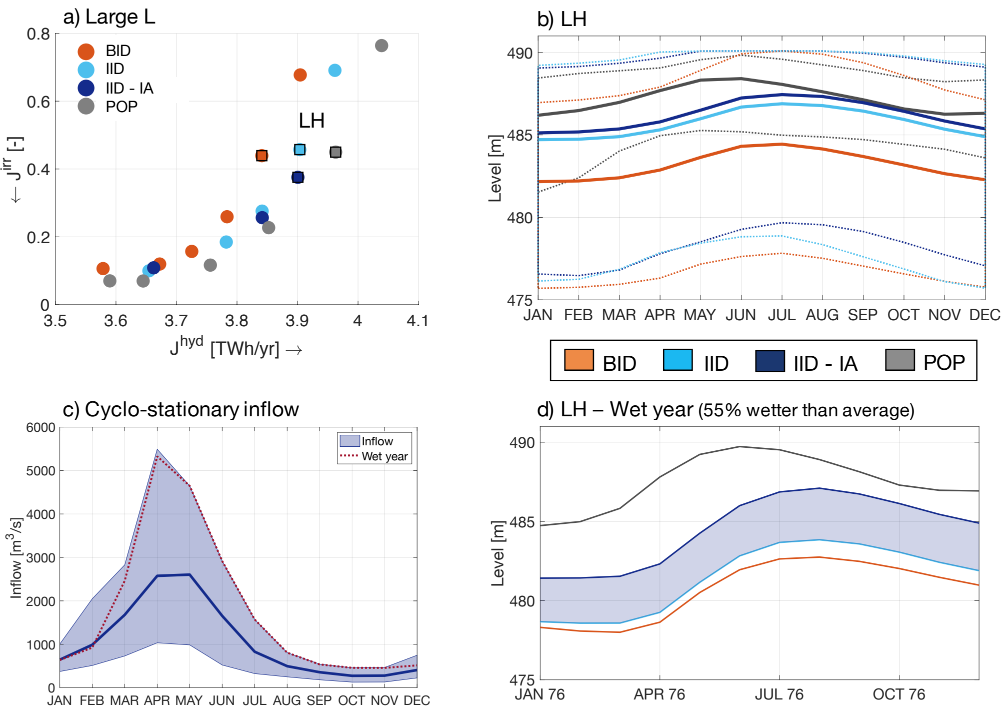

Figure 6: Panel a: Two objective performance trade-offs for the large L dam size, where the *BID* (orange), *IID* (cyan), *IID* under inter-annual forecasts (blue) and *POP* (grey) solutions associated to a hydropower focused operating policy H are squared in black. Panel b: monthly cyclo-stationary level trajectories for the four solutions highlighted in panel a. Dotted lines bound the 5-th and 95-th percentiles of the monthly levels, whereas bold lines identify the monthly cyclo-stationary average. Panel c: monthly cyclo-stationary inflow trajectory of the Kariba dam. The shaded area is bounded by the 5-th and 95-th percentiles of the monthly inflows, whereas the bold line identifies the monthly cyclo-stationary average. The red dotted line corresponds to the inflow trajectory of a wet year (i.e., 1976). Panel d: monthly level trajectories for the four solutions highlighted in panel a during a wet year (i.e., 1976). The blue shaded area covers the distance between the level trajectories of the *IID* and *IID* under inter-annual forecasts.

whereas M dam sizes are characterized by an 8% reduction. Medium dam sizes are therefore more sensitive to an over-estimation in the streamflow forecasts. Since they do not have enough storage capacity to store large flood water volumes, they must be carefully operated in order not to spill and thus waste both hydropower potential and water for irrigation supply. If future streamflows are over-estimated and less water actually enters the reservoir, the boundary conditions under which these dams have been optimized change. Their operation is therefore not optimal anymore with respect to what actually unfolds, preventing these reservoirs to be carefully operated in order to exploit their full active storage capacity. As for large and small dam sizes, they are less sensitive to over-estimated forecasts. The former present enough active storage capacity to compensate over-estimated flows. The latter are characterized by a small operation discretion space that does not allow for substantially different operating strategies, since large water volumes are already spilled under perfect streamflow forecasts.

Since the sets of *IID*s are more sensitive to over-estimated forecasts, it is interesting to analyze the effects of such biased forecasts on the system dynamics for both medium and large dam sizes when operated under a hydropower-prone operating policy (i.e., MH and LH in Figure 7b and 7c respectively; for small dam sizes, refer to Section S6 in the Supplementary Material). Figure 8a and 8b display the system dynamics of LH and MH respectively in terms of level and release cyclo-stationary trajectories, attained under basic infrastructure design (orange) and perfect operating policy (grey), along with informed infrastructure design under over-estimated (*IID* (Over-est) - yellow) and perfect (*IID* (Perfect) - cyan) seasonal streamflow forecasts. When comparing *IID* (Over-est) and *IID* (Perfect) solutions, both LH and MH solutions present the same dynamics. In particular, over-estimated forecasts, which over-estimate the wet season flood peaks, force both reservoirs to release more water from January to March in order not to spill, while lowering releases during the dry season (September-October) in order not to excessively lower the reservoir levels and still maintain a satisfactory level of hydropower production. On the contrary, perfect seasonal forecasts enable both reservoirs to release less during the wet season, saving water for the dry season when irrigation demand is higher. As for the level dynamics, both *IID* (Perfect) and *IID* (Over-est) present the same trajectories under both LH and MH solutions, keeping the levels about 2 meters higher on average than the corresponding basic trajectories. These dynamics allow both informed infrastructure designs to attain the same hydropower production regardless of dam size, yet over-estimated forecasts attain a 8% and 20% higher irrigation deficit under LH and MH solutions respectively.

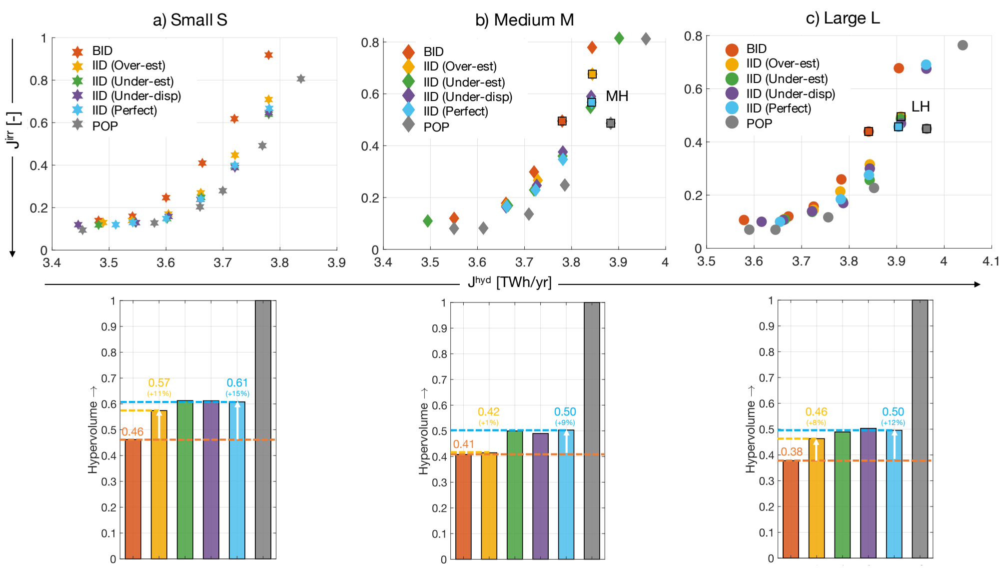

Figure 7: Two objective performance trade-offs where the basic infrastructure designs (orange) are compared against the corresponding informed infrastructure designs and perfect operating policies (grey) for each of the three dam sizes selected, namely small S (stars, panel a), medium M (diamonds, panel b), and large L (circles, panel c). For each dam size, the sets of *IID*s are identified under perfect (cyan), over-estimated (yellow), under-estimated (green), and under-dispersed (purple) forecasts and the associated value of information is quantified in terms of hypervolume. Arrows indicate the direction of preference in the objectives.

5 Conclusions and Future Work

This paper investigates the value of streamflow forecasts in informing the coupled design of a water reservoir size and its operations, exploring their interdependencies

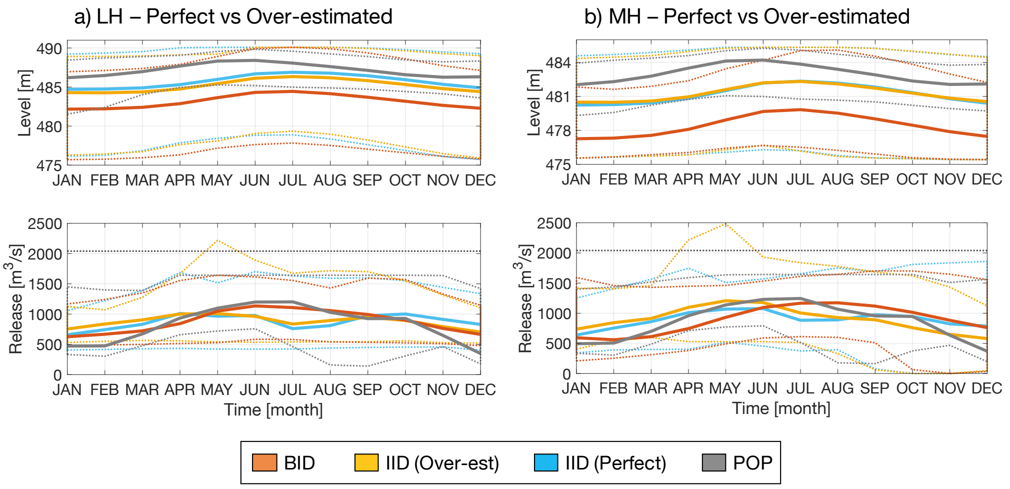

Figure 8: Monthly cyclo-stationary level and release trajectories for large L (panel a) and medium M (panel b) dam sizes associated to a hydropower-prone operating policy H, corresponding to the solutions highlighted in panels b and c of Figure 7. Dotted lines bound the 5-th and 95-th percentiles of the monthly levels and releases, whereas bold lines identify the monthly cyclo-stationary average.

and how streamflow forecasts shape them to enhance the final infrastructure design. Our main contribution in this paper is to explicitly map how informative streamflow forecasts shape multi-objective dam design trade-offs. The main goal of this study is to assess if the inclusion of forecast information can lead to less costly, more efficient water reservoirs with respect to less informed systems. We demonstrate the potential of our contribution through an ex-post design analysis of the Kariba dam in the Zambezi river basin.

Our results show that the operation of large dams characterized by a wide operational discretion space and aimed at maximizing hydropower production is expected to benefit more from the information from seasonal forecasts than smaller dams serving irrigated agriculture. In particular, the same hydropower production levels of a forecast uninformed dam can be generated with a 20% reduction in capital costs (i.e. 20% smaller dam). Dam size being the same, forecasts allow to increase of 60 GWh/yr hydropower production in large dam, corresponding to more than 25% of the yearly average electricity consumption by the agriculture sector in Zambia, where Kariba dam is located (IEA, 2019). This hydropower improvement is attained at no additional cost for irrigation, as the irrigation deficit remains unchanged.
Extrapolating these figures to the new planned dams (Mupata 1200 MW, Mhpanda 1350MW, Batoka 1600 MW), which will cumulatively add more than double the installed power in Kariba (1830 MW), might increase this rate to 75 %. In the end, when tested over realistic streamflow forecasts characterized by different biases, the Kariba reservoir system has proven to be more sensitive to an over-estimation in the seasonal streamflow forecasts used to inform the infrastructure design, leading to an 8% maximum loss in the value of information.

Given the positive outcomes attained using perfect streamflow forecasts, a future research direction will be to test the effectiveness of our contribution using an existing forecast system, such as the Multi-model and Multi-product seasonal hydrological streamflow forecasting platform available for the Upper Zambezi (Roy et al., 2017; SERVIR Water Africa-Arizona Team (SWAAT), 2019) or the seasonal forecasts global system (SMHI, 2019). Our work can also be extended to other catchments located in different hydro-climatic regions and characterized by different planning and management challenges, in order to assess further interrelations between dam sizes, operational trade-offs, and forecast value. In the end, our study could be repeated with multiple, stochastic streamflow realizations in order to design reservoir systems that are less vulnerable to the intrinsic, stationary variability of the hydrologic processes.

6 Acknowledgments

This work was partially supported by the DAFNEDecision Analytic Framework to explore the waterenergyfood Nexus in complex transboundary water resource systems of fast developing countries research project funded by the Horizon 2020 programme WATER 2015 of the European Union, GA 690268. All the data used in this study are from the Zambezi River Authority (ZRA) and were collected during the DAFNE project (http://dafne.ethz.ch/). Because the model and all the figures contain sensitive information on hydropower plant, demand, and streamflow that is protected by a nondisclosure agreement with the ZRA authority, it cannot be made public.

Acknowledgments

This work was partially supported by the *DAFNE - Decision Analytic Framework to explore the water-energy-food Nexus in complex transboundary water resource systems of fast developing countries* research project funded by the Horizon 2020 programme WATER 2015 of the European Union, GA no. 690268. All the data used in this study are from the Zambezi River Authority (ZRA) and were collected during the DAFNE project (`http://dafne.ethz.ch/`). Because the model contains sensitive information on hydropower plant, demand, and streamflow that is protected by a nondisclosure agreement with the ZRA authority, it cannot be made public.

References

Afshar, A., Masoumi, F., & Solis, S. S. (2015). Reliability based optimum reservoir design by hybrid ACO-LP algorithm. *Water Resources Management*, *29*(6), 2045–2058.

Afzali, R., Mousavi, S. J., & Ghaheri, A. (2008). Reliability-based simulation-optimization model for multireservoir hydropower systems operations: Khersan experience. *Journal of Water Resources Planning and Management*, *134*(1), 24–33.

Anghileri, D., Voisin, N., Castelletti, A., Pianosi, F., Nijssen, B., & Lettenmaier, D. P. (2016). Value of long-term streamflow forecasts to reservoir operations for water supply in snow-dominated river catchments. *Water Resources Research*, *52*(6), 4209–4225.

Baker, S. A., Wood, A. W., & Rajagopalan, B. (2019). Developing subseasonal to seasonal climate forecast products for hydrology and water management. *JAWRA Journal of the American Water Resources Association*, *55*(4), 1024–1037.

Bellman, R. (1957). *Dynamic programming.* Princeton University Press.

Bertoni, F., Castelletti, A., Giuliani, M., & Reed, P. M. (2019). Discovering dependencies, trade-offs, and robustness in joint dam design and operation: An ex-post assessment of the Kariba dam. *Earth's Future*, *7*, 1367 1390.

Bertsekas, D. (1976). *Dynamic programming and stochastic control.* Academic Press, New York.

Block, P. (2011). Tailoring seasonal climate forecasts for hydropower operations. *Hydrology and Earth System Sciences*, *15*(4), 1355–1368.

Cassagnole, M., Ramos, M.-H., Thirel, G., Gailhard, J., & Garçon, R. (2017). Is the economic value of hydrological forecasts related to their quality? Case study

of the hydropower sector. In *EGU General Assembly Conference Abstracts* (Vol. 19, p. 9073).
Celeste, A. B., & Billib, M. (2009). Evaluation of stochastic reservoir operation optimization models. *Advances in Water Resources*, *32*(9), 1429–1443.
Cervigni, R., Liden, R., Neumann, J. E., & Strzepek, K. M. (2015). *Enhancing the climate resilience of Africa's infrastructure: The power and water sectors.* The World Bank.
Choudhury, D., Mehrotra, R., Sharma, A., Sen Gupta, A., & Sivakumar, B. (2019). Effectiveness of CMIP5 decadal experiments for interannual rainfall prediction over Australia. *Water Resources Research*.
Coello Coello, C. A., Lamont, G. B., Van Veldhuizen, D. A., et al. (2007). *Evolutionary algorithms for solving multi-objective problems* (Vol. 5). Springer.
Denaro, S., Anghileri, D., Giuliani, M., & Castelletti, A. (2017). Informing the operations of water reservoirs over multiple temporal scales by direct use of hydro-meteorological data. *Advances in Water Resources*, *103*, 51–63.
Doblas-Reyes, F., Balmaseda, M., Weisheimer, A., & Palmer, T. (2011). Decadal climate prediction with the European Centre for Medium-range Weather Forecasts coupled forecast system: Impact of ocean observations. *Journal of Geophysical Research: Atmospheres*, *116*(D19).
Eastman, J., & ReVelle, C. (1973). Linear decision rule in reservoir management and design: 3. Direct capacity determination and intraseasonal constraints. *Water Resources Research*, *9*(1), 29–42.
Faber, B. A., & Stedinger, J. (2001). Reservoir optimization using sampling sdp with ensemble streamflow prediction (esp) forecasts. *Journal of Hydrology*, *249*(1-4), 113–133.
Franz, K. J., Hartmann, H. C., Sorooshian, S., & Bales, R. (2003). Verification of National Weather Service ensemble streamflow predictions for water supply forecasting in the Colorado River basin. *Journal of Hydrometeorology*, *4*(6), 1105–1118.
Fuchs, M., Kling, H., Wipplinger, B., & Faber, R. (2018). Inflow forecasting for improved dam safety management on the Niger river. In *Twenty-sixth international congress on large dams.* CRC Press.
Galelli, S., & Castelletti, A. (2013a). Assessing the predictive capability of randomized tree-based ensembles in streamflow modelling. *Hydrology and Earth System Sciences*, *17*(7), 2669–2684.
Galelli, S., & Castelletti, A. (2013b). Tree-based iterative input variable selection for hydrological modeling. *Water Resources Research*, *49*(7), 4295–4310.
Galelli, S., Humphrey, G. B., Maier, H. R., Castelletti, A., Dandy, G. C., & Gibbs, M. S. (2014). An evaluation framework for input variable selection algorithms for environmental data-driven models. *Environmental Modelling & Software*, *62*, 33–51.
Gandolfi, C., & Togni, D. (1997). Optimal flow allocation in the Zambezi River system. *Water Resources Management*, *11*, 377–393.
Georgakakos, A., Yao, H., Kistenmacher, M., Georgakakos, K., Graham, N., Cheng, F.-Y., ... Shamir, E. (2012). Value of adaptive water resources management in northern california under climatic variability and change: Reservoir management. *Journal of Hydrology*, *412*, 34–46.
Georgakakos, K., Graham, N., Carpenter, T., & Yao, H. (2005). Integrating climate-hydrology forecasts and multi-objective reservoir management for Northern California. *EOS, Transactions American Geophysical Union*, *86*(12), 122–127.
Geressu, R. T., & Harou, J. J. (2015). Screening reservoir systems by considering the efficient trade-offsinforming infrastructure investment decisions on the Blue Nile. *Environmental Research Letters*, *10*(12), 125008.
Geurts, P., Ernst, D., & Wehenkel, L. (2006). Extremely randomized trees. *Machine Learning*, *63*(1), 3–42.

Giuliani, M., Castelletti, A., Pianosi, F., Mason, E., & Reed, P. M. (2016). Curses, tradeoffs, and scalable management: Advancing evolutionary multiobjective direct policy search to improve water reservoir operations. *Journal of Water Resources Planning and Management*, *142*(2), 04015050.

Giuliani, M., Pianosi, F., & Castelletti, A. (2015). Making the most of data: An information selection and assessment framework to improve water systems operations. *Water Resources Research*, *51*(11), 9073–9093.

Giuliani, M., Quinn, J. D., Herman, J. D., Castelletti, A., & Reed, P. M. (2018). Scalable multiobjective control for large-scale water resources systems under uncertainty. *IEEE Transactions on Control Systems Technology*, *26*(4), 1492–1499.

Graham, N. E., & Georgakakos, K. P. (2010). Toward understanding the value of climate information for multiobjective reservoir management under present and future climate and demand scenarios. *Journal of Applied Meteorology and Climatology*, *49*(4), 557–573.

Guariso, G., Rinaldi, S., & Soncini-Sessa, R. (1986). The management of Lake Como: A multiobjective analysis. *Water Resources Research*, *22*(2), 109–120.

Hadka, D., & Reed, P. (2013). Borg: An auto-adaptive many-objective evolutionary computing framework. *Evolutionary Computation*, *21*(2), 231–259.

Hadka, D., & Reed, P. (2015). Large-scale parallelization of the borg multiobjective evolutionary algorithm to enhance the management of complex environmental systems. *Environmental Modelling & Software*, *69*, 353–369.

Hall, W., Askew, A., Yeh, W., et al. (1969). Use of the critical period in reservoir analysis. *Water Resources Research*, *5*(6), 1205–1215.

Ham, Y.-G., Rienecker, M. M., Suarez, M. J., Vikhliaev, Y., Zhao, B., Marshak, J., ... Schubert, S. D. (2014). Decadal prediction skill in the GEOS-5 forecast system. *Climate Dynamics*, *42*(1-2), 1–20.

Hamlet, A. F., Huppert, D., & Lettenmaier, D. P. (2002). Economic value of long-lead streamflow forecasts for Columbia River hydropower. *Journal of Water Resources Planning and Management*, *128*(2), 91–101.

Hejazi, M. I., Cai, X., & Ruddell, B. L. (2008). The role of hydrologic information in reservoir operation – Learning from historical releases. *Advances in Water Resources*, *31*(12), 1636–1650.

Houck, M. H., Cohon, J. L., & ReVelle, C. S. (1980). Linear decision rule in reservoir design and management: 6. incorporation of economic efficiency benefits and hydroelectric energy generation. *Water Resources Research*, *16*(1), 196–200.

IEA. (2019). *World Energy Balances.* Retrieved from `https://webstore.iea.org/world-energy-balances-2019`

IRENA. (2012). *Renewable power generation costs in 2012: An overview* (Tech. Rep.). IRENA.

Jeuland, M., & Whittington, D. (2014). Water resources planning under climate change: Assessing the robustness of real options for the Blue Nile. *Water Resources Research*, *50*(3), 2086–2107.

Kelman, J., Stedinger, J. R., Cooper, L. A., Hsu, E., & Yuan, S.-Q. (1990). Sampling stochastic dynamic programming applied to reservoir operation. *Water Resources Research*, *26*(3), 447–454.

Kim, Y.-O., & Palmer, R. N. (1997). Value of seasonal flow forecasts in bayesian stochastic programming. *Journal of Water Resources Planning and Management*, *123*(6), 327–335.

Koutsoyiannis, D., & Economou, A. (2003). Evaluation of the parameterization-simulation-optimization approach for the control of reservoir systems. *Water Resources Research*, *39*(6).

Lall, U., & Miller, C. W. (1988). An optimization model for screening multipurpose reservoir systems. *Water Resources Research*, *24*(7), 953–968.

Lu, M., Lall, U., Robertson, A. W., & Cook, E. (2017). Optimizing multiple reliable forward contracts for reservoir allocation using multitime scale streamflow forecasts. *Water Resources Research*, *53*(3), 2035–2050.

Matrosov, E. S., Huskova, I., Kasprzyk, J. R., Harou, J. J., Lambert, C., & Reed, P. M. (2015). Many-objective optimization and visual analytics reveal key trade-offs for londons water supply. *Journal of Hydrology*, *531*, 1040–1053.

Maurer, E. P., & Lettenmaier, D. P. (2004). Potential effects of long-lead hydrologic predictability on Missouri River main-stem reservoirs. *Journal of Climate*, *17*(1), 174–186.

Meehl, G. A., Goddard, L., Murphy, J., Stouffer, R. J., Boer, G., Danabasoglu, G., . . . others (2009). Decadal prediction: Can it be skillful? *Bulletin of the American Meteorological Society*, *90*(10), 1467–1486.

Montaseri, M., & Adeloye, A. (1999). Critical period of reservoir systems for planning purposes. *Journal of Hydrology*, *224*(3-4), 115–136.

Moriasi, D. N., Arnold, J. G., Van Liew, M. W., Bingner, R. L., Harmel, R. D., & Veith, T. L. (2007). Model evaluation guidelines for systematic quantification of accuracy in watershed simulations. *Transactions of the ASABE*, *50*(3), 885–900.

Mousavi, H., & Ramamurthy, A. (2000). Optimal design of multi-reservoir systems for water supply. *Advances in Water Resources*, *23*(6), 613–624.

Murphy, A. H. (1993). What is a good forecast? an essay on the nature of goodness in weather forecasting. *Weather and Forecasting*, *8*(2), 281–293.

Nayak, M. A., Herman, J. D., & Steinschneider, S. (2018). Balancing flood risk and water supply in California: Policy search integrating short-term forecast ensembles with conjunctive use. *Water Resources Research*, *54*(10), 7557–7576.

Oliveira, R., & Loucks, D. P. (1997). Operating rules for multireservoir systems. *Water Resources Research*, *33*(4), 839–852.

Owens, R. G., & Hewson, T. D. (2018). *ECMWF Forecast user guide*. Reading: ECMWF.

Paseka, S., Kapelan, Z., & Marton, D. (2018). Multi-objective optimization of resilient design of the multipurpose reservoir in conditions of uncertain climate change. *Water*, *10*(9), 1110.

Perelman, L., Housh, M., & Ostfeld, A. (2013). Robust optimization for water distribution systems least cost design. *Water Resources Research*, *49*(10), 6795–6809.

Quinn, J. D., Reed, P. M., Giuliani, M., Castelletti, A., Oyler, J. W., & Nicholas, R. E. (2018). Exploring how changing monsoonal dynamics and human pressures challenge multireservoir management for flood protection, hydropower production, and agricultural water supply. *Water Resources Research*, *54*(7), 4638–4662.

Raso, L., Schwanenberg, D., van de Giesen, N., & van Overloop, P. (2014). Short-term optimal operation of water systems using ensemble forecasts. *Advances in Water Resources*, *71*, 200–208.

Reed, P. M., Hadka, D., Herman, J. D., Kasprzyk, J. R., & Kollat, J. B. (2013). Evolutionary multiobjective optimization in water resources: The past, present, and future. *Advances in Water Resources*, *51*, 438–456.

Rodriguez, R. A., Becker, S., & Greiner, M. (2015). Cost-optimal design of a simplified, highly renewable pan-European electricity system. *Energy*, *83*, 658–668.

Roy, T., Serrat-Capdevila, A., Gupta, H., & Valdes, J. (2017). A platform for probabilistic Multimodel and Multiproduct Streamflow Forecasting. *Water Resources Research*, *53*(1), 376–399.

Saaty, T., & Gass, S. (1954). Parametric objective function (part 1). *Journal of the Operations Research Society of America*, *2*(3), 316–319.

Saavedra Valeriano, O. C., Koike, T., Yang, K., Graf, T., Li, X., Wang, L., & Han, X. (2010). Decision support for dam release during floods using a distributed biosphere hydrological model driven by quantitative precipitation forecasts. *Water Resources Research*, *46*(10).

SADC. (2012). *Zambezi River Basin Atlas of the Changing Environment.* Southern African Development Community (SADC), Southern African Research and Documentation Centre (SARDC), United Nations Environment Programme (UNEP), GRID-Arendal, Zambezi Watercourse Commission (ZAMCOM). Gaborone, Harare and Arendal.

Salazar, J. Z., Reed, P. M., Quinn, J. D., Giuliani, M., & Castelletti, A. (2017). Balancing exploration, uncertainty and computational demands in many objective reservoir optimization. *Advances in Water Resources*, *109*, 196–210.

Sankarasubramanian, A., Lall, U., Souza Filho, F. A., & Sharma, A. (2009). Improved water allocation utilizing probabilistic climate forecasts: Short-term water contracts in a risk management framework. *Water Resources Research*, *45*(11).

Satishkumar, U., Balakrishnan, P., Ramaswamy, K., et al. (2010). Application of chance-constrained linear programming model for reservoir sizing of a cascade system of tanks in northern karnataka. *Karnataka Journal of Agricultural Sciences*, *23*(2), 286–292.

Schuster, M., Grieger, J., Richling, A., Schartner, T., Illing, S., Kadow, C., ... Ulbrich, U. (2019). Improvement in the decadal prediction skill of the North Atlantic extratropical winter circulation through increased model resolution. *Earth System Dynamics*, *10*, 901–917.

SERVIR Water Africa-Arizona Team (SWAAT). (2019). *Streamflow forecasts in African basins.* Retrieved from `http://www.swaat.arizona.edu/index.html`

Shourian, M., Mousavi, S., & Tahershamsi, A. (2008). Basin-wide water resources planning by integrating PSO algorithm and MODSIM. *Water Resources Management*, *22*(10), 1347–1366.

SMHI. (2019). *SMHI Hypeweb - Seasonal forecasts global.* Retrieved from `https://hypeweb.smhi.se/explore-water/forecasts/seasonal-forecasts-global/`

Smith, D., Eade, R., Scaife, A., Caron, L.-P., Danabasoglu, G., DelSole, T., ... others (2019). Robust skill of decadal climate predictions. *npj Climate and Atmospheric Science*, *2*(1), 1–10.

Soils Incorporated Ltd. (2000). *Kariba dam case study* (Tech. Rep.). Prepared as input to the World Commission on Dams, Cape Town.

Soncini-Sessa, R., Weber, E., & Castelletti, A. (2007). *Integrated and participatory water resources management-theory.* Elsevier.

Sreekanth, J., Datta, B., & Mohapatra, P. K. (2012). Optimal short-term reservoir operation with integrated long-term goals. *Water Resources Management*, *26*(10), 2833–2850.

Steinschneider, S., & Brown, C. (2012). Dynamic reservoir management with real-option risk hedging as a robust adaptation to nonstationary climate. *Water Resources Research*, *48*(5).

Stephenson, D., & Petersen, M. S. (1991). Water resources development in developing countries. Chapter 7: Reservoir sizing. *Developments in Water Science*, *41*, 130-160.

Thomas Jr, H. A., & Burden, R. P. (1963). *Operations research in water quality management* (Tech. Rep.). Division of Engineering and Applied Physics, Harvard University: Cambridge, Massachusetts.

Turner, S. W., Bennett, J. C., Robertson, D. E., & Galelli, S. (2017). Complex relationship between seasonal streamflow forecast skill and value in reservoir operations. *Hydrology and Earth System Sciences*, *21*(9), 4841–4859.

U.S. Army Corps of Engineers. (1975). *Hydrologic Engineering Methods for Water Resources Development: Reservoir Yield, vol. 8.* Davis, CA: Hydrologic Engineering Centre.

U.S. Army Corps of Engineers. (1977). *Hydrologic Engineering Methods for Water Resources Development: Reservoir System Analysis for Conservation, vol. 9.* Davis, CA: Hydrologic Engineering Centre.

Voisin, N., Hamlet, A. F., Graham, L. P., Pierce, D. W., Barnett, T. P., & Lettenmaier, D. P. (2006). The role of climate forecasts in western US power planning. *Journal of Applied Meteorology and Climatology*, *45*(5), 653–673.

Wang, F., Wang, L., Zhou, H., Saavedra Valeriano, O. C., Koike, T., & Li, W. (2012). Ensemble hydrological prediction-based real-time optimization of a multiobjective reservoir during flood season in a semiarid basin with global numerical weather predictions. *Water Resources Research*, *48*(7).

World Bank. (2010). *The Zambezi River Basin-A Multi-sector Investment Opportunities Analysis.* The World Bank Washington, DC.

Xu, B., Zhong, P.-A., Stanko, Z., Zhao, Y., & Yeh, W. W.-G. (2015). A multiobjective short-term optimal operation model for a cascade system of reservoirs considering the impact on long-term energy production. *Water Resources Research*, *51*(5), 3353–3369.

You, J.-Y., & Cai, X. (2008). Determining forecast and decision horizons for reservoir operations under hedging policies. *Water Resources Research*, *44*(11).

Zhao, T., Zhao, J., Lund, J. R., & Yang, D. (2014). Optimal hedging rules for reservoir flood operation from forecast uncertainties. *Journal of Water Resources Planning and Management*, *140*(12), 04014041.

Zitzler, E., Thiele, L., Laumanns, M., Fonseca, C. M., & da Fonseca, V. G. (2003). Performance assessment of multiobjective optimizers: An analysis and review. *IEEE Transactions on Evolutionary Computation*, *7*(2), 117-132.