Water Resources Research

RESEARCH ARTICLE
10.1029/2020WR028112

Key Points:

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

Supporting Information:

- Supporting Information S1

Correspondence to:
A. Castelletti,
andrea.castelletti@polimi.it

Citation:
Bertoni, F., Giuliani, M., Castelletti, A., & Reed, P. M. (2021). Designing with information feedbacks: Forecast informed reservoir sizing and operation. *Water Resources Research*, *57*, e2020WR028112. https://doi.org/10.1029/2020WR028112

Received 9 JUN 2020
Accepted 24 JAN 2021

Designing With Information Feedbacks: Forecast Informed Reservoir Sizing and Operation

F. Bertoni[1], M. Giuliani[1], A. Castelletti[1], and 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

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 supporting the implementation of flexible operating policies informed by streamflow forecasts to enable the design of less costly reservoirs relative to alternatives that do not rely on forecast information. Our approach initially explores the maximum potential gain attainable by searching and using the most valuable forecast information and lead time. We then analyze the results? sensitivities relative to existing and synthetic biased forecasts. 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 enables attaining the same hydropower production of the existing dam, while reducing infrastructure size and associated capital costs by 20%. A forecast-informed operation of the existing system can instead facilitate an annual average increase of 60 GWh in hydropower production. This finding, extrapolated to the new planned dams in the basin, suggests that forecast informed policies could yield power production benefits equal to 75% of the current annual electricity consumption of the Zambian agricultural sector. The use of biased forecasts substantially reduces this gain, showing that the ESP forecasts value is marginal and that informed infrastructure designs are particularly vulnerable to forecast overestimation. Advancing information feedbacks may therefore 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, the traditional engineering approach (e.g., U.S. Army Corps of Engineers, 1975, 1977) relies on the Rippl method to identify the optimal dam size based on a sequence of prespecified, desired releases or simple abstractions of predefined operating policies (e.g., Hall et al., 1969; Montaseri & Adeloye, 1999; Stephenson & Petersen, 1991). More recent examples (e.g., Bertoni et al., 2019; Geressu & Harou, 2015) show that the joint design of reservoir size and operation can benefit from the use of state aware control actions that are dynamic and adaptive to the system conditions being observed. However, these operating rules are generally conditioned on the reservoir level/storage only, while they do not account for other available information feedbacks, such as streamflow forecasts, which support more flexible and adaptive operations (e.g., Giuliani et al., 2019; Libisch-Lehner et al., 2019; Nayak et al., 2018).

The value of employing forecasts to enhance the operations of existing infrastructures has long been acknowledged (e.g., Faber & Stedinger, 2001; Kelman et al., 1990; Kim & Palmer, 1997). 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 (Denaro et al., 2017). In a water reservoir system primarily operated to satisfy short-term operating objectives (e.g., flood control), short-term forecasts are likely the most informative because they provide the system operator with anticipation capacity to create a buffer storage for mitigating the upcoming flood peak and thus minimizing flood damages (e.g., Raso et al., 2014; Saavedra Valeriano et al., 2010; Wang et al., 2012; Zhao et al., 2014). Alternatively, reservoirs operated with respect to long-term objectives (e.g., irrigation water supply) might primarily benefit from seasonal forecasts (e.g., Anghileri et al., 2016; Block, 2011; K. Georgakakos et al., 2005; Hamlet et al., 2002; Maurer & Lettenmaier, 2004;

Steinschneider & Brown, 2012; Voisin et al., 2006). When dealing with a multipurpose 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., Denaro et al., 2017; Fuchs et al., 2018; A. Georgakakos et al., 2012; Lu et al., 2017; Nayak et al., 2018; Sreekanth et al., 2012; Xu et al., 2015).

Prior studies have also shown that forecast value can change with dam size (e.g., Anghileri et al., 2016; Graham & Georgakakos, 2010; Sankarasubramanian et al., 2009; Turner et al., 2017; You & Cai, 2008). For example, the water supply operations of both undersized and oversized dams may become trivial, since in the former case the reservoir operation cannot cancel the structural deficit of the system, whereas in oversized context the operators are generally able to always 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 few 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 exploration of the coupled planning/operation design space. Given the complex interdependencies between reservoir dynamics, information feedbacks, and infrastructure design, the discrete problem decompositions tacit to these sensitivity analyses may bias the resulting forecast value estimates.

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 valuable forecast information and lead times for perfect streamflow forecasts across different dam sizes and operational trade-offs. Then, the selected forecasts are used 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 more realistic forecasts characterized by different biases. We demonstrate the value of our methodological contribution through an ex post design analysis of the Kariba Dam in the Zambezi River Basin (ZRB), 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 interannual carry over water volumes that could benefit from streamflow forecasts to mitigate seasonal and interannual drought anomalies.

In summary, the main contributions of the paper include a methodology for (1) assessing the maximum potential gain generated by informing infrastructure design with perfect streamflow forecasts; (2) selecting the most informative forecast information and lead times for different dam sizes and operational trade-offs; (3) understanding the sensitivity of the forecast-informed infrastructure designs on specific forecast biases and providing recommendations for improving existing forecast systems.

2. Case Study Description

2.1. Kariba Dam

The Kariba reservoir is a regulated lake in the transboundary ZRB in southeastern Africa (Figure 1a). Draining 1.37 million km^2, the ZRB is the fourth 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 under the supervision of the Zambezi Watercourse Commission (ZAMCON), an international river basin authority established in 2014 with the goal of promoting an equitable and reasonable utilization of the water resources of the Zambezi Watercourse, including both long-term planning and operational decisions. The Kariba Dam system is complemented by two irrigation districts, located respectively upstream and downstream the reservoir, where major cultivated crops are sugar cane, rice, wheat, and maize

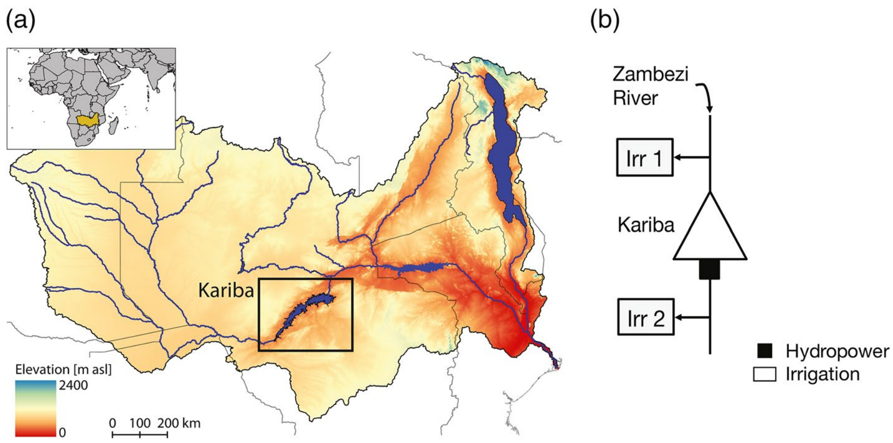

Figure 1. Panel (a) map of the Zambezi River Basin with Kariba Dam squared in black. Panel (b) schematic representation of the Kariba multipurpose reservoir system.

(Payet-Burin et al., 2019). Further downstream, the Zambezi River flows in Mozambique and is dammed at Cahora Bassa, which has an installed capacity equal to 2,075 MW. Two other dams, Ithezi-Thezi and Kafue Gorge, are instead located in Zambia on the tributary Kafue River and contribute an additional installed capacity of 1,100 MW.

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 nonlinear 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 nonlinear, 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 $\overline{u}_t(\alpha)$ and minimum $\underline{u}_t(\alpha)$ feasible release functions, due to the presence of physical (e.g., spillway activation) constraints. Such release functions directly depend upon the dam size $\alpha \in A$, so 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, and streamflow forecasts). The policy $p_{\theta_{res}}$ belongs to a predefined class of functions, according to which it is parameterized within the space of the parameters $\theta_{\mathrm{res}} \in \Theta_{\mathrm{res}}(\alpha)$. Note that the interdependency between dam size and operation is expressed in terms of the direct dependence of the feasibility set $\Theta_{\mathrm{res}}(\alpha)$ 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.

10.1029/2020WR028112

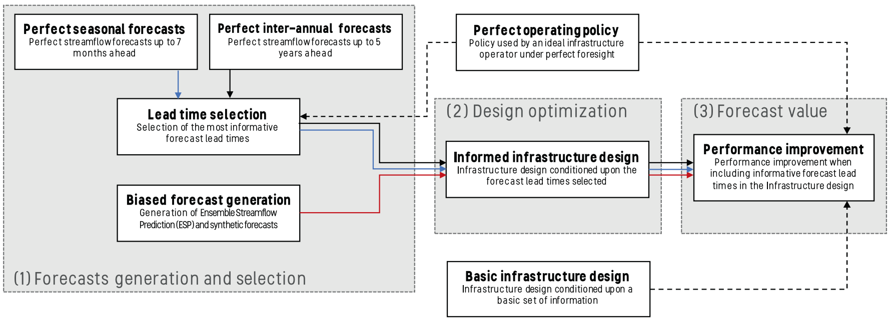

Figure 2. Flowchart of the methodology employed in this study. Each line is colored differently based on the set of streamflow forecasts it refers to, namely perfect seasonal (blue), perfect interannual (black), and realistic seasonal (red) forecasts.

As for the two irrigation districts (id = 1,2), they can abstract water a_{t+1}^{id} from the river through regulated diversion channels. The total demand of the two district is larger than the mean annual inflow entering the system, which generates a structural irrigation deficit as in average conditions there is no enough water to cover the demand regardless of the adopted system operation. The volume of water they can abstract is calculated according to a nonlinear hedging rule (Celeste & Billib, 2009), which allows diverting into the irrigation canals less water than their corresponding demands in order to account for downstream users:

$$a_{t+1}^{id} = \begin{cases} \min\left(q_{t+1}^{id}, w_t^{id} \cdot \left[\frac{q_{t+1}^{id}}{h^{id}}\right]^{m^{id}}\right) & \text{if } q_{t+1}^{id} \leq h^{id} \\ \min\left(q_{t+1}^{id}, w_t^{id}\right) & \text{else} \end{cases} \tag{2}$$

where q_{t+1}^{id} is the volume of water available in the river upstream of the id-th irrigation district, and w_t^{id} is the monthly water demand taken from World Bank (2010). 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 $\theta_{\text{irr}} = [h^1, m^1, h^2, m^2] \in \Theta_{\text{irr}}$. These parameters are optimized along with reservoir size and its operations (see Equation 3).

3. Methods and Tools

Our methodology, illustrated in Figure 2, is composed of the following three methodological steps, which are detailed in the next subsections.

The first step is the generation of streamflow forecasts and the selection of the most informative forecast information and lead times to be included within the dam design phase (see Section 3.1). Three different sets of streamflow forecasts are considered, perfect seasonal (blue arrows), perfect interannual (black arrows), and biased seasonal (red arrows) forecasts. The perfect seasonal and interannual forecasts are used to evaluate the upper bound theoretical value of information at these timescales. The biased seasonal forecasts are intended to reproduce a more realistic decision making environment, and to disentangle the influence of different forecast biases on the associated informed infrastructure designs. Then, we select the most informative seasonal and interannual 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 future hydrologic conditions (i.e., future reservoir inflows).

The second step of our methodology consists of a joint optimization of reservoir size and operations (see Section 3.2). In particular, we identify the informed infrastructure design (IID), where operations are informed by the forecast information selected in the previous step. The resulting 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 procedure is the estimation of the forecast value, namely the performance improvement that could be attained when including the selected information feedbacks in the infrastructure design (see Section 3.3). 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, which is computed for all the considered forecasts, namely perfect seasonal, perfect interannual, and biased seasonal predictions.

3.1. Generation and Selection of Forecasts

The generation and selection of forecasts step consists of first identifying two arrays of perfect streamflow forecasts at both seasonal and interannual time scale, from which the most valuable forecast information and lead times are selected to inform the search for design and operation alternatives. The resulting solutions represent an upper limit of the system performance and allows understanding if the infrastructure design problem can benefit from forecast information. To assess the sensitivity of the informed infrastructure design on more realistic forecasts, we then generate a set of biased seasonal streamflow forecasts including ensemble streamflow prediction (ESP; Day, 1985) as well as synthetic forecasts with different levels of overestimation, underestimation, and underdispersion (i.e., underestimation of high flows and overestimation of low flows) (Cassagnole et al., 2020). The comparison of the results obtained with biased forecasts with the reference provided by the (ideal) solutions relying on perfect forecasts will also provide recommendations for improving existing forecast systems in order to make them suitable for informing infrastructure design. For details about the biased forecast generation, see Section S7 of the supporting information.

The selection of the most valuable forecast information is however not a straightforward process. Since forecast skill usually degrades with long lead times (Doblas-Reyes et al., 2011), shorter and more accurate forecast lead times are more often used. Yet, the operational value of forecasts is strictly related to the structural characteristics of the system considered (e.g., reservoir size) and large dams with an annual carry over capacity might also benefit from longer, less precise forecast lead times to further enhance their operations.

In order to support this information selection problem, following the Information Selection and Assessment (ISA) procedure proposed by Giuliani et al. (2015) we use an input variable selection technique. We initially 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 of existing seasonal forecast systems, including the European Center for Medium-Range Weather Forecasts (ECMWF) SEAS5 (Owens & Hewson, 2018) and the NOAA's National Weather Service (Franz et al., 2003). Then, 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 (Galelli & Castelletti, 2013a; Geurts et al., 2006). The IIS algorithm is a hybrid model-based/model-free approach characterized by modeling flexibility (i.e., ability to approximate strongly nonlinear 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 $\mathbf{I}_t^s \in \Xi_t^s$ 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 the simulation of an ideal operating policy under the assumption of 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 of the supporting information.

Interannual and decadal time frames are particularly relevant for infrastructure planners and water resources managers, potentially bringing added socioeconomic benefits to infrastructure operations (Choudhury et al., 2019). Decadal forecasts aim at modeling future climatic conditions over a longer time horizon (i.e., the next 10–30 years) and their potential value has being recently investigated in the literature (e.g., Ham et al., 2014; Schuster et al., 2019; Smith et al., 2019). As a means of theoretically bounding the value of interannual 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 interannual anomalies related to global climate oscillations (e.g., El Niño Southern Oscillation). When informing their operations with interannual 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 interannual forecasts, in order to identify their most informative lead times $\mathbf{I}_t^i \in \Xi_t^i$ for different dam sizes.

3.2. Design Optimization

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

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

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_π^{cost} and two management J_π^{hyd}, J_π^{irr} objectives, formulated as follows:

- Minimization of dam construction costs J_π^{cost} [\$] discounted over the lifespan of the project

$$J_\pi^{cost} = c(\alpha) \cdot \frac{r}{1 - (1 + r)^{-L}} \tag{4}$$

  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_π^{hyd} [TWh/yr]

$$J_\pi^{hyd} = \frac{1}{H} \sum_{t=0}^{H-1} \delta g \gamma (\eta^N \overline{h}_t^N q_{t+1}^N + \eta^S \overline{h}_t^S q_{t+1}^S) \tag{5}$$

  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, $\gamma = 1{,}000$ [kg/m^3] is the water density, $\overline{h}_t^N$ and $\overline{h}_t^S$ [m] are the net hydraulic heads of the North and South Bank (i.e., reservoir level minus tailwater level), while q_{t+1}^N and q_{t+1}^S [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: $q_{t+1}^N = \min(r_{t+1}^N, \overline{q}^N)$, where $\overline{q}^N$ is the maximum capacity of the turbines at the North Bank and $r_{t+1}^N = \Delta \cdot r_{t+1}$ 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_π^{irr} [−] normalized with respect to the squared irrigation demand of each district

$$J_\pi^{irr} = \frac{1}{H} \sum_{t=0}^{H-1} \sum_{id=1}^{2} \left(\frac{\max(w_t^{id} - a_{t+1}^{id}, 0)}{w_t^{id}} \right)^2 \tag{6}$$

  where w_t^{id} and a_{t+1}^{id} are the monthly irrigation water demand and abstraction for the id-th irrigation district respectively.

At first, we identify the 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, q_t), consisting of the reservoir storage s_t (state of the system), the month of the year t, and the previous month inflow q_t, namely $u_t = p_{\theta_{res}}(s_t, t, q_t)$. By conditioning their operations on a minimum number of variables that can be observed at time t, the resulting basic infrastructure designs represent our lower bound system performance.

Then, we solve problem 3 to identify the 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, previous month inflow q_t, along with the forecast information $\mathbf{I}_t^f \in \Xi_t^f$ (with $f = [s, i]$) selected in the previous step of our procedure. 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 an enlarged set of policy input conditioning the resulting decision u_t, namely $u_t = p(s_t, t, q_t, \mathbf{I}_t^f)$.

We solve both the basic and informed infrastructure design problems by adopting the approach proposed in Bertoni et al. (2019), where evolutionary multiobjective direct policy search (EMODPS; Giuliani et al., 2016) is expanded to search for optimal reservoir sizing in addition to the optimal operations of both the reservoir and the two irrigation diversion channels. EMODPS indeed allows solving the multiobjective joint design of infrastructure size and operations in a single optimization run. This makes EMODPS preferable with respect to other approaches, including model predictive control, that can benefit from forecast information but would be computationally more demanding as they would require multiple optimization runs to explore different dam sizes in a nested approach (Bertoni et al., 2020) and different tradeoffs by using different scalarization values to reduce the dimensionality of the objective space (Soncini-Sessa et al., 2007). EMODPS is a parameterization-simulation-optimization approach (Guariso et al., 1986; Koutsoyiannis & Economou, 2003; Oliveira & Loucks, 1997) that searches candidate parameterized operating policies $p_{\theta_{res}}$ in the space of the parameters $\theta_{\mathrm{res}} \in \Theta_{\mathrm{res}}$ via multiobjective 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}. The coupled design of reservoir size and operation problem (see Equation 3) 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 as a nonlinear network of Gaussian radial basis functions (for further details about the mathematical formulation of the parameterized operating policy, refer to Section S1 of the supporting information).

The search tasks are 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 multiobjective 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 multimaster Borg MOEA (Hadka & Reed, 2015), which has proven to be successful in complex reservoir control problems (e.g., Giuliani et al., 2018; Quinn et al., 2018; Salazar et al., 2017). The multimaster 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).

In addition to the BID and IID formulations, we have to determine the target output of the IIS procedure described in Section 3.1. This is achieved by solving the management side of the joint optimization problem 3 with respect to the two management objectives (i.e., J_π^{hyd} in Equation 5, J_π^{irr} in Equation 6), for a fixed dam size $\overline{\alpha}$ and irrigation diversion parameters $\overline{\theta}_{irr}$. For further details about the mathematical formulation of the optimization problem, refer to Section S2 of the supporting information. 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 IIS 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, for further details, refer to Section S2 in the supplementary material). 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 multiobjective problems, the resolution of problem 3 yields a set of Pareto optimal or approximate alternatives. Following Zitzler et al. (2003), we use finite set theory to quantify the candidate space for multiobjective 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 EVPI is calculated as the difference in hypervolume between the ideal optimal Pareto front (i.e., the POP performance) and the approximation set attained using only information on the reservoir storage (i.e., the BID performance), where the Pareto front associated to the higher hypervolume is the better. As a rule, the larger the EVPI, the more the system can benefit from including more forecast information during its joint optimization of candidate designs and operations.

The EVPI is then 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 informed and 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 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

Both the BID and IID formulations are solved using the multimaster 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 runup to one million function evaluations, proved to be sufficient by visual inspection of search progress, with little variability across seeds (refer to Section S3 of the supporting information). The remainder of the multimaster 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 nondominated 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 E5-2,680 CPUs at 2.7 GHz, 128 GB of RAM, using 192 cores per island for a total of 3,840 computational hours.

4. Results and Discussion

4.1. Influence of Dam Size on the EVPI

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 EVPI. 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 (gray). For each dam size, the corresponding EVPI is calculated as the difference in hypervolume between basic and Perfect solutions (gray shaded area) and reported in the bottom right corner of each panel. The EVPI 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 intramonthly to interannually). 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 gray 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 (see Section 2.2), the information on the amount of water entering the reservoir in subsequent months is less valuable and 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 nine solutions.

4.2. Forecast Informed Infrastructure Design Using Perfect Seasonal Forecasts

After assessing the EVPI as well as the potential benefits that are achievable when hydropower strongly shapes solution preferences, we now identify via IIS the most informative seasonal forecasts $\mathbf{I}_t^s \in \Xi_t^s$ for the three target trade-offs for each of the three dam sizes (Figure 3). Independent of dam size and operational performance trade-offs, the most informative variables selected always provide information on future streamflow extremes (i.e., maximum or minimum future streamflows) 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 $\mathrm{qM7}_t$ for all three dam sizes, whereas the minimum streamflow over the same period $\mathrm{qm7}_t$ is selected in the case of the irrigation-prone solutions (more details on the results of this information selection phase can be found in Section S4 of the supporting information). We employ these two variables separately as the first, most informative input to be included in the informed infrastructure design phase in addition to reservoir storage and time. Results show that the

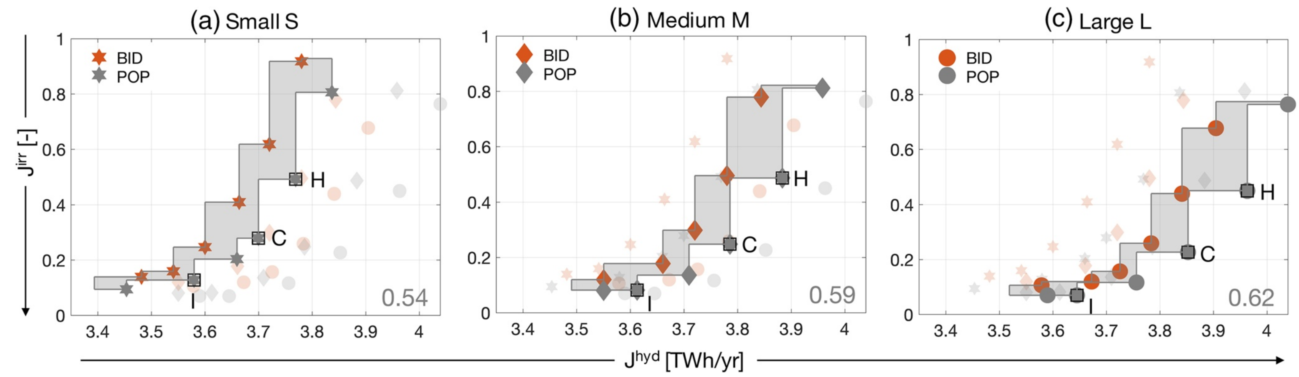

Figure 3. The objective space performance comparisons of the basic infrastructure designs (orange) and the corresponding perfect operating policies (gray) for each of the three dam sizes selected, namely small S (stars, panel (a)), medium M (diamonds, panel (b)), and large L (circles, panel (c)). The gray 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.

solutions informed with the maximum inflow outperform the ones with the minimum inflow (see Table S2 of the supporting information).

- *Perfect interannual streamflow forecasts*: in addition to the set of seasonal forecasts in Table 1, we also consider interannual 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 multiyear time resolution associated to interannual 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
- *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 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 of the supporting information). 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 POPs associated to three different target trade-offs between the two management objectives are used as target outputs of the IIS procedure to identify the most informative forecast lead times. Since the Pareto front extremes are not particularly interesting in a realistic decision making problem as they focus on a single objective only, we selected a hydropower-prone and irrigation-prone tradeoffs that attain a performance in hydropower production and irrigation deficit close to the corresponding extreme solutions while also partially accounting for the other objective
- *Information selection*: for each target POP trade-off to be explained and each dam size, the IIS algorithm is used to select the most informative forecast lead times that mostly explain the target sequence of optimal releases. At first, we perform a regression on a sample data set 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

Table 1
Set of Perfect Seasonal Streamflow Forecasts Calculated Over Different Lead Times

Name	Description	Period
$q1_t$,...,$q7_t$	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

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 interannual 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 interannual 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. 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

Figure4 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} (the full Pareto approximate set of IID is illustrated in Figure S4 of the supporting information). For each dam size, the forecast value is computed as the difference in hypervolume between basic and informed solutions (cyan shaded area), whereas the gray shaded area represents the residual space for improvement that is, not yet explained. Regardless of dam size, the informed system performs better than the basic, moving closer to the set of POPs.

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 gray shaded area shrinks accordingly. In particular, if we fix a specific high level of hydropower production (e.g., $J^{\text{hyd}} \simeq 3.84$ TWh/yr), a 20% reduction in capital costs could be attained by designing a medium dam operated with forecast information (i.e., cyan diamond

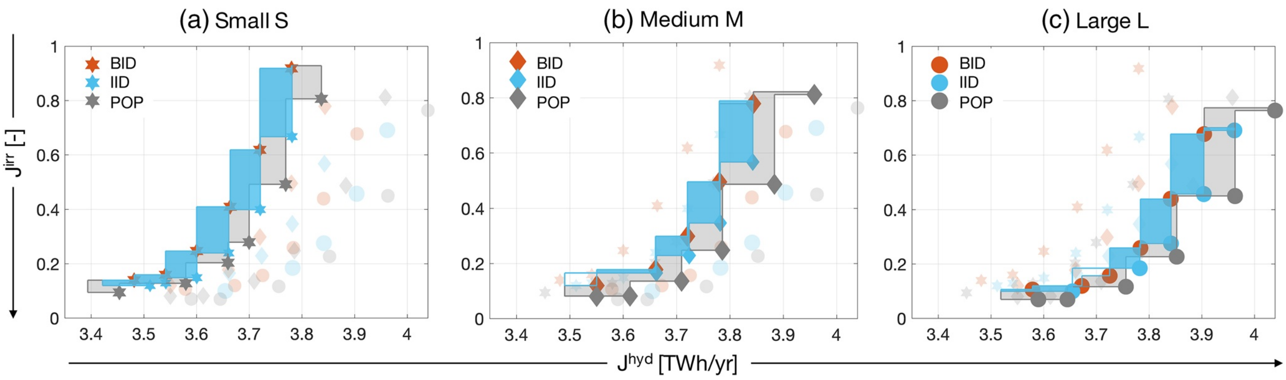

Figure 4. Comparison of the two objective tradeoffs that result from the basic infrastructure designs (orange), the informed infrastructure designs (cyan) and perfect operating policies (gray) for each of the three dam sizes selected, namely small S (stars, panel (a)), medium M (diamonds, panel (b)),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

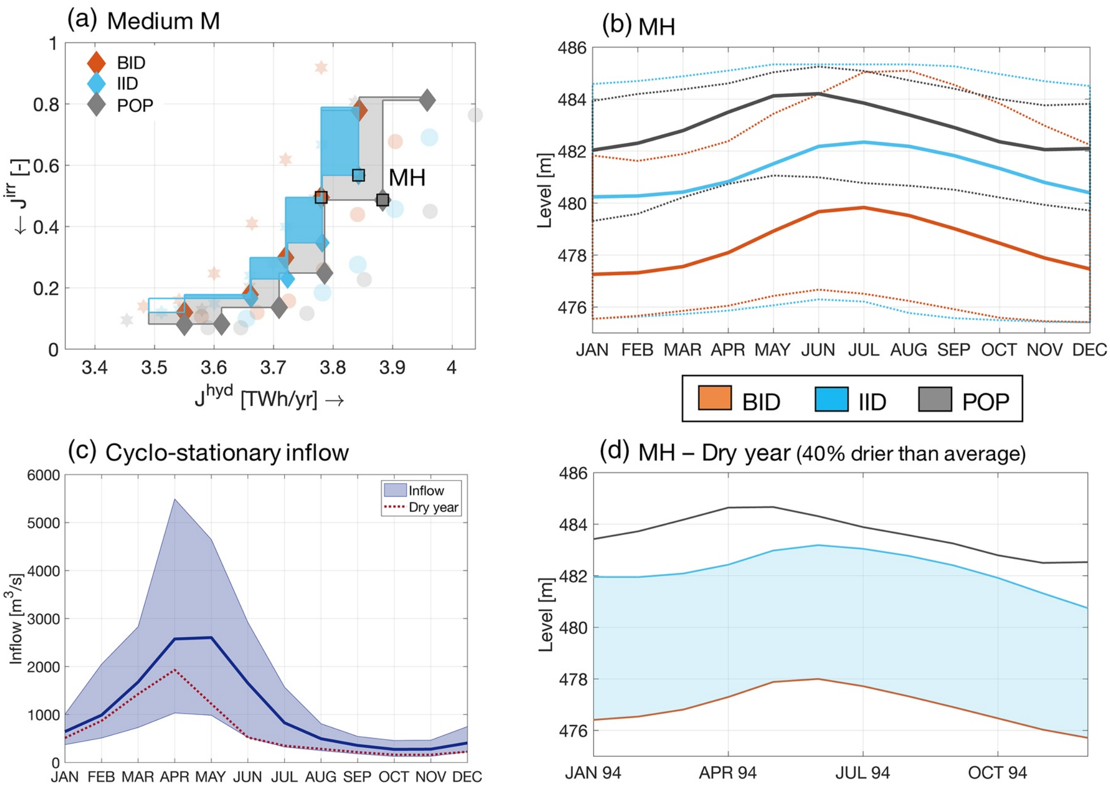

Figure 5. Panel (a) Two objective performance trade-offs for the medium M dam size, where the BID (orange), IID (cyan), and POP (gray) solutions associated to a hydropower-prone operating policy H are squared in black. Panel (b) monthly cyclo-stationary level trajectories for the three solutions highlighted 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.

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 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). It is also important to notice that this hydropower improvement is attained at no additional cost for irrigation, as the irrigation deficit remains unchanged.

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 (gray) solutions MH highlighted in panel a (for other dam sizes, refer to Section S5 in the supplementary material). In particular, the selected IID solution is the one closer to the BID and POP solutions according to the Euclidean distance metric. 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 m 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

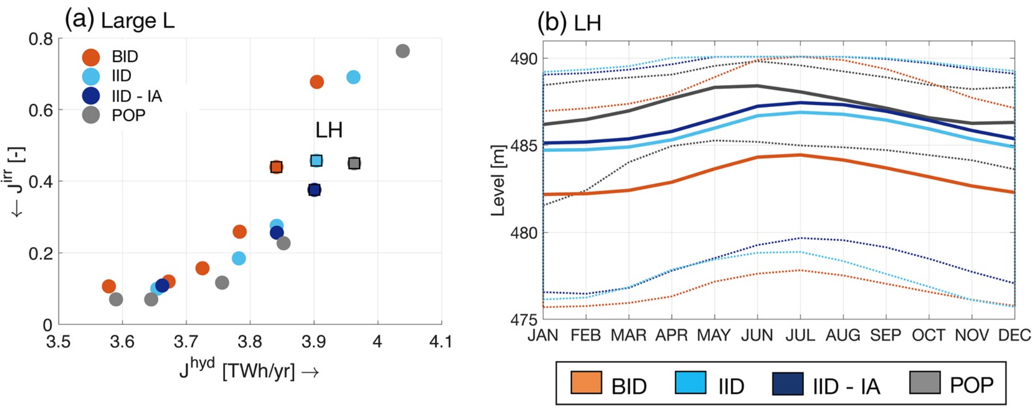

Figure 6. Panel (a) Two objective performance trade-offs for the large L dam size, where the BID (orange), IID (cyan), IID under interannual forecasts (blue), and POP (gray) solutions associated to a hydropower focused operating policy H are squared in black. Panel (b) monthly cyclo-stationary level 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.

(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 Interannual 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 of seasonal forecasts provided by weather forecast centers. However, it is also interesting to assess whether interannual streamflow forecasts contribute any further benefit particularly for large dam sizes, which can carry over large water volumes year-to-year. Even if interannual forecasts are not yet very accurate, their skill is expected to considerably increase in the near future (e.g., Nicolì et al., 2020; Redolat et al., 2020). To this end, we employ the IIS algorithm to identify the most informative lead times $\mathbf{I}_t^i$ out of an additional set of perfect interannual streamflow forecasts Ξ_t^i. The full details for our analysis of the interannual 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 interannual forecasts IID—IA (blue) infrastructure designs and perfect operating policies (gray). Interannual 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 interannual forecasts are valuable as they can store significant water volumes and carry them over interannually. Adding interannual forecasts does not bring any additional benefits in the irrigation-prone area, where the potential improvement is extremely limited.

Figure 6b investigates the reservoir dynamics under the LH solutions marked in panel a and conditioned over distinct information. As expected, under IID—IA the reservoir levels are about 2.5 and 0.5 m higher than the basic and seasonally informed respectively on average, moving closer to the perfect trajectories. 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. The value of the interannual forecasts becomes particularly evident when transitioning from wet to normal years, while the performance of BID is comparable to the both IID and IID—IA during periods with limited interannual variability (see Figure S6 of the supporting information).

4.4. Forecast Informed Infrastructure Design Using Biased Forecasts

This section explores how much the results discussed so far, which rely on the assumption of having perfect forecasts, degrade when using seasonal streamflow forecasts characterized by different biases, including ESP forecasts as well as synthetic forecasts with different levels of overestimation, underestimation, and underdispersion. These three specific biased were selected because they are expected to impact on the informed infrastructure design relying on the maximum streamflow over the next 7 months. For the infrastructure design, we employ the same settings of the IID analysis presented in Section 4.2, where the information used in the optimal infrastructure design is limited to s_t, t, and $qM7_t$.

Figure 7a displays the skill of the different biased forecasts. Here, skill is computed as $1 - RMSE_{\mathcal{S}} / RMSE_{\mathcal{C}}$, where $\mathcal{S}$ is the considered forecast system, while $\mathcal{C}$ is the streamflow climatology. This skill analysis shows that all the considered forecasts have a positive skill, i.e., they all outperform the streamflow climatology in terms of RMSE. Overestimated forecasts are the least skilful, while both underestimated and underdispersed forecasts attain relatively high values of skill. The skill of ESP forecasts is instead in between, suggesting the need of investigating the value of this realistic forecasts in informing the infrastructure design.

A summary of the 10 IID experiments using biased forecasts is reported in Figure 7b, which illustrates the values of hypervolume metric for the different forecast systems for small (left panel), medium (middle panel) and large (right panel) dam sizes. Valuable forecasts should be characterized by hypervolume values higher than the ones of BID solutions, possibly getting as close as possible to the POP ones. Similarly to the results obtained with perfect seasonal forecasts (Figure 4), the highest forecast value is again obtained for small dams, with an average 0.12 increase in hypervolume relative to the BID solutions (+27%). The average forecast value for medium and large dams is instead equal to 0.05 (+12%) and 0.07 (+19%), respectively. Medium dam sizes are therefore the most impacted by streamflow forecast biases, probably because they have not a sufficient storage capacity to buffer large flood water volumes and must be carefully operated in order to avoid spilling and thus wasting water that could have been used for both hydropower production and irrigation supply.

Interestingly, these results show that the ranking of the forecast systems based on skill (panel a) differs from the one based on forecast value (panel b). While underdispersed forecasts with Pbias equal to −6% are the most skilful forecasts, they never generates the highest hypervolume metric, which is obtained either using underestimated forecasts with −30% Pbias (small and medium dams) or underdispersed forecasts with −10% Pbias. Conversely, overestimated forecasts that have the lowest skill do not necessarily attain the lowest hypervolume metric that is, obtained using the ESP forecasts for all dam sizes. This result can be explained by the learning ability of the operating policy that allows a partial detection of the systematic biases of the synthetic forecasts and the mitigation of their negative impacts on operational decisions. The larger variability of ESP forecast errors (see the scatterplots in Figure S7 of the supporting information) is instead less interpretable by the operating policy and gen-

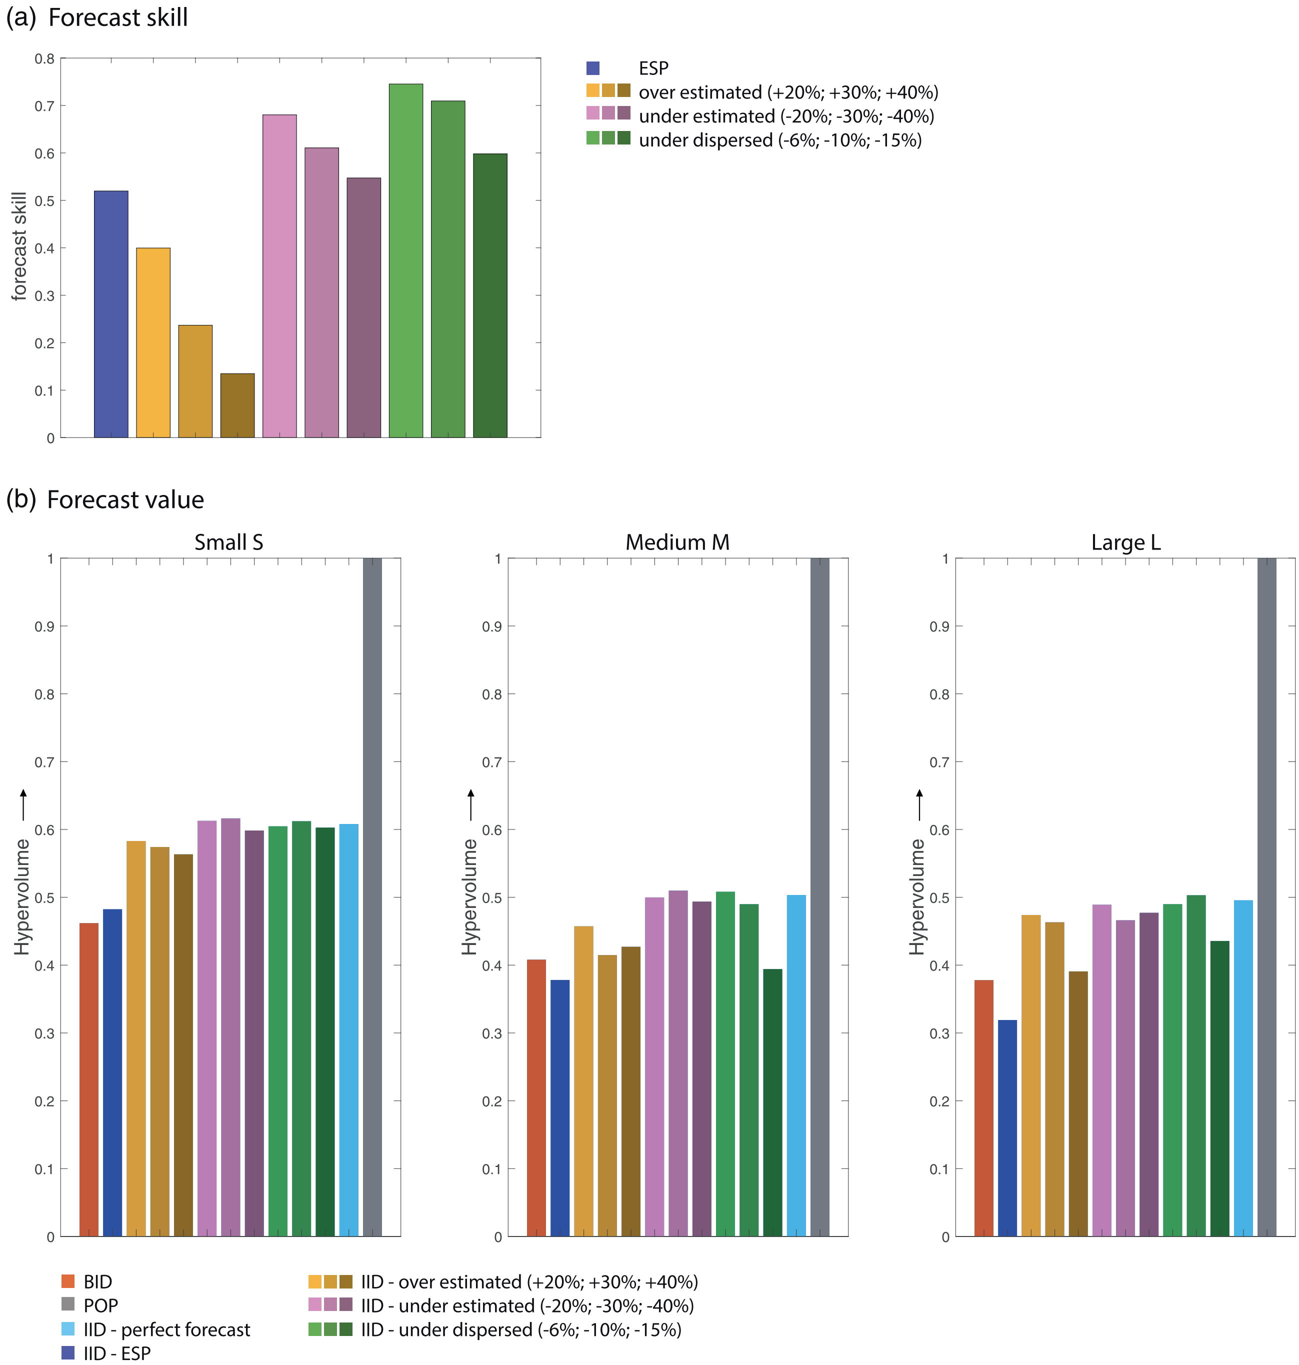

Figure 7. Analysis of forecast skill (panel a) and forecast value (panel b) for ESP forecasts and synthetic forecasts with different levels of overestimation, underestimation, and underdispersion.

erates lower system performance. Finally, it is worth noticing that the forecast value of ESP is positive only in the case of small dams, while this system does not yield any advantage for medium and large dams, confirming the need of advancing existing forecast systems to allow making the most of seasonal streamflow forecasts.

5. Discussion and Future Research

Despite the proposed methodology is general and transferable to other contexts, the numerical results illustrated in the previous section refer to the specific model and experimental settings described in Section 3. In particular, we demonstrate the value of forecast information for infrastructure design assuming a stationary future, where both inflow and demands as well as the preferences of the system operator remain unvaried over time. While in the context of this work we decided to perform the analysis over historical conditions to allow comparing our findings with the existing Kariba Dam, in a future work it would be interesting to assess the robustness of our results against future, deep uncertainties, which Bertoni et al. (2019) showed to be highly impacting on the system dynamics, in order to provide recommendations regarding the ongoing planning of new dams in the basin.

Moreover, our study focuses on the sizing of a single reservoir according to the widely adopted project-by-project planning approach Schmitt et al. (2018). Kariba is however part of a multireservoir network controlling a large share of the water flowing along the Zambezi River, which will be further expanded in the next decades. Extending our analysis to assess the value of forecast information in planning new reservoirs as part of a coordinated network is also warranted.

Another future research direction will be to test the value of the information provided by an existing forecast system, such as the multimodel and multiproduct seasonal hydrological streamflow forecasting platform available for the Upper Zambezi (Roy et al., 2017; SERVIR Water Africa-Arizona Team (SWAAT), 2019) or the Global Flood Awareness System - seasonal (Emerton et al., 2018), which might improve the performance of the ESP forecasts. Since most existing systems provide probabilistic forecasts, it could be relevant to assess the sensitivity of the resulting informed infrastructure designs on how the decision maker is interpreting the forecast ensemble depending on its level of risk aversion (Giuliani et al., 2020).

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. However, any application of our approach will require a continuous monitoring of the forecast skill to timely capture potential changes in forecast accuracy induced by the evolution of the large-scale climatic teleconnections that are the main source of predictability at the seasonal time scale (Dutta & Maity, 2018; Kumar et al., 1999; Zhang et al., 2019).

6. 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 and how information feedbacks shape the resulting infrastructure designs. The approach is demonstrated through an ex post design analysis of the Kariba Dam in the ZRB, where we investigated if forecast information could allow the discovery of a less costly and more efficient design solutions with respect to the existing reservoir.

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 seasonal forecasts than smaller dams serving irrigated agriculture. In particular, the same hydropower production levels of a dam operated with no forecasts can be designed with a 20% reduction in capital costs (i.e., 20% smaller reservoir). Dam size being the same, forecasts allow an increase of 60 GWh/yr hydropower production in a dam large as the existing Kariba, corresponding to more than 25% of the yearly average electricity consumption by the agriculture sector in Zambia (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 1,200 MW, Mhpanda 1,350 MW, and Batoka 1,600 MW), which will cumulatively add more than double the installed power in Kariba (1,830 MW), might increase this rate to 75 %. It is worth mentioning that our approach is portable to the design of other infrastructures, including the planned expansion of the irrigation districts in the ZRB that could also benefit from a more flexible and efficient operation of the water supply system.

In the end, when tested over realistic streamflow forecasts characterized by different biases, the design of the Kariba reservoir system is shown to not significantly benefit from ESP forecasts. At the same time, our results suggest that the informed infrastructure design is particularly sensitive to overestimation biases. This finding, combined with the large potential value obtained with perfect streamflow forecasts, represents a valuable insight for driving future research efforts aimed at advancing existing forecast systems.

Data Availability Statement

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.

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.

References

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. https://doi.org/10.1002/2015wr017864

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*(12), 1367–1390. https://doi.org/10.1029/2019ef001235

Bertoni, F., Giuliani, M., & Castelletti, A. (2020). Integrated design of dam size and operations via reinforcement learning. *Journal of Water Resources Planning and Management*, *146*(4), 04020010. https://doi.org/10.1061/(ASCE)WR.1943-5452.0001182

Block, P. (2011). Tailoring seasonal climate forecasts for hydropower operations. *Hydrology and Earth System Sciences*, *15*(4), 1355–1368. https://doi.org/10.5194/hess-15-1355-2011

Cassagnole, M., Ramos, M.-H., Zalachori, I., Thirel, G., Garçon, R., Gailhard, J., & Ouillon, T. (2020). Impact of the quality of hydrological forecasts on the managementand revenue of hydroelectric reservoirs – A conceptual approach. *Hydrology and Earth System Sciences Discussions*, 1–36. https://doi.org/10.5194/hess-2020-410

Celeste, A. B., & Billib, M. (2009). Evaluation of stochastic reservoir operation optimization models. *Advances in Water Resources*, *32*(9), 1429–1443. https://doi.org/10.1016/j.advwatres.2009.06.008

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*, *55*(8), 7400–7418. https://doi.org/10.1029/2018wr024462

Coello Coello, C. A., Lamont, G. B., & Van Veldhuizen, D. A. (2007). Evolutionary algorithms for solving multi-objective problems (Vol. 5). Springer.

Day, G. N. (1985). Extended streamflow forecasting using NWSRFS. *Journal of Water Resources Planning and Management*, *111*(2), 157–170. https://doi.org/10.1061/(asce)0733-9496(1985)111:2(157)

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. https://doi.org/10.1016/j.advwatres.2017.02.012

Doblas-Reyes, F. J., Balmaseda, M. A., Weisheimer, A., & Palmer, T. N. (2011). Decadal climate prediction with the European centre for medium-range weather forecasts coupled forecast system: Impact of ocean observations. *Journal of Geophysical Research*, *116*(D19). https://doi.org/10.1029/2010jd015394

Dutta, R., & Maity, R. (2018). Temporal evolution of hydroclimatic teleconnection and a time-varying model for long-lead prediction of Indian summer monsoon rainfall. *Scientific Reports*, *8*(1). https://doi.org/10.1038/s41598-018-28972-z

Emerton, R., Zsoter, E., Arnal, L., Cloke, H. L., Muraro, D., Prudhomme, C., et al. (2018). Developing a global operational seasonal hydro-meteorological forecasting system: GloFAS-Seasonal v1.0. *Geoscientific Model Development*, *11*(8), 3327–3346. https://doi.org/10.5194/gmd-11-3327-2018

Faber, B. A., & Stedinger, J. R. (2001). Reservoir optimization using sampling SDP with ensemble streamflow prediction (ESP) forecasts. *Journal of Hydrology*, *249*(1-4), 113–133. https://doi.org/10.1016/s0022-1694(01)00419-x

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. https://doi.org/10.1175/1525-7541(2003)004<1105:vonwse>2.0.co;2

Fuchs, M., Kling, H., Wipplinger, B., & Faber, R. (2018). Inflow forecasting for improved dam safety management on the Niger river. *Twenty-Sixth International Congress on Large Dams*, CRC Press.

Galelli, S., & Castelletti, A. (2013). Assessing the predictive capability of randomized tree-based ensembles in streamflow modelling. *Hydrology and Earth System Sciences*, *17*(7), 2669–2684. https://doi.org/10.5194/hess-17-2669-2013

Galelli, S., & Castelletti, A. (2013). Tree-based iterative input variable selection for hydrological modeling. *Water Resources Research*, *49*(7), 4295–4310. https://doi.org/10.1002/wrcr.20339

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. https://doi.org/10.1016/j.envsoft.2014.08.015

Gandolfi, C., Guariso, G., & Togni, D. (1997). *Water Resources Management*, 11(5), 377–393. https://doi.org/10.1023/a:1007964732399

Georgakakos, A. P., Yao, H., Kistenmacher, M., Georgakakos, K. P., Graham, N. E., Cheng, F.-Y., et al. (2012). Value of adaptive water resources management in Northern California under climatic variability and change: Reservoir management. *Journal of Hydrology*, *412-413*, 34–46. https://doi.org/10.1016/j.jhydrol.2011.04.038

Georgakakos, K. P., Graham, N. E., Carpenter, T. M., & Yao, H. (2005). Integrating climate-hydrology forecasts and multi-objective reservoir management for northern California. *Eos, Transactions American Geophysical Union*, *86*(12), 122. https://doi.org/10.1029/2005eo120002

Geressu, R. T., & Harou, J. J. (2015). Screening reservoir systems by considering the efficient trade-offs—informing infrastructure investment decisions on the Blue Nile. *Environmental Research Letters*, *10*(12), 125008. https://doi.org/10.1088/1748-9326/10/12/125008

Geurts, P., Ernst, D., & Wehenkel, L. (2006). Extremely randomized trees. *Machine Learning*, *63*(1), 3–42. https://doi.org/10.1007/s10994-006-6226-1

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. https://doi.org/10.1061/(asce)wr.1943-5452.0000570

Giuliani, M., Crochemore, L., Pechlivanidis, I., & Castelletti, A. (2020). From skill to value: Isolating the influence of end-user behaviour on seasonal forecast assessment. *Hydrology and Earth System Sciences*, *24*(12), 5891–5902. https://doi.org/10.5194/hess-2019-659

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. https://doi.org/10.1002/2015wr017044

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. https://doi.org/10.1109/tcst.2017.2705162

Giuliani, M., Zaniolo, M., Castelletti, A., Davoli, G., & Block, P. (2019). Detecting the state of the climate system via artificial intelligence to improve seasonal forecasts and inform reservoir operations. *Water Resources Research*, *55*(11), 9133–9147. https://doi.org/10.1029/2019WR025035

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. https://doi.org/10.1175/2009jamc2135.1

Guariso, G., Rinaldi, S., & Soncini-Sessa, R. (1986). The management of Lake Como: A multiobjective analysis. *Water Resources Research*, *22*(2), 109–120. https://doi.org/10.1029/wr022i002p00109

Hadka, D., & Reed, P. (2013). Borg: An auto-adaptive many-objective evolutionary computing framework. *Evolutionary Computation*, *21*(2), 231–259. https://doi.org/10.1162/evco_a_00075

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. https://doi.org/10.1016/j.envsoft.2014.10.014

Hall, W. A., Askew, A. J., & Yeh, W. W.-G. (1969). Use of the critical period in reservoir analysis. *Water Resources Research*, *5*(6), 1205–1215. https://doi.org/10.1029/wr005i006p01205

Ham, Y.-G., Rienecker, M. M., Suarez, M. J., Vikhliaev, Y., Zhao, B., Marshak, J., et al. (2014). Decadal prediction skill in the GEOS-5 forecast system. *Climate Dynamics*, *42*(1-2), 1–20. https://doi.org/10.1007/s00382-013-1858-x

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. https://doi.org/10.1061/(asce)0733-9496(2002)128:2(91)

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

IRENA. (2012). *Renewable power generation costs in 2012: An overview* (Technical Report). IRENA.

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. https://doi.org/10.1029/wr026i003p00447

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. https://doi.org/10.1061/(asce)0733-9496(1997)123:6(327)

Koutsoyiannis, D., & Economou, A. (2003). Evaluation of the parameterization-simulation-optimization approach for the control of reservoir systems. *Water Resources Research*, *39*(6). https://doi.org/10.1029/2003wr002148

Kumar, K. K. (1999). On the weakening relationship between the Indian Monsoon and ENSO. *Science*, *284*(5423), 2156–2159. https://doi.org/10.1126/science.284.5423.2156

Libisch-Lehner, C. P., Nguyen, H. T. T., Taormina, R., Nachtnebel, H. P., & Galelli, S. (2019). On the value of ENSO state for urban water supply system operators: opportunities, trade-offs, and challenges. *Water Resources Research*, *55*(4), 2856–2875. https://doi.org/10.1029/2018wr023622

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. https://doi.org/10.1002/2016wr019552

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. https://doi.org/10.1175/1520-0442(2004)017<0174:peolhp>2.0.co;2

Montaseri, M., & Adeloye, A. J. (1999). Critical period of reservoir systems for planning purposes. *Journal of Hydrology*, *224*(3-4), 115–136. https://doi.org/10.1016/s0022-1694(99)00126-2

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. https://doi.org/10.1175/1520-0434(1993)008<0281:wiagfa>2.0.co;2

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. https://doi.org/10.1029/2018wr023177

Nicolì, D., Bellucci, D., Iovino, D., Ruggieri, P., & Gualdi, S. (2020). The impact of the AMV on Eurasian summer hydrological cycle. *Scientific Reports*, *10*(1). https://doi.org/10.1038/s41598-020-71464-2

Oliveira, R., & Loucks, D. P. (1997). Operating rules for multireservoir systems. *Water Resources Research*, *33*(4), 839–852. https://doi.org/10.1029/96wr03745

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

Payet-Burin, R., Kromann, M., Pereira-Cardenal, S., Strzepek, K. M., & Bauer-Gottwein, P. (2019). WHAT-IF: An open-source decision support tool for water infrastructure investment planning within the water–energy–food–climate nexus. *Hydrology and Earth System Sciences*, *23*(10), 4129–4152. https://doi.org/10.5194/hess-23-4129-2019

Perelman, L., Housh, M., & Ostfeld, A. (2013). Robust optimization for water distribution systems least cost design. *Water Resources Research*, *49*(10), 6795–6809. https://doi.org/10.1002/wrcr.20539

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. https://doi.org/10.1029/2018wr022743

Raso, L., Schwanenberg, D., van de Giesen, N. C., & van Overloop, P. J. (2014). Short-term optimal operation of water systems using ensemble forecasts. *Advances in Water Resources*, *71*, 200–208. https://doi.org/10.1016/j.advwatres.2014.06.009

Redolat, D., Monjo, R., Paradinas, C., Pórtoles, J., Gaitán, E., Prado-Lopez, C., & Ribalaygua, J. (2020). Local decadal prediction according to statistical/dynamical approaches. *International Journal of Climatology*, *40*(13), 5671–5687. https://doi.org/10.1002/joc.6543

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. https://doi.org/10.1016/j.advwatres.2012.01.005

Rodriguez, R. A., Becker, S., & Greiner, M. (2015). Cost-optimal design of a simplified, highly renewable pan-European electricity system. *Energy*, *83*, 658–668. https://doi.org/10.1016/j.energy.2015.02.066

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. https://doi.org/10.1002/2016wr019752

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). https://doi.org/10.1029/2010wr009502

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. https://doi.org/10.1016/j.advwatres.2017.09.014

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). https://doi.org/10.1029/2009wr007821

Schmitt, R. J. P., Bizzi, S., Castelletti, A., & Kondolf, G. M. (2018). Improved trade-offs of hydropower and sand connectivity by strategic dam planning in the Mekong. *Nature Sustainability*, *1*(2), 96–104. https://doi.org/10.1038/s41893-018-0022-3

Schuster, M., Grieger, J., Richling, A., Schartner, T., Illing, S., Kadow, C., et al. (2019). Improvement in the decadal prediction skill of the North Atlantic extratropical winter circulation through increased model resolution. *Earth System Dynamics*, *10*(4), 901–917. https://doi.org/10.5194/esd-10-901-2019

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

Smith, D. M., Eade, R., Scaife, A. A., Caron, L.-P., Danabasoglu, G., DelSole, T. M., et al. (2019). Robust skill of decadal climate predictions. *npj Climate and Atmospheric Science*, *2*(1). https://doi.org/10.1038/s41612-019-0071-y

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. https://doi.org/10.1007/s11269-012-0051-z

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). https://doi.org/10.1029/2011wr011540

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

Turner, S. W. D., 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. https://doi.org/10.5194/hess-21-4841-2017

U.S. Army Corps of Engineers. (1975). Hydrologic Engineering Methods for Water Resources Development: Reservoir Yield (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 U.S. power planning. *Journal of Applied Meteorology and Climatology*, *45*(5), 653–673. https://doi.org/10.1175/jam2361.1

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). https://doi.org/10.1029/2011wr011366

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

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. https://doi.org/10.1002/2014wr015964

You, J.-Y., & Cai, X. (2008). Determining forecast and decision horizons for reservoir operations under hedging policies. *Water Resources Research*, *44*(11). https://doi.org/10.1029/2008wr006978

Zhang, T., Wang, T., Krinner, G., Wang, X., Gasser, T., Peng, S., et al. (2019). The weakening relationship between Eurasian spring snow cover and Indian summer monsoon rainfall. *Science Advances*, *5*(3), eaau8932. https://doi.org/10.1126/sciadv.aau8932

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. https://doi.org/10.1061/(asce)wr.1943-5452.0000432

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. https://doi.org/10.1109/tevc.2003.810758