

Variability, Attributes, and Drivers of Optimal Forecast-Informed Reservoir Operating Policies for Water Supply and Flood Control in California

William Taylor, P.E.¹; Zachary P. Brodeur, Ph.D.²; Scott Steinschneider, Ph.D.³; John Kucharski⁴; and Jonathan D. Herman, Ph.D.⁵

Abstract: Reservoirs balance multiple conflicting objectives, including flood control and water supply. In California, shifts in seasonal hydrologic patterns under climate change will amplify the difficulties in balancing flood control with water supply. Current flood control policies are based on fixed seasonal rule curves determined by the observed timing and magnitude of floods in the record. These rule curves generally require the release of wet season inflows, reducing the available stored water for use during the dry season. Here we investigate the potential for forecast-informed reservoir operations (FIRO) to increase water supply availability while minimizing additional flood risk at 14 reservoirs in the Sacramento, San Joaquin, and Tulare river basins. We use a differential evolution algorithm to train risk-based reservoir operation policies with an ensemble of historical forecasts over the period 2013–2023. Results show an average 8.1% increase in storage normalized by capacity, though this varies across reservoirs. The forecast-informed policies also reduce the occurrence of high-magnitude releases throughout the system. The accumulation of benefits is sensitive to the timing and magnitude of flood events, and most of the cumulative benefit is obtained during a few years. Under cross-validation, we find that large floods are needed in the training data to avoid overfitting the policy. We further examine the relationship between reservoir properties and FIRO benefits, finding that the ratio of peak inflow magnitude to maximum safe release correlates with increased storage under the FIRO policy, while the ratio of mean inflow to capacity correlates to the reduction of high-magnitude releases. This study highlights how adaptive reservoir management policies can yield water supply benefits without an increase in flood risk, given adequate historical data for policy training. These policies may be a valuable adaptation to climate change but require careful validation and out-of-sample testing. DOI: 10.1061/JWRMD5.WRENG-6471. © 2024 American Society of Civil Engineers.

Introduction

Reservoirs face the challenging task of satisfying multiple objectives at once: flood protection, water supply, environmental flows, and hydropower generation, among others. Several of these objectives often directly conflict. For example, reservoir operators can choose to leave more space in a reservoir empty to buffer against flood events, but at the expense of water supply storage. In many reservoirs across California and other regions of the Western US, current water supply and flood control policies are based on fixed seasonal rule curves, which are determined by the observed timing

and magnitude of historical streamflows (US Army Corps of Engineers 2017). Climate change will amplify the difficulties in balancing flood protection with water supply. Warmer temperatures will bring more extreme precipitation events (Das et al. 2011), requiring larger reservoir flood pools to mitigate increased flood flows. Additionally, earlier and decreased snowmelt from regional mountain ranges (Pathak et al. 2018) will diminish natural storage capacity, further increasing water supply stress.

One way to improve the trade-off between flood protection and water supply storage objectives without expensive investments in new infrastructure is to adjust current operating policies through the use of short-to-medium range (1–14 day) forecasts. While some reservoir operators are already using forecasts to guide their decisions (Hejazi et al. 2008), this use is based on operator expertise and is not formalized in the operating rules of major, federally operated reservoirs (US Army Corps of Engineers 2017). Short-to-medium range hydrologic forecasts allow release decisions to be made dynamically based on current and forecasted data instead of average historical patterns (Nayak et al. 2018), and can provide operational benefits beyond water supply and flood control objectives. A growing body of literature highlights the benefits of formalized forecast-informed reservoir operations (FIRO), although these benefits often depend on the objectives, attributes, and constraints of the system (Doering et al. 2021; Semmendinger and Steinschneider 2024; Turner et al. 2017; Yang et al. 2020).

The application of FIRO has gained traction in California and the Western US in recent years, supported by the high predictability of large precipitation events in the region (Jasperse et al. 2020; Woodside et al. 2022). Much of California's precipitation comes from atmospheric rivers (ARs) (AMS 2022; Dettinger et al. 2011),

¹Ph.D. Student, Dept. of Civil and Environmental Engineering, Univ. of California, Davis, CA 95616 (corresponding author). ORCID: <https://orcid.org/0009-0000-4033-6615>. Email: wltaylor@ucdavis.edu

²Postdoctoral Researcher, Dept. of Biological and Environmental Engineering, Cornell Univ., Ithaca, NY 14853. ORCID: <https://orcid.org/0000-0001-5242-7696>

³Associate Professor, Dept. of Biological and Environmental Engineering, Cornell Univ., Ithaca, NY 14853. ORCID: <https://orcid.org/0000-0002-8882-1908>

⁴Ph.D. Student, Dept. of Land, Air and Water Resources, Univ. of California, Davis, CA 95616; Researcher, Engineer Research and Development Center, US Army Corps of Engineers, Vicksburg, MS.

⁵Associate Professor, Dept. of Civil and Environmental Engineering, Univ. of California, Davis, CA 95616. ORCID: <https://orcid.org/0000-0002-4081-3175>

Note. This manuscript was submitted on November 2, 2023; approved on May 20, 2024; published online on August 5, 2024. Discussion period open until January 5, 2025; separate discussions must be submitted for individual papers. This paper is part of the *Journal of Water Resources Planning and Management*, © ASCE, ISSN 0733-9496.

which are narrow bands of intense moisture transport originating from the Pacific Ocean that result in heavy precipitation upon landfall. While the exact landfall location of ARs can be difficult to forecast even at short lead times (1–3 days) (Lavers et al. 2018), their occurrence and magnitude are more predictable than convective storms that dominate extreme precipitation in other regions of the United States (Feng et al. 2016). The dynamic nature of ARs creates challenges for seasonal forecasts (Warner et al. 2015), but short-to-medium range forecasts (1–14 days) may be accurate enough to provide decision makers with valuable information to conserve water and decrease flood risk. Further, machine learning and postprocessing techniques continue to improve the accuracy of AR and subsequent streamflow forecasts (Badrinath et al. 2023; Chapman et al. 2019), which is important considering ARs are predicted to increase in intensity and frequency in the Western US (Corringham et al. 2022).

In California, there is an ongoing effort to conduct FIRO viability assessments and implementation for several reservoirs with water supply, flood control, and ecological objectives (Jasperse et al. 2020). The first viability assessment was conducted for Lake Mendocino on the Russian River, where a FIRO policy was able to increase median storage by 33% for the period 1985–2010 without an increase in flood risk (Delaney et al. 2020). The authors used ensemble streamflow forecasts to represent uncertainty and developed a probabilistic decision support system called the ensemble forecast operations (EFO) policy to implement release decisions that achieved a targeted level of flood risk reduction. EFO was also applied to a case study of Prado Dam in Southern California, where the strategy increased managed aquifer recharge by 6.1%–13.3% annually, depending on the accepted level of flood risk (Woodside et al. 2022).

Despite the increased interest in FIRO research and implementation within the last decade, key questions remain. In particular, the variability of storage and flood control benefits across reservoirs is poorly understood, as are the attributes of infrastructure,

hydrology, and forecasts that correlate with those benefits. Understanding how these attributes affect the viability of revised operations has implications for future site selection. To address this research gap, we contribute a novel optimization framework that generalizes the risk-based EFO strategy developed in prior studies of individual reservoirs in the Western US (Delaney et al. 2020; Jasperse et al. 2020; Woodside et al. 2022). We formulate the optimization problem by parameterizing the EFO risk curve and a seasonal flood pool level, enabling a computationally efficient framework to test the water supply and flood control benefits of FIRO operations across a set of reservoirs that differ in infrastructure, hydrology, and forecast skill. Using these optimized policies, we quantify the attributes of infrastructure and hydrology across sites that most strongly relate to storage and flood control benefits and the degree of risk aversion. The proposed approach and resulting insights of this work can be used to inform the selection of sites where FIRO is likely a cost-effective adaptation strategy, and develop baseline operating policies that can be further refined in more detailed viability assessments.

Methods

Case Study

The Sacramento, San Joaquin, and Tulare river basins comprise 58% of California's average annual streamflow and cover 153,000 km² (59,000 mi²) (California Department of Water Resources 2023). As shown in Fig. 1, Shasta, Oroville, New Bullards Bar, and Folsom reservoirs are located on the Sacramento River and its tributaries. Pardee, New Hogan, New Melones, Don Pedro, New Exchequer, and Millerton reservoirs are located within the San Joaquin watershed. Finally, Pine Flat, Terminus, Success, and Isabella reservoirs are within the Tulare watershed. The hydrology of this region follows a Mediterranean climate, with wet winters and springs followed by dry summers and falls (Peel et al. 2007).




Fig. 1. Location of each reservoir in the experiment. Capacities listed in thousand acre-feet (TAF).

California historically has experienced weather whiplash, shifting between drought and flood conditions with little advanced warning, a phenomenon that has increased in frequency under climate change (Swain et al. 2018). Reservoirs play a crucial role in buffering this variability (Pan et al. 2019). The majority of reservoir inflow comes during the winter and spring from storm events and when snowmelt runs off from the Sierra Nevada (Pathak et al. 2018). The reservoirs support California's Central Valley agricultural industry, where crops valued at \$17 billion are grown annually, representing 8% of the nation's agricultural revenue (USGS 2023). The reservoirs also provide municipal water supply to California's 39 million residents (US Census Bureau 2023), as well as environmental flows to the Sacramento–San Joaquin Delta (Zimmerman et al. 2018). In addition to meeting these water demands, the reservoir system must also mitigate flood events. Current reservoir operations maintain a flood pool during the wet season to buffer flood waters, which in some years may prevent storage of excess water from floods and snowmelt. This can result in a large opportunity cost from lost water supply for agricultural, urban, and environmental uses during the dry summer months.

Simulation Model and Data

A baseline, daily simulation model of the Central Valley reservoir system was developed in prior work on climate vulnerability assessments for this region (Cohen et al. 2020; Steinschneider et al. 2023). The model uses parsimonious water supply and flood control rules to approximate observed release policies, and the parameters are calibrated over the period 2009–2023. Each reservoir follows the mass balance relationship outlined in Eq. (1):

$$S_t = S_{t-1} + Q_t - R_t \quad (1)$$

where S_t and S_{t-1} = storage at the current and previous time step; Q_t = observed inflow; and R_t = system release. The value of R_t is initially based on the median historical release at that specific day of the water year. If the current storage is less than the median historical storage, an exponential hedging function is applied to decrease the release, reserving some additional water. The model also considers the top of conservation storage (TOCS) limit, which changes seasonally through the water year. If Eq. (1) predicts that the storage will exceed the allowable storage defined by the TOCS, then R_t is increased until the storage limit constraint is satisfied. The release values are also limited by a maximum safe release value, which accounts for the infrastructure limits at the dam and downstream. Finally, the release values are constrained by ramping rates, limiting the change in release between each 24-h time step. The model has an average accuracy of $R^2 = 0.865$ across reservoirs between the modeled storage values and the observed values during 2009–2023 (Steinschneider et al. 2023). The model does not consider external factors that may have influenced observed release decisions, such as changes in environmental regulations or temporary changes in operations during infrastructure repairs. It also does not consider informal use of forecasts by local expert reservoir operators.

The model inputs required for the baseline simulation model are:

- TOCS values: The top of conservation storage values for each reservoir are drawn from the California Food Energy Water System (CALFEWS) model (Zeff et al. 2021). The values represent a percentage of the reservoir capacity that is allowed to be used for water storage. The TOCS value is imposed during the flood season but not used during the dry season when there is little risk of flooding. For reservoirs where the TOCS depends on conditional factors, such as precipitation or snowmelt, we use an

average seasonal value assuming that forecast-based operations would serve the same purpose.

- Ramping rates: Ramping rates [cubic feet per second (cfs) per day] define how large of a change in reservoir release can be made between daily time steps. For this experiment, we assumed maximum ramping rate values for each reservoir based on the highest observed difference in daily release over the study period. The rate is applied to both increasing and decreasing outflow. This method may not reflect the reality of the physical infrastructure as it currently stands, or the existing ramping rate policies, but it does provide a consistent way to evaluate the policies over the study period. Without the ramping rate constraints, the policy may schedule releases that change more rapidly than is realistic. Ramping rate restrictions are partly based on physical infrastructure limitations, but are also important for downstream environmental systems (Jager and Smith 2008).
- Max release: Maximum release values ($R_{\max}$, in cfs) represent the allowable safe release from the reservoir, which is determined by downstream levee capacity. An operating policy is heavily penalized for exceeding this limit, even though most of the reservoirs have historically exceeded this value when unavoidable due to spill. Values for each reservoir are drawn from Tustison (2020). In cases where the maximum observed release exceeded this value, we use the observed value instead.
- Capacity: The reservoir capacity, K , represents the total amount of water the reservoir can store, notwithstanding any top of conservation storage limitations.
- Inflows: The historical reservoir inflows were sourced from the California Data Exchange Center, which is managed by the California Department of Water Resources. The data are collected using streamflow gauges and aggregated into a daily mean flow rate. The daily data are available from 1994 to 2023, but the model runs are limited by the length of the forecast record (2013–2023, described in “FIRO-Based Operations”).
- Seasonal release pattern: The daily water supply release for each reservoir was estimated by finding the median release volume per day of the water year over the period 1994–2023 and scaling it based on annual hydrology following the method in Steinschneider et al. (2023). This method leverages the historical seasonal release pattern as a proxy for the actual system demand, which is difficult to estimate for individual reservoirs and contains a combination of urban, agricultural, and environmental demands throughout the state.

FIRO-Based Operations

A FIRO release policy is added to the baseline model by creating a piecewise linear risk tolerance curve for each reservoir. We construct the risk curve as piecewise linear to emulate the optimized EFO curves developed by Delaney et al. (2020). A linear curve also reduces the potential for overfitting. Risk in this context is defined as the fraction of forecast ensemble members (described subsequently) where the forecasted storage will exceed the TOCS if no releases are made. The risk tolerance curve for each reservoir is defined by two parameters: lead days with no allowable risk (i.e., zero ensemble members can exceed the TOCS), and the slope of the risk curve once some risk is tolerated (% risk allowed per day lead time). Fig. 2 provides a visual representation of the FIRO policy structure. At each time step, the risk is calculated for each of the next 14 forecasted days. If the risk exceeds the risk tolerance curve at the given lead time, the model will release the minimum amount of water needed to lower the risk below that allowed by the curve. The piecewise linear risk curve is a generalization of the formulation in Delaney et al. (2020), who represented the risk curve using




Fig. 2. FIRO policy structure. Parameters 1 and 2 control the risk tolerance curve for each reservoir, while Parameter 3 adjusts the TOCS multiplier.

14 decision variables, one for each lead time. The FIRO release (R_{FIRO}) is calculated in parallel with Eq. (1) and then compared to R_t , the release required to stay below the TOCS for the current time step. The target release is chosen as the maximum of the two

$$R_t = \begin{cases} R_t, & R_{\text{FIRO}} < R_t \\ R_{\text{FIRO}}, & R_{\text{FIRO}} > R_t \end{cases} \quad (2)$$

In addition to the two risk curve parameters, a third parameter is included in the policy, the TOCS multiplier (x_3 in Fig. 2). This allows the model to set the TOCS to a higher value if the additional risk can be offset by the forecast-based operating policy, and if the hydrology of the study period allows it. The FIRO policy is constrained by the same characteristics of the individual reservoirs: a safe release volume, a ramping rate, and total capacity. If the FIRO policy can stay within those bounds, the TOCS can increase, giving

the reservoir the potential to store more water during the winter. All three parameters of the FIRO policy are optimized using the differential evolution algorithm described in “Computations Experiment.”

The policy uses 1–14 day hydrologic forecasts generated by the California Nevada River Forecast Center (CNRFC 2023b), with data available from November 1, 2013 to present (May 24, 2023). The forecasts include daily inflow values for between 40 and 58 ensemble members depending on the reservoir and year, which we restrict to a randomly sampled set of 40 for consistency. Fig. 3(a) shows an example forecast at Oroville Reservoir on February 1, 2017. The forecasts over the next 14 days predict an increase in flow, but underestimate the peak of the observed flow. In this case, the risk tolerance curve would need to be combined with a flood pool to safely manage the event. However, the forecasts become more accurate as the event approaches. We use the




Fig. 3. (a) Example forecast ensemble for Oroville Reservoir inflows on February 1, 2017; and (b) NSE for the ensemble mean forecast across the study period for all sites.

Nash-Sutcliffe efficiency (NSE) as a metric to evaluate the accuracy of the CNRFC forecasts

$$\text{NSE} = 1 - \frac{\sum_{t=1}^T (Q_o^t - Q_m^t)^2}{\sum_{t=1}^T (Q_o^t - \bar{Q}_o)^2} \quad (3)$$

where $\bar{Q}_o$ = mean of the observed inflows; Q_m^t = CNRFC modeled inflow; and Q_o^t = observed inflow at time t . The NSE is a measure of goodness of fit for hydrological models, with an NSE value of 1 indicating a perfect model (McCuen et al. 2006). The NSE of the ensemble mean forecasts across lead times is shown in Fig. 3(b).

Computational Experiment

We tested four different operating policies to analyze the factors influencing their performance across reservoirs. Each policy builds off the baseline model.

- **Baseline:** The baseline model is simulated using all inputs except for the forecasts. It creates storage and release decisions based on seasonal rule curves derived from the median storage and release variables and calibrated to historical releases.
- **Baseline-TOCS:** This version of the model is the same as the baseline, except that parameter x_3 , the TOCS multiplier, is allowed to increase if the storage and release decisions stay within the constraints (safe release, ramping rate, and reservoir capacity). This policy isolates the benefit of using the forecasts from the additional benefit of the TOCS multiplier. The multiplier value is found by iteration for each reservoir, where x_3 is incrementally increased until the constraints cannot be met.
- **Forecast:** The FIRO model uses the piecewise linear risk tolerance curve to determine releases in advance of flood events. This is the operating policy that could be realistically implemented using forecast data and is the primary focus of this experiment.
- **Perfect:** This policy replaces the forecast ensemble with perfect forecast (observed) data over the next 14 days. This shows how the FIRO policy could perform under perfect forecast conditions. A large difference between this policy and the Forecast policy would indicate that an increase in forecast skill could further improve FIRO-related benefits.

We used an optimization approach to find the best-performing policy parameters for the Forecast and Perfect cases. The objective function seeks to maximize the normalized mean storage across reservoirs (0–1), with a penalty of 1.0 subtracted for each occurrence of a release over the maximum safe release value

$$\max J = \sum_i \left(\frac{\bar{S}_i}{K_i} - \sum_t I(R_t > R_{\max}) \right) \quad (4)$$

where K_i = capacity; $\bar{S}_i$ = average storage over the full time series for the i th reservoir; and I = indicator function. The value $\bar{S}_i$ is normalized by K_i to avoid the larger reservoirs dominating the objective function. The optimization is performed with the differential evolution algorithm (Storn and Price 1997) in Python's SciPy package (Virtanen et al. 2020) for approximately 30,000 function evaluations and 10 random seeds. Although not as efficient as gradient-based optimization, the differential evolution algorithm can evaluate a large solution space of noisy and discontinuous solutions. In general, evolutionary algorithms have been shown to be an effective method for reservoir policy search problems (Gupta et al. 2020; Reed et al. 2013; Zatarain Salazar et al. 2016). After finding the optimal parameters using the algorithm, we use those parameters to run additional simulations and analyze the results.

We first compare the mean storage and frequency of maximum releases for each reservoir under each of the four policies to

determine how the policies rank in terms of their primary performance goals. The frequency of maximum releases gives insight into which of the four policies may be operationally viable because an increased frequency of maximum releases is highly undesirable for flood risk mitigation, even if such releases stay below damaging downstream release thresholds ($R_{\max}$). We then diagnose the FIRO policies in more detail, including their optimal parameter values and the site-specific factors (e.g., reservoir capacity, forecast accuracy, maximum safe release) that correlate with storage and flood control benefits across reservoirs. Additionally, we quantify the variability of policy parameters across random seeds, which provides insight about the policy sensitivity to the training period, and has important implications for operational implementation. Finally, we conclude with a leave-one-out cross-validation experiment, showing how the performance of the optimized policies may change in water years held out of the training data set.

Results

In the results that follow, Figs. 4–9 are based on policies trained to the entire available record, while Table 1 and Fig. 10 show results from the cross-validation experiment.

Storage and Flood Control

Fig. 4(a) shows the optimized policy results in terms of normalized mean storage. The Forecast policy increases mean storage by 8.1% of capacity over the baseline policy, averaged across all reservoirs. For an additional measure of aggregate results, we find that the Forecast policy increases the median storage at the end of the flood management season (defined as May 1) by 10.9%, averaged across all reservoirs. The positive objective function values listed in the legend of Fig. 4(a) indicate that all policies keep releases below $R_{\max}$, avoiding large flood penalties. The largest gains relative to capacity occur in the San Joaquin and Tulare reservoirs because these have lower safe release rates and thus benefit more from forecasts of large floods (investigated in detail in “Factor Analysis”). In absolute terms, the total additional systemwide storage created by the Forecast policy relative to the Baseline is 1.35 billion cubic meters (1.1 million acre-feet), averaged across the simulation period. Further, a comparison of the Forecast and Perfect policies suggests that the current forecast accuracy is sufficient to achieve nearly the same benefit as perfect information during this period. The forecasts are highly accurate at 3–5 day lead times [Fig. 3(b)], and nearly all of the reservoirs can safely release their flood pool volumes in this time. We note in Fig. 4(a) that the Perfect policy has the highest objective value, but does not have the highest storage for some reservoirs due to the aggregate formulation of the objective function. Finally, while Fig. 4(a) shows that the Forecast policy increases average storage, it also suggests that most of this storage benefit arises from raising the TOCS value during a simulation period (2013–2023) with relatively few large flood events, rather than from the use of forecasts. This is shown by the comparison of the Baseline-TOCS policy and the Forecast policies in Fig. 4(a).

However, the Baseline-TOCS policy cannot achieve this performance without increasing the frequency of maximum safe releases [Fig. 4(b)], which introduces undesirable flood risk. Fig. 4(b) also shows that the Forecast policy reduces the frequency of the maximum release for most reservoirs. The Baseline-TOCS is forced to make more frequent, higher-volume releases to avoid spills, while the Forecast policy spreads the releases out over several days, reducing the magnitude of releases and decreasing the risk to downstream infrastructure. This indicates the flood control benefits of FIRO in addition to the potential surplus storage, and these two




Fig. 4. (a) Mean normalized storage per policy, represented as a percentage of total capacity. The optimized objective function values for one random seed are displayed in the legend. (b) Number of maximum release occurrences (days) by policy. Under certain policies, some of the reservoirs are never forced to make a release at the maximum safe value.




Fig. 5. (a) Storage plot for the entire study period at Oroville Reservoir. The Baseline-TOCS, Forecast, and Perfect policies overlap at many time steps. (b) Release plot for Oroville Reservoir during a 2017 flood event. The Baseline-TOCS policy closely follows the Baseline policy for most of the event.

improvements arise from the ensemble risk curve combined with a higher TOCS.

The time series of reservoir storage and releases highlights the differences in how each policy makes release decisions. Fig. 5(a) shows the storage level at Oroville Reservoir across the entire study period, with the Baseline-TOCS, Forecast, and Perfect policies all keeping consistently higher storage levels than the Baseline policy, especially during wet and average hydrologic years. Fig. 5(b) depicts the releases under each policy during the largest flood event in

the study period, which is associated with the 2017 Oroville spillway failure (Koskinas et al. 2019) (note the model does not simulate the spillway failure). Fig. 5(b) shows that the Perfect policy anticipates the large inflows well in advance, and determines a steady release rate through the entire flood period. The Forecast policy starts to ramp up releases prior to the flood peak as the forecasts become more accurate and more ensemble members exceed the risk tolerance curve. When the flood peak occurs, the Baseline and Baseline-TOCS policies are forced to react and release a large




Fig. 6. Risk curve and parameter values for each reservoir as trained under the best-performing random seed: (a) risk curve constructed with the optimal parameters; (b) x_1 (days with no allowable risk); (c) x_2 (slope of risk curve); and (d) TOCS multiplier. (b and c) Only show the Forecast policy parameters because the other policies do not use the risk tolerance curve to determine release values.

amount of water in a short period to avoid spill. The Forecast and Perfect policies have two advantages in this situation: they reduce the maximum release, and they also reduce the ramping rate. Both features highlight the ability of the forecast-informed policies to maintain a higher storage level without increasing flood risk.

Policy Parameters

The optimized policy parameters are shown in Fig. 6. The Forecast policy contains all three parameters, while the others include only the TOCS multiplier. In the case of the Perfect policy, this is because the ensemble is collapsed to a single member, and the risk curve parameters x_1 and x_2 do not apply (i.e., the Perfect policy never allows the forecasted storage to exceed the TOCS). The Forecast and Perfect policies can raise the TOCS value from the Baseline without an increase in flood risk. A higher TOCS multiplier does not guarantee higher mean storage, as evidenced by Fig. 4(a). In some instances, the Forecast TOCS multiplier is greater than the Perfect TOCS multiplier, yet the Perfect policy still yields higher mean storage [Shasta (SHA), New Melones (NML), Exchequer (EXC), Lake Success (SCC)]. These optimized policies should be interpreted in the context of the hydrology during this relatively dry period. For example, the return period of the largest inflow at Oroville Reservoir during this period (February 2017) was only 33 years (Koskinas et al. 2019), and the other reservoirs experienced similar hydrologic conditions due to heavy precipitation in the months of January and February 2017 (CNRFC 2023a).

Factor Analysis

We analyzed whether certain reservoir characteristics are correlated with the surplus storage achieved by the Forecast policy. Fig. 7 compares the storage benefit to four attributes of each reservoir. In all panels, the y-axis shows the difference between the mean storage of the Forecast and Baseline policies, with a larger value indicating that the Forecast policy stores more water. The x-axes

of each panel in Fig. 7 show a different reservoir attribute: the ratio of the mean inflow to capacity [Fig. 7(a)], the ratio of maximum inflow to reservoir safe release [Fig. 7(b)], the ratio of reservoir safe release to capacity [Fig. 7(c)], and the NSE of forecasts at a 4-day lead time, which is approximately the time needed to release the Baseline flood pool volume for most of the reservoirs [Fig. 7(d)]. The data in Fig. 7(a) show no correlation, while Fig. 7(b) shows a positive relationship between storage benefit and the ratio of the peak inflow to the maximum safe release. This results from the fact that the Forecast policy can manage large inflows with much lower maximum releases than Baseline policies [Figs. 5(b) and 6], thereby reducing the flood buffer requirement and increasing available storage. This relationship has a correlation coefficient of $r = 0.55$ and is significant at the 95% level. Fig. 7(c) suggests that reservoirs with a lower ratio of safe release to capacity may achieve larger storage benefits under FIRO because these reservoirs cannot release water as quickly relative to their capacity. However, the outliers make this conclusion difficult to state with confidence ($r = -0.33$ and insignificant). Finally, Fig. 7(d) shows a small and insignificant positive relationship ($r = 0.25$) between the storage benefit and forecast skill, represented here by the NSE at a 4-day lead time. Other lead times showed a similar noisy and weak relationship. This result suggests that while there may be some room for additional improvements in short-term forecast accuracy to support additional storage, improved forecast skill at this lead time does not immediately translate to higher achievable storage.

We also analyze whether the same reservoir characteristics shown in Fig. 7 are related to flood management benefits. Fig. 8 compares these four characteristics to the reduction in the frequency of maximum safe releases from the Baseline to the Forecast policy. A larger y-axis value indicates that the Forecast policy provides a larger reduction in flood risk relative to the Baseline. The x-axes of each panel in Fig. 8 are the same as in Fig. 7. Reservoirs that do not reach the safe release threshold at least 10 times in either the Baseline or Forecast policy are omitted from this analysis [SHA, Oroville (ORO), New Bullards Bar (BUL), Folsom (FOL),




Fig. 7. Surplus storage of the Forecast policy relative to Baseline (y-axis) compared to system characteristics (x-axis): (a) the ratio of the reservoir mean inflow to the reservoir capacity; (b) the ratio of the largest inflow the reservoir received to the reservoir safe release capacity; (c) the ratio of the reservoir safe release to the reservoir capacity; and (d) the NSE at 4 days lead time.

New Hogan (NHG)] because their small sample sizes bias the results. The data in Fig. 8(a) show a statistically significant negative relationship ($r = -0.79$) between the ratio of mean inflow to capacity and a reduction in flood risk. This suggests that reservoirs with larger storage relative to inflow are better able to buffer large flood events with the use of forecasts. Fig. 8(b) shows no correlation, while Fig. 8(c) depicts an insignificant negative relationship ($r = -0.42$) driven largely by one outlier site (SCC). Finally, Fig. 8(d) shows a weak and insignificant positive relationship ($r = 0.20$), which is consistent with the result for storage benefit [Fig. 7(d)]. Considering the omitted reservoirs and relatively weak correlations, it is difficult to conclude definitively that better forecasts would significantly improve policy performance. We performed a regression analysis using all the factors as predictors and the y-axes of Figs. 7 and 8 as the response (surplus storage, % reduction of max release frequency) and found similar results to the individual correlations; these results are included in Tables S1 and S2. Overall, the results of Figs. 7 and 8 highlight how the benefits of forecast information depend on complex dynamics of hydrology and release decisions at each reservoir over the study period, more so than can be predicted from these static factors alone.

Diagnostics

In Figs. 6(b and c), the risk curve parameters x_1 and x_2 do not have a discernible relationship with the storage results, likely because different combinations can achieve similar reductions in risk. To further analyze this point, Fig. 9 shows the mean normalized

storage and the accompanying parameter values for the 10 random seed trials of the Forecast policy optimization. The objective function values in the legend highlight that the differential evolution algorithm converges reliably to within 1% of the best objective value, indicating that the overall results are still similar between seeds. The policy from Seed 4 (star) was chosen for the results in Figs. 4–8 because it had the highest objective function value. A few of the reservoirs have a larger difference in storage across trials, including Pardee (PAR), Millerton (MIL), and Terminus (TRM) [Fig. 9(a)]. By contrast, SHA, ORO, and NML have little to no change in storage across trials. We find that while the trials converge to roughly the same objective function value, they do so with different policy parameters. This shows the different approaches to risk management that a policy could adopt to achieve the same goal: either a high TOCS with more risk-averse releases, or vice versa. Between these two outcomes we would expect a lower TOCS to be preferred by the operating agencies because it is a more modest change from the existing policy.

Finally, we performed a leave-one-out cross-validation (LOOCV) experiment to evaluate the optimized policy performance in out-of-sample conditions. We implemented this LOOCV on an annual basis across all 10 years of data, leaving one year out in each iteration. Table 1 gives the results of this experiment by reservoir.

In Table 1, we note that the mean storage averaged across the entire study period is similar between the LOOCV and fully trained results, with differences mostly within 1%–2% (see the last column). However, the LOOCV shows some instances of higher flood risk compared to the fully trained results for FIRO policies




Fig. 8. Reduction in frequency of maximum safe releases of the Forecast policy relative to Baseline (y-axis) compared to system characteristics (x-axis): (a) the ratio of the reservoir mean inflow to the reservoir capacity; (b) the ratio of the largest inflow the reservoir received to the reservoir safe release capacity; (c) the ratio of the reservoir safe release to the reservoir capacity; and (d) the NSE at 4 days lead time.




Fig. 9. Random seed optimization results: (a) mean normalized storage for the Forecast policy at each reservoir for 10 random seeds. As in (a), but for (b) x_1 (days with no allowable risk) at each reservoir for 10 random seeds; (c) x_2 (slope of risk curve) at each reservoir for 10 random seeds; and (d) TOCS multiplier at each reservoir for ten random seeds. In all instances, the star denotes the trial that was chosen for the results shown in Figs. 4–8.

Table 1. LOOCV results by reservoir


Reservoir	Days exceeding max release	Change in max release frequency training versus validation	Change in mean storage training versus validation (%)
SHA	0	2	-1.77
ORO	0	0	1.09
BUL	0	1	-1.30
FOL	0	0	0.27
PAR	7 (2017)	-4	-0.14
NHG	0	0	0.57
NML	0	2	0.31
DNP	0	1	0.00
EXC	1 (2023)	-13	0.39
MIL	0	-8	1.39
PNF	12 (2017)	-5	5.66
TRM	0	-1	1.14
SCC	0	0	1.86
ISB	0	0	-0.09

Note: Results are aggregated across left-out years for each reservoir and show the number of days (and out-of-sample years) when releases exceed the maximum allowable safe release; the number of days when releases equal the maximum allowable safe release, expressed as a difference between the LOOCV (validation) and fully trained results; and the percent change in mean storage between the LOOCV (validation) and fully trained results. SHA = Shasta; ORO = Oroville; BUL = New Bullards Bar; FOL = Folsom; PAR = Pardee; NHG = New Hogan; NML = New Melones; DNP = Don Pedro; EXC = Exchequer; MIL = Millerton; PNF = Pine Flat; TRM = Terminus; SCC = Lake Success; and ISB = Isabella.

at three sites (PAR, EXC, and PNF), in terms of the number of days when the maximum allowable safe release is exceeded (second column). Yet the LOOCV and fully trained results do not differ substantially in the total frequency of releases that equal the maximum release. Overall, the results of Table 1 highlight challenges when

using an optimized model to manage out-of-sample events. To maximize storage, the model finds the most aggressive combination of risk curve and TOCS multiplier possible on the training data. When exposed to larger, out-of-sample floods, it is unsurprising that some reservoir policies fail to maintain releases below the allowable maximum. It is notable, however, that this only occurs at three of the reservoirs, and only 20 days in total, with little difference in the frequency of releases that equal the allowable maximum. This suggests some degree of inherent robustness in the FIRO policies while also highlighting the need for careful design of optimization experiments.

In all cases of policy failure, the maximum safe release exceedances occur in a short time frame, during periods with multiple floods in short succession. Once the reservoir is full or close to capacity, the policy struggles to manage another incoming flood unless an event of similar magnitude has been previously seen in the training data. Fig. 10 shows an example of one such policy failure at PAR in 2017, when the reservoir had to draw down following a major flood peak in mid-February in anticipation of the second flood peak in late February, leading to releases that exceeded the allowable maximum threshold. Moreover, two of the reservoirs that failed the LOOCV (PAR and EXC) have a high ratio of largest inflow to safe release [Fig. 7(b)]. Accounting for this information in the development of FIRO policies may be important for managing flood risk at sites with limited saved forecasts for policy training. We test this concept by performing a second leave-one-out cross-validation experiment in which the upper bound of the TOCS multiplier was scaled based on the hydrologic record from 1994 to 2013 (i.e., when forecasts are not available for all sites). The scaling factor was created by dividing the largest 3-day inflow sum in the training period by the largest 3-day inflow sum in the historical period. With this approach to reduce the TOCS upper bound, no reservoir exceeded the maximum safe release in cross-validation, and we observe only a 4% reduction in mean storage per reservoir.

**Fig. 10.** Pardee (a) storage; and (b) release time series in 2017, with a FIRO policy fit to data excluding 2017.

Discussion and Conclusions

This study developed an optimization framework to generalize the risk-based EFO approach (Delaney et al. 2020; Woodside et al. 2022). We quantify the storage and flood control benefits of this approach in a multireservoir system, using a differential evolution algorithm to find optimal policy parameters for each of the 14 reservoirs in the Sacramento, San Joaquin, and Tulare river basins. The results show an 8.1% increase in water supply storage normalized by capacity, averaged across reservoirs and over the period 2013–2023, as well as a 10.9% increase in median storage at the end of the flood management season (May 1). This is lower than the 33% increase in median storage at the end of the flood management season found in a prior study of Lake Mendocino (Delaney et al. 2020), which had a longer and wetter study period on average (1985–2010). The Lake Mendocino experiment showed that flood risk could be successfully managed using forecasts while raising the TOCS limit. Our findings agree that much of the storage benefit in this simulation period comes from raising the TOCS value, and that the forecast information provides additional flood control benefit in terms of maintaining lower overall release magnitudes. We found that perfect forecasts further improve performance, particularly with respect to reducing the frequency of large releases. However, the room for improvement may be limited at some locations because existing forecasts appear to be sufficiently accurate to capture most of the potential benefit of the perfect forecast policy.

The results also show that many different combinations of policy parameters can be used to achieve similar levels of water supply benefits and flood risk, providing some flexibility in interpretation and implementation. Among these, policies with a lower TOCS multiplier and/or longer periods of zero allowable risk (i.e., the x_1 parameter) may be preferable for risk-averse operators. However, this equifinality obscures the degree to which these policies may transfer to new scenarios. The study also highlights reservoir characteristics that correlate with the benefit of FIRO policies. While many characteristics do not show significant correlation benefits, the ratio of the largest inflow to the maximum safe release and the ratio of mean inflow to capacity have significant relationships to surplus storage and reductions in the frequency of maximum safe releases, respectively. This finding suggests that certain site attributes may be useful for screening and selection of candidate reservoirs for implementation, even if such attributes alone cannot completely predict the benefits of revised operating policies. The results of the LOOCV experiment underscore that a longer training period is needed to ensure the selected FIRO policies do not increase flood risk. However, these results also show that optimized FIRO policies can still adapt to out-of-sample flood events in many cases (11 of the 14 reservoirs). Despite this promise, some restrictions on flood pool size must ultimately be implemented to avoid these infrequent but severe failures, especially at reservoirs with small safe release rates.

Several avenues for future work could advance the findings of this study. First, the length of the experiment is short and does not include a large variation in hydrology. The dry conditions during our 2013–2023 study provide little opportunity for the forecasts to provide value because most of the increased storage comes from the opportunity to raise the TOCS, and most reservoirs faced few large flood events. However, with a few exceptions (ORO, BUL, NHG), the majority of sites in this study do not have longer hindcast data sets required to expand the study period. A potential solution to this issue is the use of synthetic forecasts (Brodeur et al. 2024; Brodeur and Steinschneider 2021; Lamontagne and Stedinger 2018), which can be generated anywhere observational data are available. Synthetic forecast samples could be used in model optimization to prevent overfitting and produce more robust policies (Brodeur et al.

2020; Nayak et al. 2018). They could then be used to assess optimized policies against out-of-sample data containing a much richer set of hydrologic extremes spanning the observed historical record. This would allow for a more thorough analysis of FIRO-based benefits and risks and would provide higher confidence in the conclusions of this study, such as the aggregated economic benefits of FIRO. In addition, we expect that a longer and more variable hydrology (and associated synthetic forecasts) would likely show even greater benefits for the FIRO strategies than we found in this study.

One important limitation of our study was the use of an empirical approach to estimate reservoir maximum release rates and ramping rates from the observed record of releases. These details are specific to the spillway infrastructure and downstream conditions of each reservoir, and these assumptions can influence the model results. Another limitation is that the CNRFC forecasts are saved forecasts, not hindcasts, and therefore they do not reflect the most current methods and technologies. One practical aspect of water management not addressed in our study is the consideration of environmental effects, although these are represented implicitly through the use of historical releases to estimate demands, as well as constraining releases by ramping rates. Exploring alternative risk curve shapes could also prove fruitful, e.g., a quadratic or exponential risk curve may provide better performance than the piecewise linear curve assessed in this work. However, more complex risk curves have a higher likelihood of being overfit to training data, which would need to be evaluated in any future work. Moreover, future work could investigate the applicability of FIRO to a multireservoir system with shared downstream objectives, where joint release decisions and network routing would need to be considered to provide a more complete picture of systemwide benefit.

Data Availability Statement

All data, models, or code generated or used during the study are available in a repository online in accordance with funder data retention policies: <https://github.com/wltaylor/SSJRB-FIRO-Publishing>.

Acknowledgments

This work was partially supported by the National Science Foundation Grant 2205239. William Taylor also received support from the US Air Force. All conclusions are those of the authors.

Supplemental Materials

Figs. S1–S10 are available online in the ASCE Library (www.ascelibrary.org).

References

- AMS (American Meteorological Society). 2022. "Atmospheric river." Accessed October 16, 2023. https://glossary.ametsoc.org/wiki/Atmospheric_river.
- Badrinath, A., L. D. Monache, N. Hayatbini, W. Chapman, F. Cannon, and M. Ralph. 2023. "Improving precipitation forecasts with convolutional neural networks." *Weather Forecast.* 38 (2): 291–306. <https://doi.org/10.1175/WAF-D-22-0002.1>.
- Brodeur, Z. P., C. Delaney, B. Whitin, and S. Steinschneider. 2024. "Synthetic forecast ensembles for evaluating forecast informed reservoir operations." *Water Resour. Res.* 60 (2): e2023WR034898. <https://doi.org/10.1029/2023WR034898>.

- Brodeur, Z. P., J. D. Herman, and S. Steinschneider. 2020. "Bootstrap aggregation and cross-validation methods to reduce overfitting in reservoir control policy search." *Water Resour. Res.* 56 (Jun): e2020WR027184. <https://doi.org/10.1029/2020WR027184>.
- Brodeur, Z. P., and S. Steinschneider. 2021. "A multivariate approach to generate synthetic short-to-medium range hydro-meteorological forecasts across locations, variables, and lead times." *Water Resour. Res.* 57 (6): e2020WR029453. <https://doi.org/10.1029/2020WR029453>.
- California Department of Water Resources. 2023. *California water plan 2023 update*. Sacramento, CA: California Department of Water Resources.
- Chapman, W. E., A. C. Subramanian, L. Delle Monache, S. P. Xie, and F. M. Ralph. 2019. "Improving atmospheric river forecasts with machine learning." *Geophys. Res. Lett.* 46 (17–18): 10627–10635. <https://doi.org/10.1029/2019GL083662>.
- CNRFC (California Nevada River Forecast Center). 2023a. "Heavy precipitation events: California and northern Nevada: January and February 2017." Accessed August 24, 2023. https://www.cnrfc.noaa.gov/storm_summaries/janfeb2017storms.php.
- CNRFC (California Nevada River Forecast Center). 2023b. "Long range daily ensemble CSV file download." Accessed August 29, 2023. <https://www.cnrfc.noaa.gov/ensembleProductCSV.php>.
- Cohen, J. S., H. B. Zeff, and J. D. Herman. 2020. "Adaptation of multi-objective reservoir operations to snowpack decline in the Western United States." *J. Water Resour. Plann. Manage.* 146 (12): 04020091. [https://doi.org/10.1061/\(ASCE\)WR.1943-5452.0001300](https://doi.org/10.1061/(ASCE)WR.1943-5452.0001300).
- Corringham, T. W., J. McCarthy, T. Shulgina, A. Gershunov, D. R. Cayan, and F. M. Ralph. 2022. "Climate change contributions to future atmospheric river flood damages in the western United States." *Sci. Rep.* 12 (1): 13747. <https://doi.org/10.1038/s41598-022-15474-2>.
- Das, T., M. D. Dettinger, D. R. Cayan, and H. G. Hidalgo. 2011. "Potential increase in floods in California's Sierra Nevada under future climate projections." *Clim. Change* 109 (Dec): 71–94. <https://doi.org/10.1007/s10584-011-0298-z>.
- Delaney, C. J., et al. 2020. "Forecast informed reservoir operations using ensemble streamflow predictions for a multipurpose reservoir in Northern California." *Water Resour. Res.* 56 (9): e2019WR026604. <https://doi.org/10.1029/2019WR026604>.
- Dettinger, M. D., F. M. Ralph, T. Das, P. J. Neiman, and D. R. Cayan. 2011. "Atmospheric rivers, floods and the water resources of California." *Water* 3 (2): 445–478. <https://doi.org/10.3390/w3020445>.
- Doering, K., J. Quinn, P. M. Reed, and S. Steinschneider. 2021. "Diagnosing the time-varying value of forecasts in multiobjective reservoir control." *J. Water Resour. Plann. Manage.* 147 (7): 04021031. [https://doi.org/10.1061/\(ASCE\)WR.1943-5452.0001386](https://doi.org/10.1061/(ASCE)WR.1943-5452.0001386).
- Feng, Z., L. R. Leung, S. Hagos, R. A. Houze, C. D. Burleyson, and K. Balaguru. 2016. "More frequent intense and long-lived storms dominate the springtime trend in central US rainfall." *Nat. Commun.* 7 (1): 13429. <https://doi.org/10.1038/ncomms13429>.
- Gupta, R. S., A. L. Hamilton, P. M. Reed, and G. W. Characklis. 2020. "Can modern multi-objective evolutionary algorithms discover high-dimensional financial risk portfolio tradeoffs for snow-dominated water-energy systems?" *Adv. Water Resour.* 145 (Nov): 103718. <https://doi.org/10.1016/j.advwatres.2020.103718>.
- Hejazi, M. I., X. Cai, and B. L. Ruddell. 2008. "The role of hydrologic information in reservoir operation—Learning from historical releases." *Adv. Water Resour.* 31 (12): 1636–1650. <https://doi.org/10.1016/j.advwatres.2008.07.013>.
- Jager, H. I., and B. T. Smith. 2008. "Sustainable reservoir operation: Can we generate hydropower and preserve ecosystem values?" *River Res. Appl.* 24 (Jun): 340–352. <https://doi.org/10.1002/rra.1069>.
- Jasperse, J., et al. 2020. "Lake Mendocino forecast informed reservoir operations final viability assessment" Accessed October 4, 2023. <https://escholarship.org/uc/item/3b63q04n>.
- Koskinas, A., A. Tegos, P. Tsira, P. Dimitriadis, T. Iliopoulou, P. Papanicolaou, D. Koutsoyiannis, and T. Williamson. 2019. "Insights into the Oroville dam 2017 spillway incident." *Geosciences* 9 (Jan): 37. <https://doi.org/10.3390/geosciences9010037>.
- Lamontagne, J. R., and J. R. Stedinger. 2018. "Generating synthetic stream-flow forecasts with specified precision." *J. Water Resour. Plann. Manage.* 144 (4): 04018007. [https://doi.org/10.1061/\(ASCE\)WR.1943-5452.0000915](https://doi.org/10.1061/(ASCE)WR.1943-5452.0000915).
- Lavers, D. A., M. J. Rodwell, D. S. Richardson, F. M. Ralph, J. D. Doyle, C. A. Reynolds, V. Tallapragada, and F. Pappenberger. 2018. "The gauging and modeling of rivers in the sky." *Geophys. Res. Lett.* 45 (15): 7828–7834. <https://doi.org/10.1029/2018GL079019>.
- McCuen, R. H., Z. Knight, and A. G. Cutter. 2006. "Evaluation of the Nash–Sutcliffe efficiency index." *J. Hydrol. Eng.* 11 (6): 597–602. [https://doi.org/10.1061/\(ASCE\)1084-0699\(2006\)11:6\(597\)](https://doi.org/10.1061/(ASCE)1084-0699(2006)11:6(597)).
- Nayak, M. A., J. D. Herman, and S. Steinschneider. 2018. "Balancing flood risk and water supply in California: Policy search integrating short-term forecast ensembles with conjunctive use." *Water Resour. Res.* 54 (10): 7557–7576. <https://doi.org/10.1029/2018WR023177>.
- Pan, B., K. Hsu, A. AghaKouchak, S. Sorooshian, and W. Higgins. 2019. "Precipitation prediction skill for the West Coast United States: From short to extended range." *J. Clim.* 32 (1): 161–182. <https://doi.org/10.1175/JCLI-D-18-0355.1>.
- Pathak, T. B., M. L. Maskey, J. A. Dahlberg, F. Kearns, K. M. Bali, and D. Zaccaria. 2018. "Climate change trends and impacts on California agriculture: A detailed review." *Agronomy* 8 (3): 25. <https://doi.org/10.3390/agronomy8030025>.
- Peel, M. C., B. L. Finlayson, and T. A. McMahon. 2007. "Updated world map of the Köppen–Geiger climate classification." *Hydrol. Earth Syst. Sci.* 11 (Jun): 1633–1644. <https://doi.org/10.5194/hess-11-1633-2007>.
- Reed, P. M., D. Hadka, J. D. Herman, J. R. Kasprzyk, and J. B. Kollat. 2013. "Evolutionary multiobjective optimization in water resources: The past, present, and future." *Adv. Water Resour.* 51 (Jan): 438–456. <https://doi.org/10.1016/j.advwatres.2012.01.005>.
- Semmendinger, K., and S. Steinschneider. 2024. "Influence of subseasonal-to-annual water supply forecasts on many-objective water system robustness under long-term change." *J. Water Resour. Plann. Manage.* 150 (5): 04024009. <https://doi.org/10.1061/JWRMD5.WRENG-6205>.
- Steinschneider, S., J. D. Herman, J. Kucharski, M. Abellera, and P. Ruggiero. 2023. "Uncertainty decomposition to understand the influence of water systems model error in climate vulnerability assessments." *Water Resour. Res.* 59 (1): e2022WR032349. <https://doi.org/10.1029/2022WR032349>.
- Storn, R., and K. Price. 1997. "Differential evolution—A simple and efficient heuristic for global optimization over continuous spaces." *J. Global Optim.* 11 (Dec): 341–359. <https://doi.org/10.1023/A:1008202821328>.
- Swain, D. L., B. Langenbrunner, J. D. Neelin, and A. Hall. 2018. "Increasing precipitation volatility in twenty-first-century California." *Nat. Clim. Change* 8 (5): 427–433. <https://doi.org/10.1038/s41558-018-0140-y>.
- Turner, S. W. D., J. C. Bennett, D. E. Robertson, and S. Galelli. 2017. "Complex relationship between seasonal streamflow forecast skill and value in reservoir operations." *Hydrol. Earth Syst. Sci.* 21 (Apr): 4841–4859. <https://doi.org/10.5194/hess-21-4841-2017>.
- Tustison, B. 2020. "Geographical trends in central valley flood control reservoirs." Accessed October 10, 2023. <https://blog.mbkapps.com/posts/res-prop/>.
- US Army Corps of Engineers. 2017. *Management of water control systems*. Rep. No. EM 1110-2-3600. Washington, DC: US Army Corps of Engineers.
- US Census Bureau. 2023. "QuickFacts: California." Accessed August 29, 2023. <https://www.census.gov/quickfacts/fact/table/CA/PST045222>.
- USGS. 2023. "California's Central Valley." Accessed August 21, 2023. <https://ca.water.usgs.gov/projects/central-valley/about-central-valley.html>.
- Virtanen, P., et al. 2020. "SciPy 1.0: Fundamental algorithms for scientific computing in Python." *Nat. Methods* 17 (Jun): 261–272. <https://doi.org/10.1038/s41592-019-0686-2>.
- Warner, M. D., C. F. Mass, and E. P. Salathé. 2015. "Changes in winter atmospheric rivers along the North American West Coast in CMIP5 climate models." *J. Hydrometeorol.* 16 (1): 118–128. <https://doi.org/10.1175/JHM-D-14-0080.1>.
- Woodside, G. D., A. S. Hutchinson, F. M. Ralph, C. Talbot, R. Hartman, and C. Delaney. 2022. "Increasing stormwater capture and recharge

- using forecast informed reservoir operations.” *Prado Dam. Groundwater* 60 (5): 634–640. <https://doi.org/10.1111/gwat.13162>.
- Yang, G., S. Guo, P. Liu, and P. Block. 2020. “Integration and evaluation of forecast-informed multiobjective reservoir operations.” *J. Water Resour. Plann. Manage.* 146 (6): 04020038. [https://doi.org/10.1061/\(ASCE\)WR.1943-5452.0001229](https://doi.org/10.1061/(ASCE)WR.1943-5452.0001229).
- Zatarain Salazar, J., P. M. Reed, J. D. Herman, M. Giuliani, and A. Castelletti. 2016. “A diagnostic assessment of evolutionary algorithms for multi-objective surface water reservoir control.” *Adv. Water Resour.* 92 (Jun): 172–185. <https://doi.org/10.1016/j.advwatres.2016.04.006>.
- Zeff, H. B., A. L. Hamilton, K. Malek, J. D. Herman, J. S. Cohen, J. Medellin-Azuara, P. M. Reed, and G. W. Characklis. 2021. “California’s food-energy-water system: An open source simulation model of adaptive surface and groundwater management in the Central Valley.” *Environ. Modell. Software* 141 (Jul): 105052. <https://doi.org/10.1016/j.envsoft.2021.105052>.
- Zimmerman, J. K. H., D. M. Carlisle, J. T. May, K. R. Klausmeyer, T. E. Grantham, L. R. Brown, and J. K. Howard. 2018. “Patterns and magnitude of flow alteration in California, USA.” *Freshwater Biol.* 63 (Sep): 859–873. <https://doi.org/10.1111/fw.13058>.