

1 **From Stream Flows to Cash Flows: Leveraging**
2 **Evolutionary Multi-Objective Direct Policy Search to**
3 **Manage Hydrologic Financial Risks**

4 **Andrew L. Hamilton^{1,2}, Gregory W. Characklis^{1,2}, and Patrick M. Reed³**

5 ¹Department of Environmental Sciences and Engineering, Gillings School of Global Public Health,
6 University of North Carolina at Chapel Hill, Chapel Hill, NC, USA

7 ²Center on Financial Risk in Environmental Systems, Gillings School of Global Public Health, UNC
8 Institute for the Environment, University of North Carolina at Chapel Hill, Chapel Hill, NC, USA

9 ³Department of Civil and Environmental Engineering, Cornell University, Ithaca, NY, USA

10 **Key Points:**

- 11 • Reservoir control and financial risk management share a common multi-objective
12 decision structure and can be optimized using similar methods
- 13 • Evolutionary Multi-Objective Direct Policy Search (EMODPS) is used to develop
14 financial risk management policies for a hydropower producer
- 15 • Information theoretic sensitivity analysis and visual analytics are used to build
16 intuition about how policies adapt to changing conditions

Corresponding author: Andrew L. Hamilton, andrew.hamilton@unc.edu

Abstract

Hydrologic variability can present severe financial challenges for organizations that rely on water for the provision of services, such as water utilities and hydropower producers. While recent decades have seen rapid growth in decision-support innovations aimed at helping utilities manage hydrologic uncertainty for multiple objectives, support for managing the related financial risks remains limited. However, the mathematical similarities between multi-objective reservoir control and financial risk management suggest that the two problems can be approached in a similar manner. This paper demonstrates the utility of Evolutionary Multi-Objective Direct Policy Search (EMODPS) for developing adaptive policies for managing the drought-related financial risk faced by a hydropower producer. These policies dynamically balance a portfolio, consisting of snowpack-based financial hedging contracts, cash reserves, and debt, based on evolving system conditions. Performance is quantified based on four conflicting objectives, representing the classic tradeoff between “risk” and “return” in addition to decision-makers’ unique preferences towards different risk management instruments. The dynamic policies identified here significantly outperform static management formulations that are more typically employed for financial risk applications in the water resources literature. Additionally, this paper combines visual analytics and information theoretic sensitivity analysis to help decision-makers better understand how different candidate policies achieve their comparative advantages through differences in how they adapt to real-time information. The methodology developed in this paper should be applicable to any organization subject to financial risk stemming from hydrology or other environmental variables (e.g., wind speed, insolation), including electric utilities, water utilities, agricultural producers, and renewable energy developers.

Keywords

hydropower, water resources, financial risk, direct policy search, reservoir control, global sensitivity analysis

1 Introduction

Reservoir control and financial risk management share strong similarities. The principal task in each is to reduce the risk of negative impacts from variable inflows (either hydrologic flows or cash flows), through the use of a buffer stock (either a reservoir or

48 a reserve fund) that is filled in times of abundance and drawn down in times of scarcity
 49 (Figure 1). Other risk management tools may also be used to limit the impact of low-
 50 flow periods, but at a cost (e.g., water desalination or demand management for stream-
 51 flow deficits, and borrowing or financial hedging for cash flow deficits). In both cases,
 52 the manager must make decisions under an array of uncertainties, and may need to nav-
 53 igate tradeoffs between conflicting objectives (e.g., flood control vs. water supply for reser-
 54 voir control, risk vs. cost for financial risk management). And in both cases, as systems
 55 dynamically evolve, managers will have to adapt to new information as it becomes avail-
 56 able. In other words, reservoir control and financial risk management can be formulated
 57 as very similar Markov Decision Processes (MDPs) (Bertsekas, 2019; Powell, 2019), whether
 58 managers attempt to solve this problem explicitly, using programmatic approaches such
 59 as stochastic dynamic programming, or implicitly, relying on expert specified rules. Ad-
 60 ditionally, reservoir control and financial risk management are strongly interdependent
 61 activities for water-reliant organizations in the Food-Energy-Water Nexus, such as hy-
 62 dropower producers, municipal water utilities, and irrigation districts (Cai, Wallington,
 63 Shafiee-Jood, & Marston, 2018; D’Odorico et al., 2018; Scanlon et al., 2017). Such or-
 64 ganizations rely on water for the provision of services, and as a result, their revenues and/or
 65 costs can be highly dependent on hydrologic inflows (Blomfield & Plummer, 2014; Lar-
 66 son, Freedman, Passinsky, Grubb, & Adriaens, 2012). This suggests that an understand-
 67 ing of complex water resource system dynamics can be used to better characterize and
 68 adaptively manage financial risks borne by water-reliant organizations.

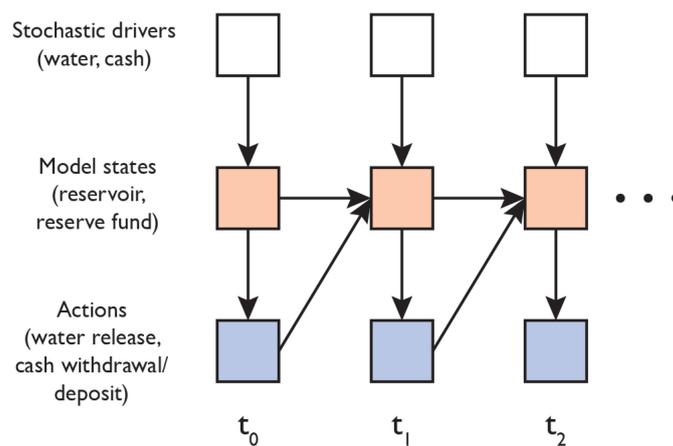


Figure 1. A simple reservoir model and a simple cash flow model share the same underlying decision structure.

69 Water resource systems researchers have developed a broad range of strategies for
70 dynamically managing reservoir operations in the face of uncertain hydrometeorology
71 and demands (see reviews by Castelletti, Pianosi, and Soncini-Sessa (2008); Labadie (2004);
72 Macian-Sorribes and Pulido-Velazquez (2019); Yeh (1985)), but Stochastic Dynamic Pro-
73 gramming (SDP) and its many derivatives have been the most popular. The problem
74 is formulated as an MDP in which a decision-maker must make sequential decisions based
75 on the stochastically evolving state of the system. Each action affects the immediate cost/reward
76 as well as the future state of the system. In SDP, this recursion is used to find optimal
77 operating rules, in the form of a discrete policy table, using the Bellman Equation (Bell-
78 man, 1957). However, despite its widespread use, SDP suffers from a number of limita-
79 tions that reduce its applicability to large, complex, multi-objective problems where op-
80 erations are evaluated using stochastic simulations (see discussion in Giuliani, Castel-
81 letti, Pianosi, Mason, and Reed (2016)).

82 A variety of approximation methods have been developed to overcome these chal-
83 lenges, such as approximate dynamic programming, reinforcement learning, and model
84 predictive control (Bertsekas, 2019). Direct Policy Search (DPS) (Rosenstein & Barto,
85 2001), or parameterization-simulation-optimization (Koutsoyiannis & Economou, 2003),
86 has become increasingly popular in the field of water resources systems analysis (Macian-
87 Sorribes & Pulido-Velazquez, 2019). DPS is an approximation in policy space (Powell,
88 2019), wherein the optimal operating policy is assumed to lie in the space of a certain
89 parametric family of functions, and the policy parameters are optimized rather than the
90 decisions themselves (i.e., optimizing state-aware adaptive rule systems instead of spe-
91 cific actions). This drastically reduces the “curse of dimensionality” that limits the tractabil-
92 ity of large SDP problems. Additionally, DPS allows for “model-free” representation of
93 stochastic inputs, meaning that observational data, synthetically generated data, and
94 process-based simulation model output can all be used in lieu of explicit probability dis-
95 tributions (Desreumaux, Côté, & Leconte, 2018; Giuliani, Quinn, Herman, Castelletti,
96 & Reed, 2018). A simulation-based approach to optimization also allows for flexible con-
97 struction of mixed multi-objective formulations (Giuliani et al., 2016; Kasprzyk, Reed,
98 & Hadka, 2016; Quinn, Reed, & Keller, 2017). In Evolutionary Multi-Objective Direct
99 Policy Search (EMODPS) (Giuliani, Herman, Castelletti, & Reed, 2014), the policies are
100 parameterized with a non-linear approximating network and optimized using a multi-
101 objective evolutionary algorithm (MOEA). EMODPS has been deployed to solve com-

plex reservoir operations problems (multiple reservoirs; multiple, mixed objectives; and model-free information) that would be untenable using a traditional SDP approach (Denaro, Anghileri, Giuliani, & Castelletti, 2017; Giuliani, Pianosi, & Castelletti, 2015; Quinn et al., 2018; Zatarain Salazar, Reed, Quinn, Giuliani, & Castelletti, 2017).

To complement algorithmic search strategies, water resources researchers have developed an assortment of computational tools to help decision-makers better understand their options. This is especially important in multi-objective contexts, where optimization results in a multitude of solutions representing the optimal tradeoffs between conflicting objectives (the Pareto set), rather than a single “best” policy. As the dimensionality of the Pareto set grows, it becomes increasingly difficult to conceptualize. High-dimensional visualization, solution brushing, and other visual analytic techniques can help decision-makers to better understand the complex tradeoffs in their system and choose the solution that best suits their needs (Herman, Zeff, Reed, & Characklis, 2014; Huskova, Matrosov, Harou, Kasprzyk, & Lambert, 2016; Kollat & Reed, 2007). These tools can also help decision-makers to refine their conceptualization of the problem through iterative reformulation (Castelletti & Soncini-Sessa, 2006; Giuliani, Herman, et al., 2014; Kasprzyk, Reed, Characklis, & Kirsch, 2012). Visual analytics are especially powerful when combined with global sensitivity analyses that probe the impacts of key uncertainties on system performance (Iooss & Lemaitre, 2015; Pianosi et al., 2016; Saltelli, Tarantola, & Campolongo, 2000). These tools can be used to “open the black box” of non-linear approximating networks and help decision-makers to better understand how the optimal operating policies adapt to changing conditions (Quinn, Reed, Giuliani, & Castelletti, 2019). In this way, visual analytics and sensitivity analysis can help to build trust between water resources modelers and real-world stakeholders. Although water resources practitioners in general have been slow to adopt computational decision support tools such as MOEAs, visual analytics, and global sensitivity analysis (Basdekas, 2014; Brown et al., 2015), a growing number of real-world use cases suggests that this may be changing (Basdekas & Hayslett, 2021; Moallemi, Kwakkel, de Haan, & Bryan, 2020; Smith, Kasprzyk, & Dilling, 2019; Wild, Reed, Loucks, Mallen-Cooper, & Jensen, 2019; Wu et al., 2016).

Many organizations such as water utilities and hydropower producers rely on water for the provision of services. During drought, these organizations can experience reduced revenues and/or increased costs (Hughes et al., 2014; Larson et al., 2012). For example, an electric utility with reduced hydropower capacity during drought will have less

135 electricity to sell (reduced revenues) and/or be forced to purchase more expensive replace-
136 ment power from other generators (increased costs). Similarly, a water utility experienc-
137 ing supply shortfalls will typically implement demand management measures (reduced
138 revenues) and/or water purchases from other utilities or irrigators (increased costs). These
139 measures can result in severe cash flow deficits that leave an organization at risk of de-
140 faulting on its obligations (e.g., debt service, operations and maintenance) (Ceres, 2017;
141 Leurig, 2010). Water utilities and hydropower-reliant electric utilities are therefore vul-
142 nerable to significant financial disruption during drought, and hydrologic financial risk
143 can have an outsized impact on the long-term viability of the utility; indeed, credit rat-
144 ing agencies have noted that the ability to manage the financial impacts of drought is
145 an important factor in determining a utility's creditworthiness (Chapman & Breeding,
146 2014; Moody's Investors Service, 2011, 2019). Tools such as reserve funds, financial hedg-
147 ing contracts, and lines of credit can be used to reduce the variability of net cash flows.
148 This, in turn, can reduce an organization's likelihood of bankruptcy, improve its credit
149 rating, and reduce its future borrowing costs (Bank & Wiesner, 2010; Pérez-González
150 & Yun, 2013), in addition to helping risk-averse staff feel more comfortable (Bodnar, Gi-
151 ambona, Graham, & Harvey, 2019; Krause & Tse, 2016). Most utilities rely heavily on
152 debt to finance infrastructure projects (Hughes & Leurig, 2013), so financial risk man-
153 agement is a key component of providing quality service at affordable rates.

154 Despite the critical role of financial risk management in water resources, decision
155 support for practitioners in this area has remained limited. There is a long history of con-
156 sidering financial objectives such as expected revenues and costs in water resources sys-
157 tems analysis (e.g., see references in Labadie (2004); Macian-Sorribes and Pulido-Velazquez
158 (2019); Yeh (1985)). However, fewer studies have explicitly accounted for variability in
159 costs and revenues, or the financial risk management actions that an organization can
160 take to combat this variability. Those that do have tended to propose static, non-adaptive
161 management strategies. For example, modeling of financial reserves is not common in
162 the water resources literature, and the limited examples tend to assume that the util-
163 ity will contribute either a fixed amount or a fixed fraction of revenues to the reserve fund
164 each year (Rehan, Knight, Unger, & Haas, 2013; Rehan, Unger, Knight, & Haas, 2015;
165 Zeff, Kasprzyk, Herman, Reed, & Characklis, 2014). Similarly, there is a growing inter-
166 est in using hydrology-based financial hedging contracts in applications such as hydropower
167 (Foster, Kern, & Characklis, 2015; Hamilton, Characklis, & Reed, 2020; Meyer, Charack-

168 lis, Brown, & Moody, 2016), water supply (Brown & Carriquiry, 2007; Maestro, Barnett,
169 Coble, Garrido, & Bielza, 2016; Zeff & Characklis, 2013), and agriculture (Denaro, Castel-
170 letti, Giuliani, & Characklis, 2020; Mortensen & Block, 2018; Turvey, 2001), but researchers
171 have generally assumed that the same contract is purchased each year, not allowing for
172 risk management to be adjusted over time as conditions change.

173 However, financial researchers have demonstrated that adaptive, state-aware ac-
174 tion is crucial to financial risk management (Bolton, Chen, & Wang, 2011; Disatnik, Duchin,
175 & Schmidt, 2014; Froot, Scharfstein, & Stein, 1993; Rampini, Sufi, & Viswanathan, 2014).
176 Just as a reservoir operator should consider current reservoir levels and expected future
177 inflows when making release decisions, so should a financial risk manager consider the
178 utility's current bank account balance and projected future revenues and costs when de-
179 ciding whether to withdraw money from the bank, or whether to hedge its drought ex-
180 posure using index contracts. A variety of optimization methods have been applied to
181 financial problems such as investment portfolio selection (Markowitz, 1952; Mulvey, 2001;
182 Pardalos, Sandström, & Zopounidis, 1994), asset-liability management (Kouwenberg &
183 Zenios, 2008; Sodhi, 2005), and cash flow management (Baumol, 1952; da Costa Moraes,
184 Nagano, & Sobreiro, 2015; Miller & Orr, 1966). As in water resources systems analysis,
185 some researchers have attempted to provide more realistic decision support using multi-
186 objective formulations (de Almeida-Filho, de Lima Silva, & Ferreira, 2020; Marqués, García,
187 & Sánchez, 2020; Salas-Molina, Pla-Santamaria, & Rodriguez-Aguilar, 2018; Zopouni-
188 dis, Galariotis, Doumpos, Sarri, & Andriosopoulos, 2015), model-free information (Sun,
189 Fang, Wu, Lai, & Xu, 2011), heuristic solution methods (Aguilar-Rivera, Valenzuela-Rendón,
190 & Rodríguez-Ortiz, 2015; da Costa Moraes & Nagano, 2013; Ponsich, Jaimes, & Coello Coello,
191 2013; Tapia & Coello Coello, 2007), and visual analytics (Flood, Lemieux, Varga, & William Wong,
192 2016; Savikhin, Lam, Fisher, & Ebert, 2011). Beyond the academic literature, the use
193 of quantitative decision support tools by financial firms (e.g., banks, hedge funds, insur-
194 ers) has proliferated in recent years, driven by growth in computing power, big data, al-
195 gorithms, and visualization software (Fabozzi, Focardi, & Jonas, 2007; Rundo, Trenta,
196 di Stallo, & Battiato, 2019; Zopounidis, Doumpos, & Niklis, 2018). However, these firms
197 generally employ proprietary and highly problem-specific technologies that are not read-
198 ily adoptable by organizations outside of the financial sector, such as water and power
199 utilities, which nevertheless face significant financial risks.

200 This paper bridges the gap between reservoir control and financial risk manage-
201 ment to show how computational tools developed for the former can be adapted to the
202 latter. This research builds on prior work by the authors dealing with drought-related
203 financial risk management by a hydropower producer. First, Hamilton et al. (2020) de-
204 veloped a hydro-financial simulation model that abstracts the hydroclimatology, hydropower
205 generation, cash flows, and financial risk management of the Power Enterprise of the San
206 Francisco Public Utilities Commission (SFPUC). The authors used this model to eval-
207 uate different static financial risk management portfolios within a Monte Carlo frame-
208 work and search for optimal portfolios using an MOEA. In related work, Gupta, Hamil-
209 ton, Reed, and Characklis (2020) introduced an adaptive EMODPS formulation of a sim-
210 plified financial risk management problem, which was used to diagnostically benchmark
211 if modern MOEAs are capable of addressing this new class of problem. The present study
212 builds on these prior works by contributing the most detailed and actionable represen-
213 tation to date of how EMODPS can be used to craft operating policies that adapt to chang-
214 ing conditions over time when managing drought-related financial risk. The advantages
215 of dynamic decision-making are demonstrated relative to a simplified static operating
216 policy akin to those commonly applied to financial risk management in the water resources
217 literature. This paper also demonstrates the value of higher-dimensional problem fram-
218 ings that explicitly account for decision-maker preferences with respect to the use of dif-
219 ferent management tools. Lastly, a framework is contributed for combining *a posteriori*
220 visual analytics with information theoretic sensitivity analysis (ITSA) in order to help
221 decision-makers better understand how complex, non-linear operating policies achieve
222 their goals by adapting to real-time information when making decisions.

223 **2 Study context**

224 **2.1 Study area**

225 San Francisco Public Utilities Commission (SFPUC) owns and operates three reser-
226 voirs (Hetch Hetchy Reservoir, Cherry Lake, and Lake Eleanor) in the upper Tuolumne
227 River basin in the Sierra Nevada mountains (Figure S1 in Supporting Information (SI)).
228 These reservoirs deliver drinking water to much of the San Francisco Bay area, and en
229 route, the water also provides hydroelectric power. SFPUC uses this hydropower to sell
230 retail electricity at fixed rates to San Francisco International Airport, municipal build-
231 ings in San Francisco, and a number of other retail customer classes within the Bay area.

232 Irrigation districts along the Tuolumne River also have the right to buy surplus hydropower,
 233 when available, at a fixed rate. When hydropower production is in excess of retail and
 234 irrigation district demands, it is sold at floating market rates into the Western Systems
 235 Power Pool (hereafter “wholesale market”). On the other hand, when hydropower is in-
 236 sufficient to meet the demand from retail customers, SFPUC is obligated to purchase
 237 the remainder on the wholesale market. Although SFPUC provides both water supply
 238 and power supply, they are operated as independent entities from a financial perspec-
 239 tive (San Francisco Public Utilities Commission, 2016), and the present work considers
 240 only the power supply enterprise.

241 **2.2 Hydro-financial simulation model**

242 This paper adopts the hydro-financial simulation model from Hamilton et al. (2020).
 243 The first component of the model is the stochastic engine, which is used to create a million-
 244 year synthetic record that can be used to drive the system. First, snow water equiva-
 245 lent depth (SWE) measurements for February 1 and April 1 (the months with the longest
 246 and most continuous datasets for the watershed) are randomly generated based on a cop-
 247 ulla model. Next, hydropower production is synthetically generated using piecewise lin-
 248 ear models for each month conditioned on SWE, combined with an autoregressive model
 249 for residual noise. Third, monthly wholesale power prices are synthetically generated us-
 250 ing a seasonal autoregressive moving average model. Lastly, monthly hydropower net rev-
 251 enues are calculated based on hydropower generation and power prices. Net revenues are
 252 defined as the total annual cash flow resulting from retail and wholesale hydropower sales,
 253 minus wholesale power purchases, minus the annual “fixed costs” (debt service payments,
 254 operations and maintenance, salaries, etc.) that must be paid each year. The synthetic
 255 records are found to closely match the historical record in terms of statistical proper-
 256 ties, while providing a wider sampling of possible outcomes than can be found in the lim-
 257 ited historical data. For more details on the methodology and validation of the stochas-
 258 tic engine, see Hamilton et al. (2020).

259 Three annual quantities are derived from this monthly synthetic dataset and used
 260 as stochastic drivers for the present study. Firstly, the SWE index (ε^S , in inches) is a
 261 weighted average of February and April SWE observations. The inflows to SFPUC’s reser-
 262 voirs are dominated by the seasonal dynamics of snow accumulation and melt, so SWE
 263 measurements taken upstream of the reservoirs in the late winter/early spring can be

264 used to predict the magnitude of streamflows during the melt period in the late spring/early
265 summer. A weighted average of February and April observations is found to improve cor-
266 relation with annual hydropower production, relative to either month in isolation, by in-
267 corporating information about the timing of snowfall and melt (Hamilton et al., 2020).
268 This correlation suggests that the index is a good candidate for financial hedging with
269 index contracts (see below). The second stochastic driver is total hydropower net rev-
270 enue over the water year (ε^R , in \$M). Lastly, the power price index (ε^P , in \$/MWh) is
271 defined as the expected value of the generation-weighted average wholesale power price
272 over the coming water year. This index takes advantage of autocorrelation in the mar-
273 ket to predict how favorable the wholesale power prices will be for the utility’s net hy-
274 dropower revenues over the coming water year. Although the correlation is relatively low
275 ($\rho = 0.35$, see SI Figure S2), the index still provides potentially valuable information
276 for making decisions regarding financial risk, and is used as one of the inputs to the dy-
277 namic control policies (Section 3.1.2). More details on ε^P can be found in SI Section S1.

278 Absent any financial risk management, the utility will experience years in which
279 costs outweigh revenues (i.e., net revenue is negative). This situation can be extremely
280 disruptive because the utility risks defaulting on its obligations (e.g., debt service or op-
281 erations and maintenance). The hydro-financial simulation model provides three tools
282 which can be used to avoid such negative outcomes. Firstly, it can purchase a snowpack-
283 based hedging contract called a capped Contract for Differences (CFD). The CFD (SI
284 Figure S3) provides payouts to the utility in low-SWE years (below 24.7 inches), when
285 it expects to have low hydropower and thus low revenue, in return for the utility mak-
286 ing payments in high-SWE years (above 24.7 inches), when the utility expects to have
287 abundant hydropower and surplus revenue. The negative correlation between hydropower
288 revenue and CFD payout has been found to significantly reduce the volatility of the com-
289 bined cash flow, suggesting its value as a financial risk management tool (Hamilton et
290 al., 2020). The second risk management tool is a reserve fund, into which the utility can
291 deposit surplus cash flows. This allows it to withdraw from the fund when hydropower
292 revenues are insufficient to pay its bills. Lastly, the utility has a letter of credit with a
293 bank, under which it can borrow money (i.e., issue short-term debt). The debt is paid
294 back each year (with interest), and is assumed to take up the slack in situations where
295 the other two tools fail to generate sufficient cash flows to avoid defaulting on the util-
296 ity’s obligations. Note that the short-term debt considered in this model is distinct from

297 longer-term debt service obligations related to past bond offerings, typically associated
 298 with infrastructure investments, and which are assumed to be part of the “fixed costs”
 299 above.

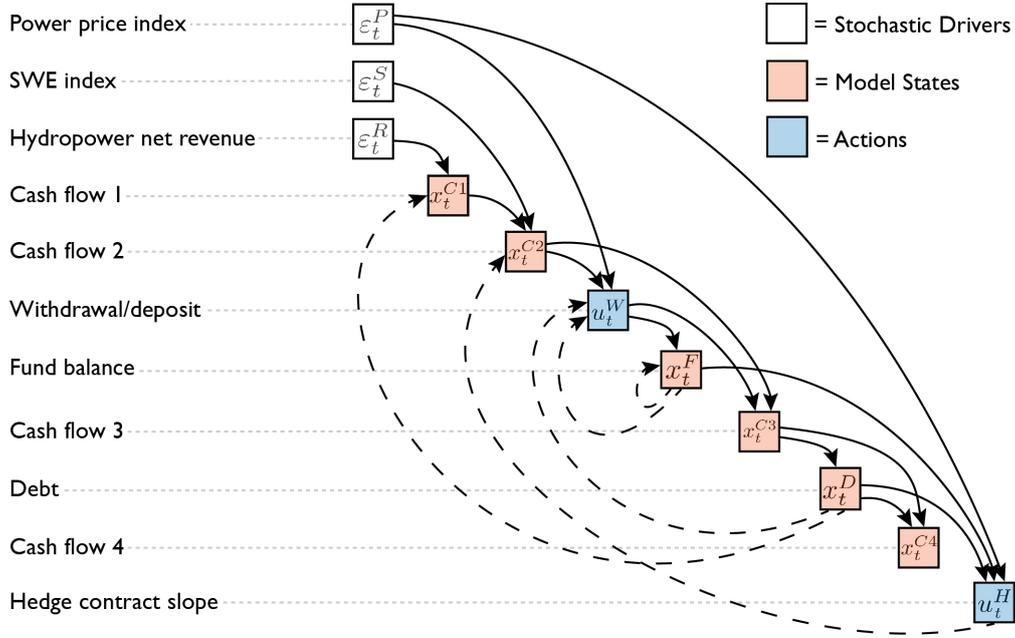


Figure 2. Annual sequence of operations in hydro-financial simulation model (moving from top left to bottom right). Solid (dashed) arrows represent the information flows from the current (previous) time step.

300 Figure 2 shows how these financial operations are abstracted in the hydro-financial
 301 simulation model (see Table 1 for a list of variable names, symbols, units, and constants).
 302 The sequence of operations occurs at the end of each water year, September 30, based
 303 on the stochastic outcomes that occur over the course of that water year, ε_t . Two state-
 304 aware “actions” each year are governed by the control policy (to be described in Section
 305 3.1): the amount of cash withdrawn from/deposited to the reserve fund (u_t^W , in \$M, where
 306 $u_t^W > 0$ represents a withdrawal and $u_t^W < 0$ represents a deposit), and the hedging
 307 contract slope (u_t^H , in \$M/inch of SWE). All other variables (“model states”) are au-

308 tomatically updated according to the following rules:

$$309 \quad x_t^{C1} = \varepsilon_t^R - r^D x_{t-1}^D \quad (1)$$

$$310 \quad x_t^{C2} = x_t^{C1} + u_{t-1}^H h(\varepsilon_t^S) \quad (2)$$

$$311 \quad x_t^F = r^F x_{t-1}^F - u_t^W \quad (3)$$

$$312 \quad x_t^{C3} = x_t^{C2} + u_t^W \quad (4)$$

$$313 \quad x_t^D = \max(-x_t^{C3}, 0) \quad (5)$$

$$314 \quad x_t^{C4} = x_t^{C3} + x_t^D \quad (6)$$

315 where x_t^{C1} , x_t^{C2} , and x_t^{C3} are intermediate cash flows and x_t^{C4} is the final cash flow in
 316 year t ; x_t^D and x_t^F are the short-term debt and reserve fund balance at the end of time
 317 step t ; r^D and r^F are the annual real interest rates on debt and reserves; and $h(\varepsilon_t^S)$ is
 318 the CFD payout function (SI Figure S3). This function converts the stochastic SWE in-
 319 dex value from the current year into a number of inches of SWE for which the utility will
 320 receive compensation (if $h(\varepsilon_t^S) > 0$) or owe payment (if $h(\varepsilon_t^S) < 0$). To get the util-
 321 ity's total payout received (or payment due), this output is multiplied by the CFD slope,
 322 u_{t-1}^H , as chosen by the control policy at the end of the previous year (Section 3.1). The
 323 reader is referred to Hamilton et al. (2020) for more details on construction of the CFD.

324 A full realization of the hydro-financial simulation model requires iterating this se-
 325 quence for $T = 20$ years, subject to a randomly sampled $(T+1)$ -year sequence of stochas-
 326 tic drivers. The multi-year simulation accounts for the path-dependent dynamics of the
 327 reserve fund and debt, as well as the autocorrelation within the stochastic power prices.
 328 The reserve fund and debt are assumed to be zero at $t = 0$ (in practice these values could
 329 be set based on circumstance). The hedging contract policy in year 0 (the slope to be
 330 used for the payout in year 1) is calculated using x_0^F , x_0^D , and ε_0^P .

331 **3 Methods**

332 Figure 3 shows how the stochastic engine and hydro-financial model are integrated
 333 into the broader framework of this study. The EMODPS methodology combines adap-
 334 tive control rules, Monte Carlo ensemble simulation, and MOEA-driven policy search.
 335 The search produces a large population of candidate policies, which can be explored us-
 336 ing optimal tradeoff analysis, many-objective visualization, and information theoretic sen-
 337 sitivity analysis. This framework is further described in what follows.

Table 1. Variables and constants for hydro-financial simulation model.

Variable	Symbol	Value	Units
Power price index	ε_t^P	-	\$/MWh
SWE index	ε_t^S	-	inches
Annual net revenue	ε_t^R	-	\$M
Cash flow 1	x_t^{C1}	-	\$M
Cash flow 2	x_t^{C2}	-	\$M
Withdrawal	u_t^W	-	\$M
Reserve fund balance	x_t^F	-	\$M
Cash flow 3	x_t^{C3}	-	\$M
Debt	x_t^D	-	\$M
Cash flow 4	x_t^{C4}	-	\$M
Hedge contract slope	u_t^H	-	\$/inch
Mean net revenue before risk management	\bar{R}	10.99	\$M
Real discount rate	r^A	0.9615	-
Real interest rate on fund	r^F	0.9825	-
Real interest rate on debt	r^D	1.0100	-
Time horizon	T	20	years
Debt sustainability constraint	ϵ	0.05	\$M
Normalization for power price index	k^P	350	\$/MWh
Normalization for hedge contract slope	k^H	4	\$/inch
Normalization for revenues & cash flows	k^R	250	\$M
Normalization for fund & debt	k^F	150	\$M

338

3.1 Control formulations

339

340

341

342

343

344

Within the hydro-financial simulation model, there are two important decisions that must be made each year: the hedging contract slope and the withdrawal from/deposit to the reserve fund. A control policy refers to a structured set of rules for making these two decisions each year. This study introduces two types of control: static (or open-loop) policies, which perform the same actions with each time step (Section 3.1.1), and dynamic (or closed-loop) policies, which adapt to changing conditions over time (Section 3.1.2).

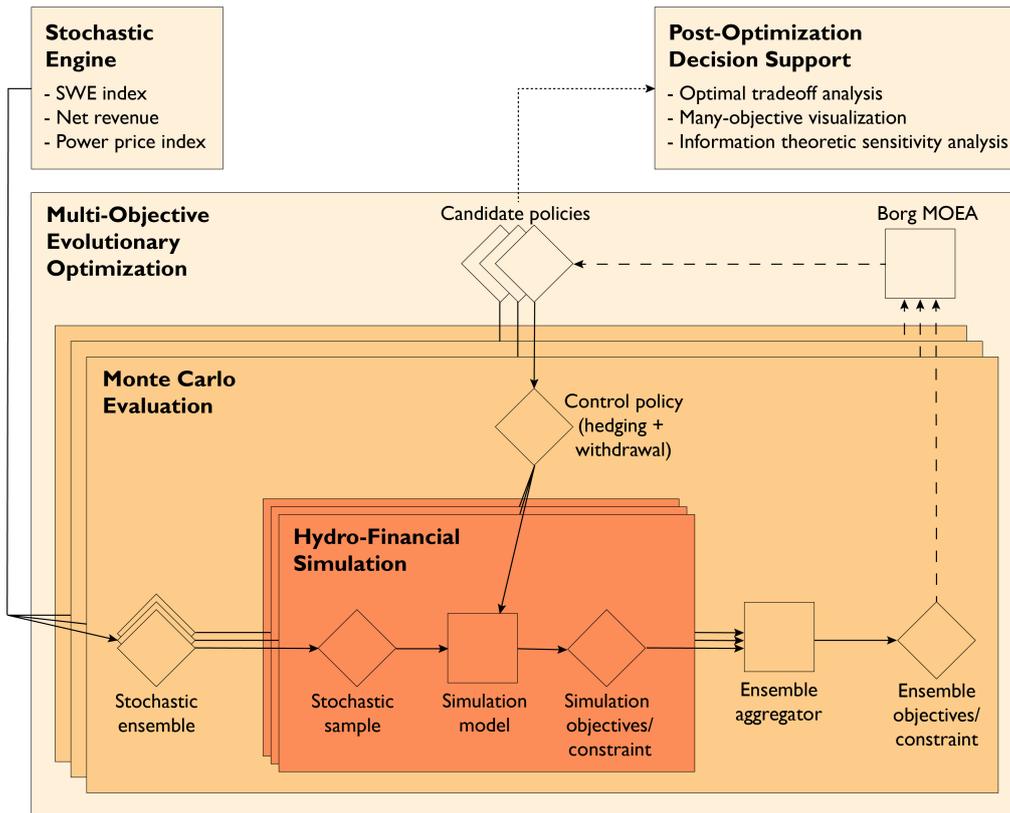


Figure 3. Schematic showing overall workflow for this study. Rectangles represent modules and diamonds represent inputs/outputs. Dashed arrows show the feedback process for the Borg MOEA, where objective and constraint values from prior control policy evaluations are used to generate new candidate policies for evaluation. The dotted arrow represents the final population output from the MOEA search, which is used as input to the post-optimization decision support.

345 Dynamic policies are considered state-aware because the decisions at each time step are
 346 conditioned on the current state of the model. Under both static and dynamic formu-
 347 lations, a policy is defined by a parameter vector which governs its operations. Multi-
 348 objective evolutionary optimization (Section 3.3) will be used to search for parameter
 349 vectors that perform well across four objectives related to the annualized cash flow, the
 350 risk of extreme debt levels, the probability of using hedging contracts, and the size of
 351 the reserve fund (Section 3.2).

3.1.1 Static policies

The static control formulation (adapted from Hamilton et al. (2020)) is given by:

$$\boldsymbol{\theta}_{stat} = [u^H, x_{max}^F] \quad (7)$$

where $\boldsymbol{\theta}_{stat}$ is the policy parameter vector and u^H and x_{max}^F are the two parameters to be optimized. u^H is the CFD slope, which is held fixed across all years in the simulation, while x_{max}^F is the maximum allowable reserve fund. Given x_{max}^F , the reserve fund operates according to the following simple rules: If the intermediate cash flow is negative ($x_t^{C2} < 0$), cash is withdrawn from the reserve fund to make up the deficit if possible. If $x_t^{C2} > 0$, the surplus is deposited into the fund, up until the fund has reached x_{max}^F . This policy is referred to as “static” because the CFD slope does not react to changing conditions (i.e., it is not state-aware). Although the withdrawal policy is quasi-state-aware via cash-balance constraints (money can neither be created nor destroyed), it is not truly dynamic in a meaningful sense (e.g., it cannot condition its reserve fund target on power price projections). Note that in Figure 2, the static formulation does not include the three input arrows into u_t^H , and only includes the two input arrows into u_t^W that relate to the cash balance constraints (x_t^{C2} and x_{t-1}^F).

3.1.2 Dynamic policies using Direct Policy Search (DPS)

The dynamic control formulation conditions the decision at each time step on the information available at that time. For a decision $u_t^{\mathcal{D}}$, with $\mathcal{D} \in \{W, H\}$ representing the withdrawal and hedging decisions, respectively:

$$u_t^{\mathcal{D}} = \mathcal{P}^{\mathcal{D}}(\mathcal{I}_{t'}^{\mathcal{D}} | \boldsymbol{\theta}_{dyn}^{\mathcal{D}}) \quad (8)$$

where $\mathcal{P}^{\mathcal{D}}$ is the mathematical form of the policy for decision \mathcal{D} (e.g., discrete policy table for SDP), $\boldsymbol{\theta}_{dyn}^{\mathcal{D}}$ is the vector of parameters to be optimized for the policy, and $\mathcal{I}_{t'}^{\mathcal{D}}$ is the information upon which the decision is conditioned. This information can be any subset of the model states, actions, and stochastic drivers. The subscript t' on each element represents either the current (t) or previous ($t - 1$) time step, based on the sequential nature of decisions (see Figure 2).

In DPS, \mathcal{P} is assumed to be a family of parametric functions (Rosenstein & Barto, 2001). This approximation drastically reduces the number of decision variables in the search relative to SDP (Bertsekas, 2019; Powell, 2019). Many parametric function fam-

382 ilies are available (e.g., piecewise linear, polynomial, artificial neural network), but ra-
 383 dial basis functions (RBFs) have been shown to be efficient universal approximators for
 384 DPS (Giuliani, Mason, Castelletti, Pianosi, & Soncini-Sessa, 2014). In this work, a sum
 385 of RBFs is paired with a constant shift parameter, along with an outer function that per-
 386 forms operations such as normalization and constraints. Equation 8 can be rewritten as:

$$387 \quad u_t^{\mathcal{D}} = \phi^{\mathcal{D}} \left(a^{\mathcal{D}} + \sum_{m=1}^M w_m^{\mathcal{D}} \varphi_m \left(\mathcal{I}_{t'}^{\mathcal{D}} \right) \right) \quad (9)$$

388 where $\phi^{\mathcal{D}}$ is the outer function, $a^{\mathcal{D}} \in [-1, 1]$ is a constant shift, and $w_m^{\mathcal{D}}$ is the weight
 389 given to the m th out of M total RBFs, φ_m . The weights must be chosen such that $\sum_{m=1}^M w_m^{\mathcal{D}} =$
 390 1, and $w_m^{\mathcal{D}} \geq 0$ for all m . The RBF is defined

$$391 \quad \varphi_m(\mathcal{I}_{t'}^{\mathcal{D}}) = \exp \left(- \sum_{l=1}^L \frac{\left([\mathcal{I}_{t'}^{\mathcal{D}}]_l - c_{l,m} \right)^2}{(b_{l,m})^2} \right) \quad (10)$$

392 where $[\mathcal{I}_{t'}^{\mathcal{D}}]_l$ is the l th out of L informational inputs, and $c_{l,m} \in [-1, 1]$ and $b_{l,m} \in (0, 1]$
 393 are the center and radius, respectively, of the m th RBF in the direction of the l th in-
 394 put. The M RBFs are shared by the two decisions in the control policy.

395 The information vector for each decision includes the combination of state variables
 396 and external drivers that might be useful for making the decision:

$$397 \quad \mathcal{I}_{t'}^W = [r^F \tilde{x}_{t-1}^F, \quad r^D \tilde{x}_{t-1}^D, \quad \tilde{\varepsilon}_t^P, \quad \tilde{x}_t^{C2}] \quad (11)$$

$$398 \quad \mathcal{I}_{t'}^H = [\tilde{x}_t^F, \quad \tilde{x}_t^D, \quad \tilde{\varepsilon}_t^P] \quad (12)$$

399 where all tildes represent values that have been normalized to lie between 0 and 1, us-
 400 ing the normalization constants in Table 1. Both decisions utilize information about the
 401 reserve fund balance and debt, but $u^{\mathcal{D}}$ uses last year's balance plus accumulated inter-
 402 est, while u^W uses the updated value from the present year (Figure 2). Both decisions
 403 also use the current power price index. Finally, the cash flow prior to withdrawal/deposit,
 404 x_t^{C2} , is used for u^W but not $u^{\mathcal{D}}$. Because the M RBFs are shared across the two deci-
 405 sions, $L = \max(L^W, L^H) = 4$.

406 The outer functions ϕ^W and ϕ^H (Equation 9) each consist of multiple nested func-
 407 tions performing specific operations. The more straightforward ϕ^H consists of a normal-
 408 ization function, ϕ^{HN} , and a constraint function, ϕ^{HC} . Let z_t be the argument to ϕ^H ,
 409 the action prescribed by the constant shift and sum of radial basis functions in Equa-
 410 tion 9 when H is substituted for \mathcal{D} . This equation can be decomposed as

$$411 \quad u_t^H = \phi^H(z_t) = \phi^{HC}(\phi^{HN}(z_t)) \quad (13)$$

412 First, ϕ^{HN} scales the hedging contract slope to the proper scale, $[0, k^H]$ (\$M/inch),
 413 where k^H is the hedging contract normalization constant in Table 1.

$$414 \quad z'_t = \phi^{HN}(z_t) = k^H \max(\min(z_t, 1), 0) \quad (14)$$

415 Next, ϕ^{HC} constrains the contract slope to be greater than or equal to a constant
 416 threshold, $k^H d^H$, where the threshold parameter $d^H \in [0, 1]$ is included in the policy
 417 parameter vector to be optimized, along with a^H , \mathbf{w}^H , \mathbf{c} , and \mathbf{b} .

$$418 \quad u_t^H = \phi^{HC}(z'_t) = \begin{cases} z'_t, & \text{if } z'_t \geq k^H d^H \\ 0, & \text{otherwise} \end{cases} \quad (15)$$

419 The outer function for the withdrawal decision, ϕ^W , consists of four nested oper-
 420 ations. Let z_t now be the sum of the constant shift and RBFs in Equation 9 when W
 421 is substituted for \mathcal{D} . Then:

$$422 \quad u_t^W = \phi^W(z_t) = \phi^{WCO}(\phi^{WCI}(\phi^{WW}(\phi^{WN}(z_t)))) \quad (16)$$

423 where ϕ^{WCO} , ϕ^{WCI} , ϕ^{WW} , and ϕ^{WN} are the outer constraint, inner constraint, with-
 424 drawal transformation, and normalization functions. First, when designing the withdrawal
 425 policy, it was discovered that the EMODPS search produces better results when z_t is
 426 defined as the prescribed post-withdrawal cash flow rather than the withdrawal itself. For
 427 this reason, the normalization function, ϕ^{WN} , transforms z_t to the scale of $[-k^R, k^R]$ (\$M),
 428 where k^R is the normalization constant for all revenues and cash flows in Table 1.

$$429 \quad z'_t = \phi^{WN}(z_t) = k^R \max(\min(2z_t - 1, 1), -1) \quad (17)$$

430 The withdrawal transformation function, ϕ^{WW} , transforms z'_t from a cash flow into
 431 a withdrawal/deposit using the relationship between incoming and outgoing cash flow:

$$432 \quad z''_t = \phi^{WW}(z'_t) = z'_t - x_t^{C2} \quad (18)$$

433 The inner constraint function, ϕ^{WCI} , ensures that the withdrawal/deposit is con-
 434 sistent with cash-balance equations:

$$435 \quad z'''_t = \phi^{WCI}(z''_t) = \begin{cases} \min(z''_t, r^F x_{t-1}^F), & \text{if } z''_t \geq 0 \\ \max(z''_t, -\max(x_t^{C2}, 0)), & \text{otherwise} \end{cases} \quad (19)$$

436 The first condition ensures that a withdrawal ($z_t'' > 0$) cannot be larger than the bal-
 437 ance in the reserve fund. The second case dictates that a deposit ($z_t'' < 0$) is only al-
 438 lowed when the available cash flow x_t^{C2} is positive, and that the deposit cannot be larger
 439 in magnitude than this cash flow.

440 Lastly, the outer constraint, ϕ^{WCO} , ensures that the reserve fund balance (after
 441 withdrawal/deposit) cannot be larger than a constant threshold, $k^F d^W$, where k^F (\$M)
 442 is the normalization constant used for the reserve fund and debt in Table 1, and $d^W \in$
 443 $[0, 1]$ is another decision variable to be optimized.

$$444 \quad u_t^W = \phi^{WCO}(z_t''') = \begin{cases} r^F x_{t-1}^F - k^F d^W, & \text{if } (r^F x_{t-1}^F - z_t''') > k^F d^W \\ z_t''', & \text{otherwise} \end{cases} \quad (20)$$

445 This threshold sets the maximum allowable reserve fund size, equivalent to x_{max}^F in the
 446 static formulation.

447 Equations 8-20 constitute the full dynamic control policy. The parameter vector
 448 to be optimized for each decision $\mathcal{D} \in \{W, H\}$ is

$$449 \quad \boldsymbol{\theta}_{dyn}^{\mathcal{D}} = [a^{\mathcal{D}}, \quad d^{\mathcal{D}}, \quad \mathbf{w}^{\mathcal{D}}, \quad \mathbf{c}, \quad \mathbf{b}] \quad (21)$$

450 where $\mathbf{w}^{\mathcal{D}} = [w_0^{\mathcal{D}}, \dots, w_M^{\mathcal{D}}]$, $\mathbf{c} = [c_{0,0}, \dots, c_{L,M}]$, and $\mathbf{b} = [b_{0,0}, \dots, b_{L,M}]$. The total pa-
 451 rameter vector to be optimized, $\boldsymbol{\theta}_{dyn}$, is the set of unique parameters,

$$452 \quad \boldsymbol{\theta}_{dyn} = [a^W, \quad a^H, \quad d^W, \quad d^H, \quad \mathbf{w}^W, \quad \mathbf{w}^H, \quad \mathbf{c}, \quad \mathbf{b}] \quad (22)$$

453 3.2 Objective formulations

454 This study uses “noisy” objective formulations to account for the uncertainty of
 455 outcomes under the stochastic drivers. Each candidate policy is evaluated using a Monte
 456 Carlo ensemble of N realizations, each representing one possible trajectory of the hydro-
 457 financial system under a T -year sample of the stochastic drivers. To convert an ensem-
 458 ble of time series into a scalar performance metric requires both a time aggregation step
 459 (e.g., taking the maximum debt over a T -year realization) and a noise filtering step (e.g.,
 460 taking the 95th percentile over N realizations in the ensemble). Four objectives are con-
 461 sidered in this study, each defined as the maximization or minimization of a particular
 462 performance metric.

463 The first objective is to maximize the expected annualized cash flow, J^{cash} , a mea-
 464 sure of “average” cash flows. A high value represents a low-cost risk management pol-

465 icy. Although public utilities are not strictly profit-maximizing firms, they nonetheless
 466 aim to maintain sufficient cash flows to keep customer rates low and/or invest in new
 467 infrastructure, and J^{cash} is used as a proxy for this type of financial health.

$$468 \quad J^{cash} \left(x_{t \in (1, \dots, T)}^{C^4}, x_T^F, x_T^D \right) = E_{\epsilon} \left[ANN_t \left(x_{t \in (1, \dots, T)}^{C^4}, x_T^F, x_T^D \right) \right] \quad (23)$$

469 where $x_t^{C^4}$ is the final cash flow for year t ; x_T^F and x_T^D are the reserve fund balance and
 470 debt at the end of the simulation; E_{ϵ} is the expectation over the stochastic drivers (ap-
 471 proximated by the mean of N Monte Carlo samples); and ANN_t is the annualization
 472 operator:

$$473 \quad ANN_t \left(x_{t \in (1, \dots, T)}^{C^4}, x_T^F, x_T^D \right) = \frac{1}{\sum_{t=1}^T (r^A)^t} \left(\sum_{t=1}^T ((r^A)^t x_t^{C^4}) + (r^A)^{T+1} (r^F x_T^F - r^D x_T^D) \right) \quad (24)$$

474 where where r^A is the real discount rate and r^F and r^D are the real interest rates on re-
 475 serves and debt (Table 1). ANN_t sums the net present value (NPV) of all discounted
 476 cash flows over T years, plus the NPV of the reserve fund and debt in year T , and di-
 477 vides this sum by a normalization factor. The normalized value represents the constant
 478 cash flow, or annuity, that is equivalent in terms of NPV to the variable cash flow. On
 479 the whole, annualization allows for a fair comparison, accounting for the time value of
 480 money, between cash flow time series resulting from different management strategies.

481 The second objective is to minimize J^{debt} , the 95th percentile of maximum debt.
 482 This is a measure of the short-term debt load that would be needed to meet fixed costs
 483 in an extremely bad year (or sequence of years). This performance metric is used as a
 484 proxy for “risk”, and a decision-maker would want to minimize this quantity in order
 485 to avoid compromising the utility’s credit rating, increasing future borrowing costs, and/or
 486 risking bankruptcy.

$$487 \quad J^{debt} \left(x_{t \in (1, \dots, T)}^D \right) = Q95_{\epsilon} \left[\max_{t \in (1, \dots, T)} [x_t^D] \right] \quad (25)$$

488 where the *max* operator takes the maximum debt over a T -year realization, and the $Q95$
 489 operator takes the 95th percentile over the Monte Carlo ensemble.

490 These first two objectives, adopted from Hamilton et al. (2020), are representative
 491 of the risk/return tradeoff analysis that is common in financial applications (Hull, 2009;
 492 Markowitz, 1952). However, financial researchers have found that higher-dimensional prob-
 493 lem framings can more accurately represent managers’ behavior in the empirical data
 494 (Spronk, Steuer, & Zopounidis, 2005; Zopounidis et al., 2015). For example, in addition

495 to maximizing return and minimizing risk, an investment portfolio manager might want
 496 to minimize the number of unique securities held because this limits the associated pa-
 497 perwork, transactions fees, etc. Similarly, in workshops designed to help water utilities
 498 integrate MOEAs into their water portfolio planning processes, Smith et al. (2019) have
 499 found that managers often weigh the decision levers (e.g., whether a new reservoir must
 500 be built) alongside more traditional measures of portfolio performance (e.g., supply re-
 501 liability) when deciding which portfolio to choose. This represents an expansion of the
 502 objective space in practice, and reflects decision-makers' expert knowledge of the trade-
 503 offs associated with various management tools. Bringing together these lines of research,
 504 a utility manager would be expected to balance tradeoffs associated with different finan-
 505 cial risk management tools in addition to performance metrics like risk and return (Bank
 506 & Wiesner, 2010; Hughes et al., 2014). Two additional objectives are now introduced
 507 in order to explore the impact of such tradeoffs.

508 The third objective is to minimize J^{hedge} , the expected hedging frequency.

$$509 \quad J^{hedge} \left(u_{t \in (0, \dots, T-1)}^H \right) = E_{\epsilon} \left[\max_{t \in (0, \dots, T-1)} \left[\mathbf{1}_{u_t^H > 0} \right] \right] \quad (26)$$

510 where the indicator function $\mathbf{1}_{u_t^H > 0}$ returns a 1 if the hedging contract slope is non-zero,
 511 and a 0 otherwise. This metric represents the likelihood that the utility will enter into
 512 at least one hedging contract over the course of 20 years. Note that each hedging con-
 513 tract does have an annual cost, a “loading” applied by the contract seller that makes the
 514 expected payout of h (SI Figure S3) negative (Hamilton et al., 2020). However, this cost
 515 is already accounted for by J^{cash} , and does not need to be double-counted. J^{hedge} , rather,
 516 relates to the significant extra costs (in time, personnel, and/or money) of having to set
 517 up the first hedging contract within a realization, assuming that this start-up cost will
 518 be significantly diminished in subsequent contract purchases. Moreover, this objective
 519 can be taken to represent the general discomfort that a utility manager may have with
 520 financial hedging contracts due to their novelty and perceived complexity or opacity (Bank
 521 & Wiesner, 2010).

522 The last objective is to minimize J^{fund} , the expected maximum reserve fund bal-
 523 ance.

$$524 \quad J^{fund} \left(x_{t \in (1, \dots, T)}^F \right) = E_{\epsilon} \left[\max_{t \in (1, \dots, T)} \left[x_t^F \right] \right] \quad (27)$$

525 This metric represents the expected value of the largest reserve fund used in a T -year
 526 realization, which a utility manager may want to minimize in order to avoid attracting
 527 regulatory scrutiny over holding large liquid reserves (Hughes et al., 2014).

528 Finally, a “debt sustainability” constraint ensures that feasible policies do not al-
 529 low debt to grow unchecked over time (on average), which would likely lead to a credit
 530 downgrade in practice:

$$531 \quad E_{\epsilon} [x_T^D - x_{T-1}^D] < \epsilon \quad (28)$$

532 where ϵ is a small constant (Table 1). This “noisy” constraint is calculated from the en-
 533 tire Monte Carlo ensemble; there is no constraint on debt use in individual extreme re-
 534 alizations.

535 **3.3 Multi-objective evolutionary optimization of control policies**

536 As described in Sections 1 and 3.1.2, DPS has a number of advantages relative to
 537 traditional methods such as SDP, especially when combined with non-linear approximat-
 538 ing networks such as RBFs. However, RBF parameterization can result in a highly non-
 539 linear and non-convex search space that is difficult to traverse with gradient-based meth-
 540 ods, especially when combined with noisy multi-objective formulations (Giuliani & Castel-
 541 letti, 2016; Giuliani, Mason, et al., 2014; Giuliani et al., 2018). These problems are bet-
 542 ter handled by MOEAs, which use evolution-inspired strategies (e.g., selection, mating,
 543 mutation) to iteratively improve a population of solutions competing on multiple objec-
 544 tives (Coello Coello, Lamont, & Van Veldhuizen, 2007). Population-based methods can
 545 approximate the entire Pareto set in a single run, rather than rerunning many single-
 546 objective optimizations, making them quite efficient on many-objective problems. Ad-
 547 ditionally, these heuristic approaches require no information on the topology of a prob-
 548 lem and are well-adapted to the types of nonlinear, non-convex, high-dimensional, and
 549 stochastic problems that are common in both water resources (Maier et al., 2014; Nick-
 550 low et al., 2010; Reed, Hadka, Herman, Kasprzyk, & Kollat, 2013) and finance (Ponsich
 551 et al., 2013; Tapia & Coello Coello, 2007).

552 This study employs the Borg Multiobjective Evolutionary Algorithm (MOEA) (Hadka
 553 & Reed, 2013), which has been particularly successful across a range of difficult prob-
 554 lems in water resources (Gupta et al., 2020; Hadka & Reed, 2012; Reed et al., 2013; Zatarain Salazar,
 555 Reed, Herman, Giuliani, & Castelletti, 2016) and engineering design (Singh et al., 2020;

556 Woodruff, Reed, & Simpson, 2013). The Borg MOEA includes novel components such
557 as adaptive search operator selection, adaptive population sizing, stagnation detection
558 via epsilon-progress, and epsilon-dominance archiving. Its self-adaptive nature makes the
559 Borg MOEA highly controllable (Hadka & Reed, 2013; Reed et al., 2013), and the master-
560 worker parallel variant used in this study is scalable on high-performance computing in-
561 frastructure (Giuliani et al., 2018; Zatarain Salazar et al., 2017).

562 **3.4 Information theoretic sensitivity analysis**

563 A sensitivity analysis (SA) is an evaluation of the effects of a model’s input fac-
564 tors on its output factors, and a wide range of methods are available to suit different pur-
565 poses. According to the taxonomy of SA introduced by Pianosi et al. (2016), the method
566 that follows would be considered a quantitative, global, “all-at-a-time” SA, based on sim-
567 ulation model output. This SA is used to explore how different policies adapt their ac-
568 tions to changing conditions; more specifically, it will probe the sensitivity of the pre-
569 scribed hedging and withdrawal decisions (Equation 8) to changing informational inputs
570 (Equations 11-12). This type of analysis can help to “open the black box” of control poli-
571 cies, helping decision-makers better understand how different policies respond to chang-
572 ing information (Quinn et al., 2019).

573 However, commonly-used variance-based methods, which decompose the variance
574 of an output variable into contributions from covariance with different input variables,
575 are inappropriate in the proposed context. First, the policies described by Equations 9-
576 20 are highly non-linear and discontinuous, so that variance and covariance are inappro-
577 priate measures of variability and relationship. Secondly, most variance decomposition
578 methods assume independence between the input variables, and can lead to misleading
579 results when this independence is violated (Borgonovo, 2007; Borgonovo, Castaings, &
580 Tarantola, 2011). This is especially problematic in the current context because most Pareto-
581 optimal solutions will impose the following relationship between the reserve fund and debt:
582 if one is large, the other is usually zero. For these reasons, moment-independent global
583 SA methods, such as entropy-based SA (Auder & Iooss, 2009; Krzykacz-Hausmann, 2001),
584 are preferred. Hejazi, Cai, and Ruddell (2008) use ITSA to study the impact of hydro-
585 logic information on historical release decisions made by reservoir operators under dif-
586 ferent conditions. A similar approach is adopted here to study how different policies along
587 the Pareto front use model state information to make decisions.

588 Shannon entropy (Shannon, 1948) quantifies how much information is needed, on
 589 average, to describe a random variable. Consider $u^{\mathcal{D}}$, $\mathcal{D} \in \{W, H\}$, the two policy-prescribed
 590 actions. $u^{\mathcal{D}}$ is a function of the information vector, $\mathcal{I}^{\mathcal{D}}$, which varies stochastically through
 591 time and across Monte Carlo realizations. As such, both the information vector and the
 592 prescribed action can be considered random variables, $\mathbf{I}^{\mathcal{D}}$ and $U^{\mathcal{D}}$. The entropy of the
 593 action is:

$$594 \quad H(U^{\mathcal{D}}) = - \sum_{u^{\mathcal{D}} \in v^{\mathcal{D}}} p(u^{\mathcal{D}}) \log_2 p(u^{\mathcal{D}}) \quad (29)$$

595 where $p(u^{\mathcal{D}})$ is the probability mass function (PMF) after discretizing the outcome to
 596 a discrete domain, $v^{\mathcal{D}}$. The entropy (in bits when written with a base-2 logarithm) can
 597 be thought of as a moment-free measure of uncertainty, or dispersion, in the probabil-
 598 ity distribution of a random variable. A variable whose outcome is known determinis-
 599 tically has zero entropy, while a uniformly distributed variable is the most uncertain and
 600 has the largest possible entropy. Although a continuous variant of entropy based on Kullback-
 601 Leibler divergence can also be used for SA (Auder & Iooss, 2009; Liu, Chen, & Sudjianto,
 602 2006; Pappenberger, Beven, Ratto, & Matgen, 2008), the discrete version is more straight-
 603 forward when the random variable's distribution is unknown.

604 The mutual information between two random variables measures the average re-
 605 duction in the entropy of one variable when the other variable's outcome is known:

$$606 \quad MI(\mathbf{I}_i^{\mathcal{D}}, U^{\mathcal{D}}) = H(U^{\mathcal{D}}) - H(U^{\mathcal{D}} | \mathbf{I}_i^{\mathcal{D}}) \quad (30)$$

$$607 \quad = - \sum_{\mathcal{I}_i^{\mathcal{D}} \in \iota_i^{\mathcal{D}}} \sum_{u^{\mathcal{D}} \in v^{\mathcal{D}}} p(\mathcal{I}_i^{\mathcal{D}}, u^{\mathcal{D}}) \log_2 \frac{p(\mathcal{I}_i^{\mathcal{D}}, u^{\mathcal{D}})}{p(\mathcal{I}_i^{\mathcal{D}})p(u^{\mathcal{D}})} \quad (31)$$

608 where $\mathbf{I}_i^{\mathcal{D}}$ is the random variable for the i th informational input (e.g., reserve fund bal-
 609 ance or power price index), $H(U^{\mathcal{D}} | \mathbf{I}_i^{\mathcal{D}})$ is the entropy of the action conditional on the in-
 610 put, $p(\mathcal{I}_i^{\mathcal{D}})$ is the PMF for the input on the discrete domain $\iota_i^{\mathcal{D}}$, and $p(\mathcal{I}_i^{\mathcal{D}}, u^{\mathcal{D}})$ is the
 611 joint PMF on the discrete domain $\iota_i^{\mathcal{D}} \times v^{\mathcal{D}}$. This mutual information is a measure how
 612 much information the outcome of one random variable contains about the outcome of
 613 the other: how much does knowledge of a particular informational input reduce the un-
 614 certainty in the prescribed action?

615 Finally, the ITSA index is defined by dividing the mutual information by the en-
 616 tropy of the prescribed action:

$$617 \quad \eta_i^{\mathcal{D}} = \frac{MI(\mathbf{I}_i^{\mathcal{D}}, U^{\mathcal{D}})}{H(U^{\mathcal{D}})} \quad (32)$$

618 where $\eta_i^{\mathcal{D}}$ is the sensitivity index for the i th input for decision \mathcal{D} . This index varies be-
 619 tween 0 and 1; $\eta_i^{\mathcal{D}} = 0$ implies that $\mathbf{I}_i^{\mathcal{D}}$ and $U^{\mathcal{D}}$ are independent random variables, while
 620 $\eta_i^{\mathcal{D}} = 1$ implies perfect dependence (knowledge of $\mathcal{I}_i^{\mathcal{D}}$ gives us perfect knowledge of $u^{\mathcal{D}}$).

621 4 Computational experiments

622 4.1 Problem formulations

623 This study considers both the static and dynamic control formulations, each of which
 624 has its own parameter vector to be optimized. The static parameter vector ($\boldsymbol{\theta}_{stat}$, Equa-
 625 tion 7) has two elements to be optimized. The dynamic parameter vector, ($\boldsymbol{\theta}_{dyn}$, Equa-
 626 tion 22) has $4 + 2M + 2ML$ elements, where $L = 4$ is the number of informational in-
 627 puts, and M is the number of RBFs in the policy. With $M = 2$ RBFs (see next sec-
 628 tion), $\boldsymbol{\theta}_{dyn}$ contains 24 elements to be optimized.

629 For each control formulation, both two-objective and four-objective problems are
 630 considered. The two-objective problem can be written:

$$631 \quad \boldsymbol{\theta}^* = \arg \min_{\boldsymbol{\theta}} [-J^{cash}(\boldsymbol{\theta}), \quad J^{debt}(\boldsymbol{\theta})] \quad (33)$$

632 while the four-objective problem can be written:

$$633 \quad \boldsymbol{\theta}^* = \arg \min_{\boldsymbol{\theta}} [-J^{cash}(\boldsymbol{\theta}), \quad J^{debt}(\boldsymbol{\theta}), \quad J^{hedge}(\boldsymbol{\theta}), \quad J^{fund}(\boldsymbol{\theta})] \quad (34)$$

634 For both problems, the feasible solution space is restricted to solutions satisfying the sus-
 635 sustainable debt constraint (Equation 28). The two-objective problem is the same as that
 636 used by Hamilton et al. (2020), allowing for a direct comparison, while the four-objective
 637 problem provides more nuanced insight into risk management tradeoffs.

638 4.2 MOEA experiments

639 An ensemble of $N = 50,000$ realizations is run for each function evaluation, bal-
 640 ancing computational demand against the need to minimize sampling error in the noisy
 641 objective/constraint evaluations (see discussions in Kasprzyk et al. (2012); Quinn, Reed,
 642 Giuliani, and Castelletti (2017); Zatarain Salazar et al. (2017)). In order to select the
 643 appropriate number of RBFs, the dynamic 4-objective formulation is repeated with 1,
 644 2, 3, 4, 8, and 12 RBFs. Due to the inherent stochasticity of evolutionary algorithms,
 645 each optimization is repeated with 10 different random seeds. Each seed is run for 150,000
 646 function evaluations (candidate policy trials). Final populations are assessed in terms

647 of hypervolume, additive epsilon indicator, and generational distance (SI Figure S4), three
 648 common metrics for assessing convergence, consistency, and diversity of multi-objective
 649 solution sets (Coello Coello et al., 2007; Hadka & Reed, 2012; Reed et al., 2013). Results
 650 are found to be relatively insensitive to the number of RBFs used in the dynamic con-
 651 trol policies, but $M = 2$ RBFs is chosen due to the robust performance across seeds.
 652 Next, 20 additional seeds are run for the dynamic 4-objective formulation with $M =$
 653 2, and 30 seeds each are also run for the dynamic 2-objective, static 2-objective, and static
 654 4-objective formulations. The best known Pareto approximate set for each formulation
 655 is the set of non-dominated solutions from across the 30 seeds. After using the same 50,000-
 656 member ensemble of 20-year simulations for all formulations/seeds in the initial optimiza-
 657 tion, each solution in the final Pareto approximate set for each formulation is rerun on
 658 a separate 50,000-member ensemble, for which results are reported. Important param-
 659 eter values for the optimization can be found in SI Table S1; all other Borg MOEA pa-
 660 rameters besides those listed are set to the default values (Hadka & Reed, 2013; Reed
 661 et al., 2013).

662 4.3 Information theoretic sensitivity analysis parameters

663 ITSA indices for each specific operating policy are calculated using a 50,000-member
 664 ensemble of 20-year simulations, yielding 1,000,000 realizations of \mathcal{I}_i^D and u^D . Each com-
 665 ponent is discretized into 50 bins in order to calculate the marginal and joint probabil-
 666 ity mass functions (Equations 29, 31). This process is repeated for each control policy
 667 in the Pareto set, yielding separate ITSA indices for each.

668 5 Results and discussion

669 5.1 Static vs. dynamic financial risk management

670 Figure 4 shows the resulting Pareto approximate sets from the 2-objective optimiza-
 671 tion problem (Equation 33), under both static and dynamic control formulations. Each
 672 point represents a different financial risk management policy. The ideal performance, de-
 673 noted by a black star, would be achieved with a cash flow metric (J^{cash}) of \$10.99M (the
 674 average net revenue in the absence of any financial risk management) and a debt met-
 675 ric (J^{debt}) of zero. However, this is not possible due to the strong tradeoff between “risk”
 676 and “return” that is standard in financial risk applications: in order to achieve higher

677 expected cash flows, the utility must forego costly risk management actions and there-
 678 fore risk more extreme debt burdens in less favorable realizations. As discussed in Sec-
 679 tion 3.2, large short-term debt in our model can be viewed as a proxy for larger finan-
 680 cial disruptions such as credit rating downgrades or bankruptcy in practice. Decision-
 681 makers will have to balance this tradeoff when selecting a particular policy for the util-
 682 ity to use, based on risk aversion, access to credit, and other organizational factors.

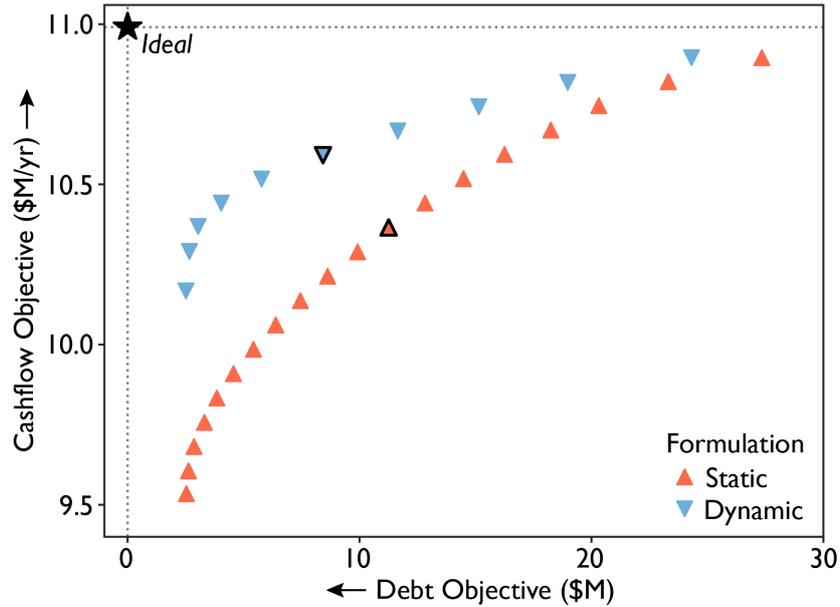


Figure 4. Comparison of 2-objective Pareto approximate sets under static and dynamic control formulations. The best compromise policy from each formulation is outlined in black and described in Table 2.

683 However, decision-makers can drastically reduce the risk management tradeoff by
 684 using adaptive operating rules that respond to changing conditions. The Pareto approx-
 685 imate set from the dynamic EMODPS control formulation is found to dominate the Pareto
 686 approximate set from the static formulation, suggesting that one can improve on both
 687 the cash flow and debt objectives simultaneously. For example, consider the two exam-
 688 ple policies outlined in black in Figure 4 and listed in Rows 1-2 in Table 2. These are
 689 chosen as the “best compromise” policies near the centers of their respective Pareto ap-
 690 proximate sets (as selected using the TOPSIS method with equal weights on each ob-
 691 jective (Behzadian, Khanmohammadi Otaghsara, Yazdani, & Ignatius, 2012; Roszkowska,
 692 2011)). The dynamic policy is found to reduce J^{debt} by \$2.83M, or 25.1%, relative to the

693 static policy. At the same time, it increases J^{cash} by \$0.23M, representing a 36.1% re-
 694 duction in risk management cost. This dual improvement highlights the value of dynamic
 695 financial risk management: the utility can improve on both objectives simultaneously
 696 without requiring any investment in its infrastructure or changes to its physical oper-
 697 ations. All that is required is to switch to a more dynamic financial risk management
 698 policy.

Table 2. Performance of six example policies referenced throughout the results sections. Rows 1 and 2 represent the best compromise policies from the static and dynamic control formulations, respectively, under the 2-objective optimization problem (Section 5.1). Row 3 represents the best compromise policy from the 4-objective optimization problem and the dynamic control formulation, after brushing with *a posteriori* constraints (Section 5.2). Rows 4-6 represent policies that are highly sensitive to information about the reserve fund balance, debt, and power price index, respectively (Section 5.3).

Row	Figure	J^{cash} (\$M/yr)	J^{debt} (\$M)	J^{hedge} (unitless)	J^{fund} (\$M)	Fund Sensitivity	Debt Sensitivity	Power Sensitivity
1	4 red	10.37	11.25	1.00	16.11	–	–	–
2	4 blue	10.59	8.42	1.00	19.31	0.74	0.11	0.12
3	8	10.75	15.90	0.77	12.01	0.36	0.72	0.01
4	9a	10.20	3.22	1.00	24.55	0.93	0.12	0.00
5	9b	10.71	15.72	0.40	16.83	0.44	0.96	0.01
6	9c	9.84	8.96	1.00	1.53	0.02	0.03	0.72

699 The dynamic formulation allows the utility to take different sequences of actions
 700 under different stochastic realizations, using parameterized control rules that allow for
 701 the actions taken at any particular time to be better tailored to the current state of the
 702 system. To elucidate the differences between static and dynamic financial risk manage-
 703 ment, the two best compromise policies are simulated under two different 20-year real-
 704 izations from the synthetic record: an unusually wet period and an unusually dry pe-
 705 riod (Figure 5). Differences in SWE (5a) lead to drastic differences in hydropower gen-
 706 eration (5b) and net revenues (5d) under the two realizations, and the dry scenario ex-
 707periences lengthy periods of drought-related cash flow deficits. The two scenarios also

708 yield very different responses in terms of the hedging policy (5e & 5i), reserve fund bal-
 709 ance (5f & 5j), debt (5g & 5k), and final cash flow (5h & 5l). In the wet scenario, the
 710 reserve funds fill up quickly and stay nearly full. Neither policy requires any significant
 711 debt, and final cash flows are generally positive and rather large. In the dry scenario,
 712 the reserve funds fluctuate up and down, including two periods in which they reach zero.
 713 During these periods, significant debt is required to overcome further cash flow deficits.
 714 The final cash flows are close to zero throughout the dry simulation, as both policies strug-
 715 gle to fill their reserve funds.

716 With respect to the hedging contract, the static policy uses the same contract each
 717 year in both the wet and dry scenarios, with a payout slope of \$0.32M/inch. The dynamic
 718 policy, on the other hand, adjusts its contract slope from year to year. In the wet sce-
 719 nario, it opts not to hedge at all after year 0, while in the dry scenario, it fluctuates be-
 720 tween \$0 and \$0.85M/inch. Comparing the hedging slope dynamics to the other model
 721 state variables suggests that this policy opts to hedge only when the reserve fund bal-
 722 ance is low and/or when debt is non-zero. This strategy allows the dynamic policy to
 723 achieve higher cash flows than the static policy in wet scenarios (Sub-Figure 5h), by fore-
 724 going the cost of hedging contracts when the utility already has sufficient protection from
 725 a large reserve fund. On the other hand, when the reserve is empty and/or there is out-
 726 standing debt (presumably after a very dry year or sequence of years), the utility pur-
 727 chases large hedging contracts in order to increase its financial risk coverage and thus
 728 reduce the risk of extreme debt levels (Sub-Figure 5k). This adaptivity allows the dy-
 729 namic policy to improve on both the cash flow objective and the debt objective simul-
 730 taneously, compared to the static policy. As will be seen in Section 5.3, there are a mul-
 731 tiplicity of ways that utilities can adapt to changing conditions to meet their goals.

732 5.2 Many-objective decision-making

733 As discussed in Section 3.2, a decision-maker choosing a financial risk management
 734 policy may actually consider other factors beyond risk (J^{debt}) and return (J^{cash}). For
 735 example, the utility might also worry about the size of the reserve fund needed to en-
 736 act a particular policy (J^{fund}), or the likelihood of needing to develop and integrate a
 737 complicated hedging program (J^{hedge}). Such decision-makers are likely to find that none
 738 of the solutions found under the 2-objective problem (Figure 4) can meet their needs.
 739 The 2-objective problem cannot adequately represent important management tradeoffs

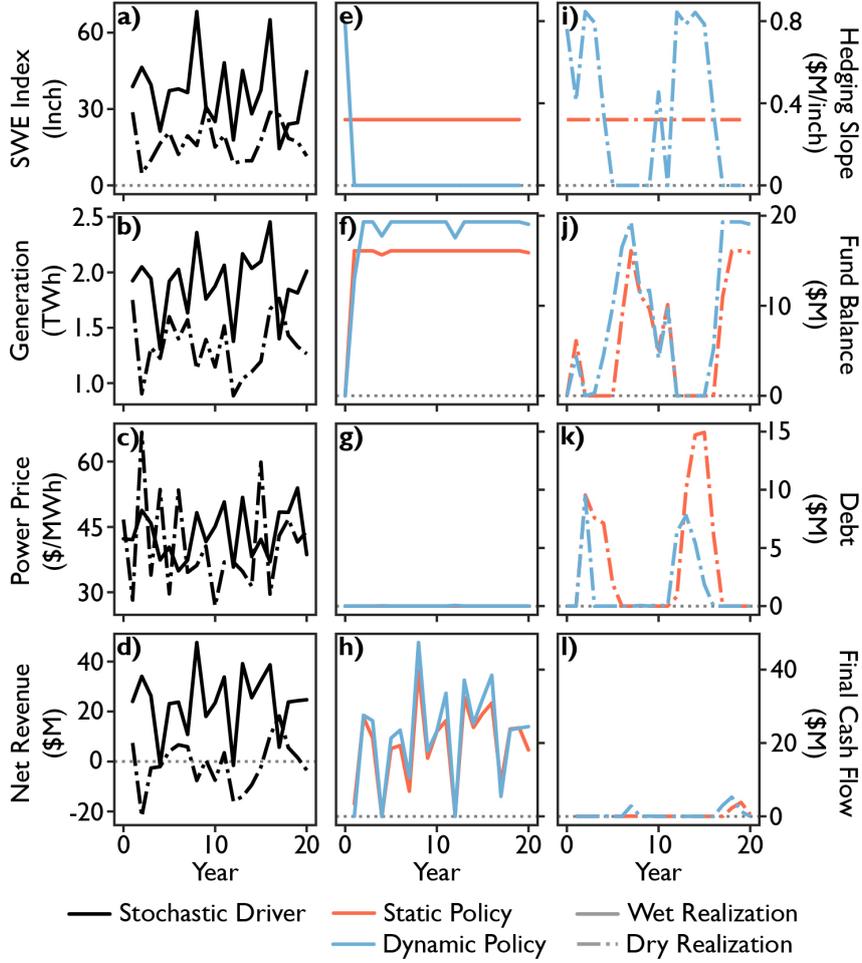


Figure 5. Trajectories for hydro-financial simulation model, over both wet and dry 20-year realizations, for the example static and dynamic policies shown in Figure 4 and Rows 1-2 of Table 2. Sub-Figures show (a) SWE index; (b) hydropower generation; (c) wholesale power price; (d) net hydropower revenue; (e & i) hedging slope action; (f & j) fund balance; (g & k) debt; and (h & l) final annual cash flow. Middle column (e-h) shares its y-axis with the right-hand column (i-l).

740 because it does not account for decision-maker preferences with respect to the use of dif-
 741 ferent risk management tools. For this reason, J^{hedge} and J^{fund} can be explicitly included
 742 in the optimization using the 4-objective problem (Equation 34).

743 Both the static and dynamic formulations produce much larger Pareto approximate
 744 sets in this higher-dimensional problem (Figure 6), representing the more complex set
 745 of tradeoffs across the four objectives. The dynamic Pareto approximate set is found to
 746 generally outperform the static Pareto approximate set, especially in terms of the over-

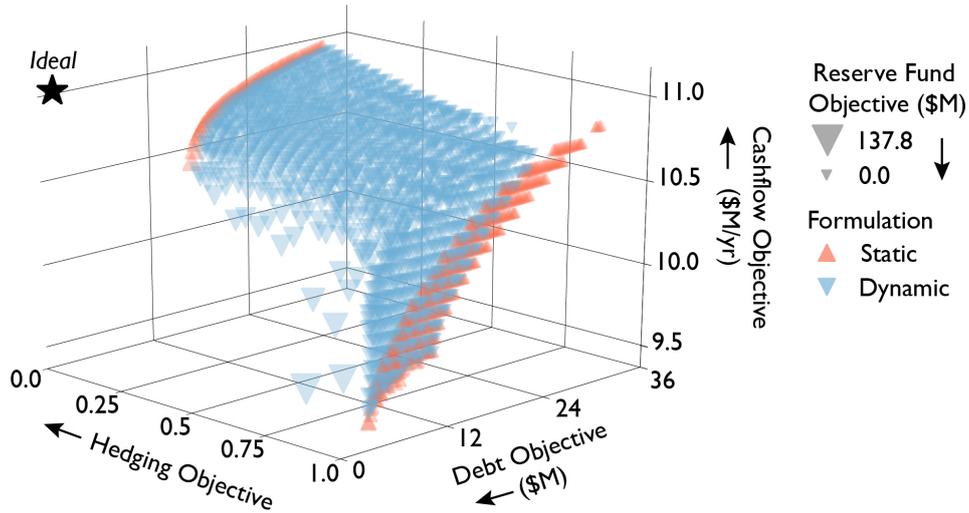


Figure 6. Comparison of 4-objective Pareto approximate sets under static and dynamic control formulations.

747 all diversity of solutions. For the static formulation, where the hedging contract slope
 748 is fixed, J^{hedge} must be equal to 1 or 0. The dynamic formulation, on the other hand,
 749 is able to find policies with J^{hedge} spanning the entire range from 0 to 1. Note that J^{hedge}
 750 is defined as the fraction of 20-year realizations that contain any hedging, not the frac-
 751 tion of years which hedge (see Equation 26). Thus, intermediate values between 0 and
 752 1 represent solutions that are unlikely to hedge in any given year, but maintain the op-
 753 tion to do so under particularly problematic circumstances. This valuable optionality
 754 is only possible with a dynamic control strategy. Additionally, the dynamic solution set
 755 occupies a much larger region within the ridge where $J^{hedge} = 1$. These policies out-
 756 perform the nearest static policies with respect to J^{cash} and J^{debt} , but may require the
 757 use of larger reserve funds. Because the dynamic control method produces a much more
 758 complete and continuous Pareto approximate set, it allows decision-makers to find con-
 759 trol policies that more precisely match their preferences.

760 A major benefit of solving the larger-dimensional problem is that the solution set
 761 will already contain all of the tradeoffs for all possible lower-dimensional problems (di
 762 Pierro, Khu, & Savić, 2007). In the present context, the 4-objective Pareto front will in-
 763 clude within it the Pareto fronts for the four 3-objective problems, six 2-objective prob-
 764 lems, and four 1-objective problems that are embedded within the 4-objective problem
 765 (Figure 7). In Sub-Figure 7a, the blue triangles show the subset of the 4-objective Pareto

766 approximate set that is non-dominated with respect to the original two objectives, J^{cash}
 767 and J^{debt} . When compared to the original 2-objective solutions (Figure 4), the 4-objective
 768 policies are very similar with respect to the first two objectives. However, they can achieve
 769 improvements with respect to the two new objectives (see SI Figure S5). In other words,
 770 it is possible to improve J^{fund} and/or J^{hedge} with no penalty in J^{cash} or J^{debt} , but they
 771 must be included in the optimization explicitly to realize this benefit.

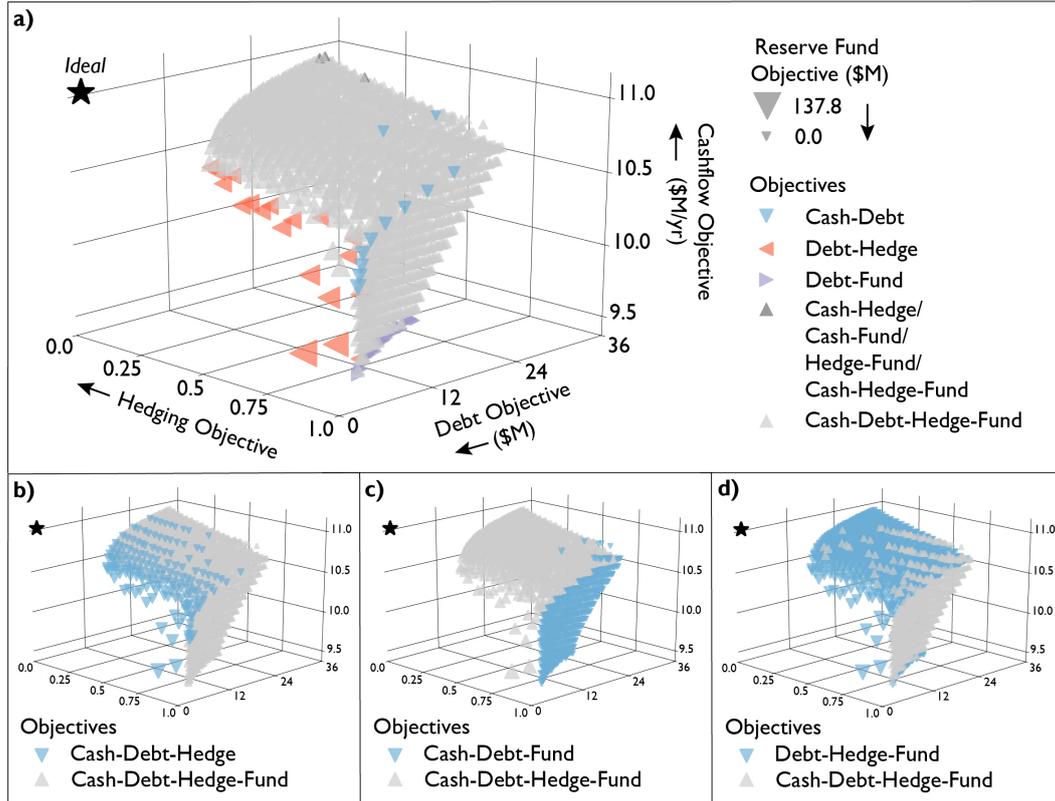


Figure 7. Visualization of Pareto approximate sets for different sub-problems. Colored points represent solutions that are non-dominated with respect to a particular sub-problem; for example, orange points in sub-figure (a) represent solutions that are non-dominated with respect to J^{debt} and J^{hedge} . Light grey points in all sub-figures represent solutions from the 4-objective problem that are not captured in the lower-dimensional problems.

772 More broadly, the lower-dimensional sub-problems tend to produce Pareto approx-
 773 imate sets that are near the extreme boundaries of the larger-dimensional problem. Sub-
 774 Figure 7a includes four sub-problems for which the Pareto approximate set consists of
 775 a single solution ($J^{cash}_{-}J^{hedge}$, $J^{cash}_{-}J^{fund}$, $J^{hedge}_{-}J^{fund}$, $J^{cash}_{-}J^{hedge}_{-}J^{fund}$). Each

776 of these sub-problems excludes debt, leading to a single optimal policy that performs es-
 777 sentially no risk management. This is consistent with prior work finding that conflicts
 778 in higher-dimensional problems can remain hidden in lower-dimensional sub-problems
 779 (Kollat & Reed, 2007; Matrosov et al., 2015; Woodruff et al., 2013). Sub-Figure 7a also
 780 shows results for the $J^{cash}\text{-}J^{debt}$, $J^{debt}\text{-}J^{hedge}$, and $J^{debt}\text{-}J^{fund}$ sub-problems. Each sub-
 781 set of solutions is concentrated along an outer border of the larger Pareto front, where
 782 performance of the two explicitly-considered objectives is optimized at the expense of
 783 the other two objectives. The same pattern is evident in the 3-objective sub-problems
 784 of Sub-Figures 7b ($J^{cash}\text{-}J^{debt}\text{-}J^{hedge}$), 7c ($J^{cash}\text{-}J^{debt}\text{-}J^{fund}$), and 7d ($J^{debt}\text{-}J^{hedge}\text{-}J^{fund}$).
 785 These solution sets are larger, but still occupy extremal regions of the overall Pareto front.
 786 Thus, by choosing to optimize a 2- or 3-objective sub-problem, decision-makers may un-
 787 wittingly produce an incomplete and biased Pareto approximate set.

788 The larger-dimensional problem leads to a fuller set of alternatives that better rep-
 789 represents the tradeoffs associated with decision-maker preferences for different financial risk
 790 management tools. However, it is a non-trivial task to select a single operating policy
 791 from among the large Pareto approximate set. Interactive visualization approaches can
 792 help with this task. One example is to allow decision-makers to apply *a posteriori* per-
 793 formance criteria and “brush away” solutions that fail to meet these constraints (Kasprzyk,
 794 Nataraj, Reed, & Lempert, 2013). The strictness of the constraints can be iteratively
 795 increased until decision-makers are relatively agnostic about the tradeoffs across the fea-
 796 sible solution set. For example, consider a utility whose financial team (perhaps in con-
 797 sultation with its regulatory commission) develops the following criteria: if $\bar{R} = \$10.99\text{M}$
 798 is the mean annual net hydropower revenue in the absence of any risk management, then
 799 (1) the risk management policy should not reduce expected annualized cash flows by more
 800 than 2.5% ($J^{cash} \geq 0.975\bar{R}$); (2) the utility should rarely be forced to borrow more than
 801 150% of mean net revenue to cover cash flow deficits ($J^{debt} \leq 1.5\bar{R}$); and (3) the util-
 802 ity should not maintain reserves larger 150% of mean net revenue ($J^{fund} \leq 1.5\bar{R}$). These
 803 constraints drastically reduce the set of feasible solutions (Figure 8). At this point, a quan-
 804 titative method such as TOPSIS (Behzadian et al., 2012; Roszkowska, 2011) can be used
 805 to select one of the remaining policies for the utility to use (e.g., the policy outlined in
 806 Figure 8 and listed in Row 3 of Table 2).

807 While these constraints could, in theory, be applied *a priori* and used to reduce
 808 the number of objectives in the optimization, it is very difficult in practice for decision-

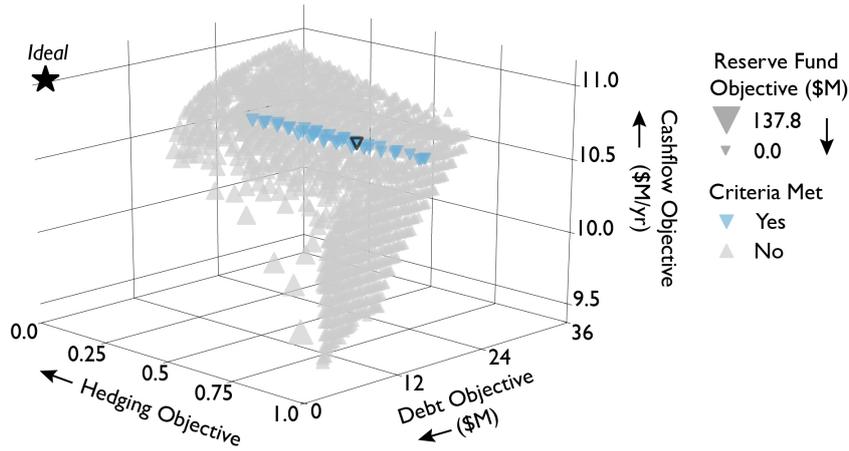


Figure 8. Set of feasible solutions after filtering for stakeholder-determined *a posteriori* constraints. The best compromise policy from the feasible set is outlined in black and described in Row 3 of Table 2

809 makers to effectively set the constraint values without first understanding the topology
 810 of the tradeoff surface (Kasprzyk et al., 2016; Spronk et al., 2005). This highlights the
 811 value of the EMODPS approach, which is scalable to extremely large problems on mod-
 812 ern high-performance computing infrastructure (Giuliani et al., 2018; Zatarain Salazar
 813 et al., 2016), suggesting that the formulation used here could be expanded to include ad-
 814 ditional objectives such as customer rates, social equity, and environmental quality. Ad-
 815 ditionally, future work should consider the effects of alternative problem framings; for
 816 example, a decision-maker may prefer a risk metric based on cash flow semi-variance (Tur-
 817 vey & Nayak, 2003), or a hedging objective that seeks to maximize year-to-year stabil-
 818 ity for planning purposes (Quinn, Reed, & Keller, 2017). In practice, researchers and stake-
 819 holders can iteratively refine the multi-objective problem in a way that matches their
 820 intuitions and goals (Smith, Kasprzyk, & Dilling, 2017; Wu et al., 2016) while balanc-
 821 ing the accuracy of the Monte Carlo estimator and the tractability of the search (Kasprzyk
 822 et al., 2012; Quinn, Reed, Giuliani, & Castelletti, 2017; Zatarain Salazar et al., 2017).

823 5.3 Value of state information for control

824 As demonstrated above, the EMODPS method can be used to develop control poli-
 825 cies that perform well across a range of stakeholder preferences. However, decision-makers
 826 may be unwilling to adopt a complex, non-linear control policy if its operating rules re-

827 main opaque; it may be necessary to “open the black box” for users if they are to ap-
 828 ply such tools in practice (Castelvecchi, 2016; Quinn et al., 2019). Each policy represents
 829 a map from a vector of inputs (e.g., reserve fund balance) to its outputs (e.g., the hedg-
 830 ing contract slope). ITSA (Section 3.4) can help decision-makers to better understand
 831 how different policies respond to changing model state information. Figure 9 shows the
 832 hedging policy sensitivity indices for each solution in the Pareto approximate set, rep-
 833 resenting the degree to which each policy adjusts its annual hedging decision based on
 834 each of the three inputs: the reserve fund balance (η_F^H , Sub-Figure 9a), the debt (η_D^H ,
 835 9b), and the power price index (η_P^H , 9c). Each index is a measure of the importance of
 836 a particular input variable for controlling a state-aware policy; $\eta = 1$ implies that the
 837 policy is entirely controlled by the input, while $\eta = 0$ means that the input has no im-
 838 pact on the policy. Interestingly, Figure 9 shows that each input has a different region
 839 of “specialization” in objective space. The reserve fund balance is the most important
 840 input for policies along the top of the ridge where $J^{hedge} = 1$. These are policies that
 841 achieve a relatively low levels of debt and high levels of cash flow, in return for frequent
 842 hedging and a relatively large reserve fund. The debt information, on the other hand,
 843 is critical for policies occupying the swath of objective space with J^{hedge} between 0 and
 844 1. The power price index is less informative overall, but does provide value for policies
 845 along the bottom edge of the Pareto front with minimal reserve funds and debt.

846 In order to better understand how these policies utilize information, it is helpful
 847 to visualize the policies themselves. One high-sensitivity policy is chosen for each input
 848 (as outlined in Figure 9, and listed in Rows 4-6 of Table 2). Each policy is used to sim-
 849 ulate 20 random 20-year trajectories. The 400 resulting decisions are visualized in state-
 850 action space using parallel-coordinate plots (Figure 10). The first three vertical axes rep-
 851 resent the three hedging policy inputs (reserve fund balance, debt, power price index).
 852 The policy output (hedging contract slope) is represented by the fourth vertical axis as
 853 well as the colorbar to aid interpretation. Each colored line connecting the four axes rep-
 854 represents one of the 400 simulated decisions. These visualizations, in combination with the
 855 sensitivity indices, can be useful in understanding how each policy operates. For exam-
 856 ple, the policy in Sub-Figure 10a appears to hedge selectively, when the reserve fund bal-
 857 ance has fallen below a certain threshold. Above the threshold, no hedging contract is
 858 purchased, and below the threshold, the hedging slope increases as the fund balance falls.
 859 The policy in Sub-Figure 10b has a similar strategy, but structured around debt; hedg-

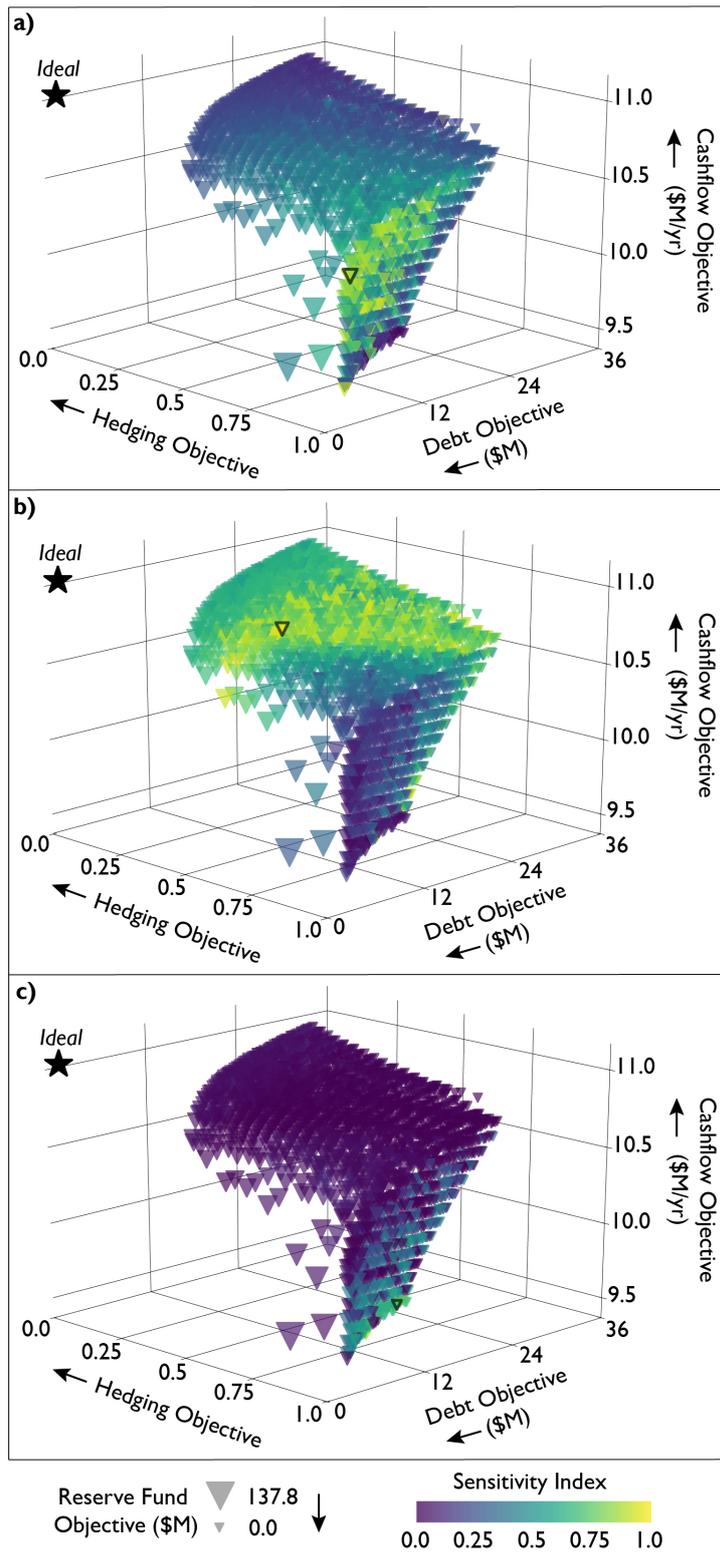


Figure 9. Information theoretic sensitivity indices, relative to hedging contract slope decision, for the (a) reserve fund balance; (b) debt; and (c) power price index. One high-sensitivity solution for each input is outlined in black and described in Rows 4-6 of Table 2

860 ing is zero below some threshold, and increases with debt above the threshold. Lastly,
861 the bottom policy always utilizes hedging contracts, the magnitude of which tend to be
862 inversely proportional to the power price index. Each of these patterns is consistent with
863 the sensitivity indices in Figure 9 and Table 2.

864 These plots can be used to build intuition about how the different risk management
865 policies achieve their competitive advantages. For example, compare the fund-sensitive
866 policy (a) to the debt-sensitive policy (b). The former maintains a relatively large re-
867 serve fund for its risk management needs, and uses hedging contracts as a substitute to
868 maintain its risk protection when the reserve fund is inadequate. This is qualitatively
869 similar behavior to the example policy simulated in Section 5.1 (Figure 5). The debt-
870 sensitive policy, on the other hand, keeps a much smaller reserve fund, which results in
871 more frequent cash flow shortfalls and debt during dry years. In order to reduce the like-
872 lihood of extreme debt spirals during longer droughts, this policy begins to use hedging
873 contracts when it has significant debt, and ceases hedging once it has paid off this debt.
874 The result is that the debt-sensitive policy is significantly more risky than the fund-sensitive
875 policy, but in return, it is less costly and requires less frequent hedging and a smaller re-
876 serve fund. The power-sensitive policy (c) takes a more consistent approach, purchas-
877 ing hedging contracts each year. This makes it the most expensive policy of the three
878 due to the cost of these contracts. However, the risk coverage from hedging allows it to
879 maintain a very small reserve fund and still avoid substantial debt. This policy also ad-
880 justs its hedging contract in response to projected wholesale power prices. If the power
881 price index is high, then the utility expects that its net revenue per unit of hydropower
882 will be higher than average, and vice versa when the index is low. By purchasing hedg-
883 ing contracts in inverse proportion to this index, the utility can dampen the overall vari-
884 ability of its combined cash flow (hydropower net revenue plus the net payout from the
885 hedging contract), and thus reduce its financial risk.

886 ITSA and policy visualization plots for the withdrawal/deposit decision can be found
887 in SI Figures S6-S7. The withdrawals and deposits are found to be much less sensitive
888 to model state information than hedging, suggesting that the gains from dynamic finan-
889 cial risk management in this study largely accrue from dynamic hedging rather than dy-
890 namic reserve fund management. This is consistent with past studies which have found
891 relatively simple optimal control rules for cash inventory problems; however, such stud-
892 ies often employ strict assumptions on the distribution and predictability of incoming

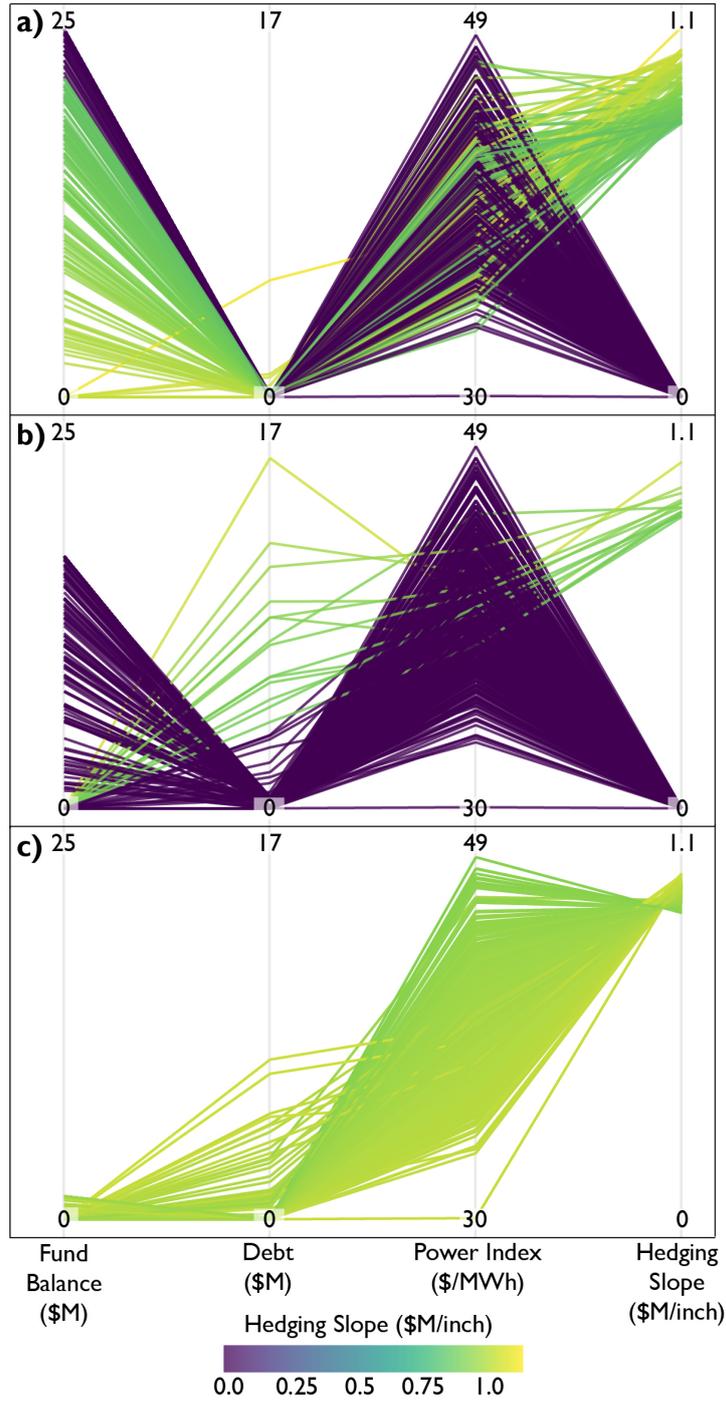


Figure 10. Hedging control policy visualization for three chosen policies in Figure 9 and Rows 4-6 of Table 2. Policies (a), (b), and (c) are highly sensitive to the reserve fund, debt, and power price index information, respectively. The first three vertical axes represent the three inputs, while the fourth axis and the colorbar represent the hedging action. Each line connecting the four axes represents one state-action combination experienced within a simulation.

893 cash flows, and their generalizability to real-world situations is uncertain (da Costa Moraes
894 et al., 2015). A major advantage of the EMODPS approach is the flexibility of the non-
895 linear approximating network used to parameterize the policies. The RBF network is found
896 to identify complex policies for the hedging decision while maintaining relatively sim-
897 ple rules for the withdrawal/deposit decision (i.e., without over-fitting). This flexibility
898 is important when the decision analyst does not know the optimal rule form for each ac-
899 tion *a priori*. In problems with a larger number of candidate actions, an iterative scheme
900 for selecting the decisions most amenable to dynamic control would be beneficial.

901 One final takeaway from Figures 9 and 10 is that the most important model states
902 to include in a state-aware control policy can vary widely across the Pareto approximate
903 set. This implies that the most important input(s) cannot be known *a priori* without
904 accounting for decision-maker preferences. This is consistent with both analytical (Gra-
905 ham & Georgakakos, 2010; Tejada-Guibert, Johnson, & Stedinger, 1995) and empirical
906 (Hejazi et al., 2008) studies in the reservoir control literature, which have found that the
907 objective(s) of the operator can affect which hydrologic factors are deemed most infor-
908 mative. However, computational constraints often require that the total set of poten-
909 tially informative data be culled to a small subset of the most important variables. The
910 results of this study confirm the importance of accounting for the multi-dimensional na-
911 ture of information value during this process (Denaro et al., 2017; Giuliani et al., 2015).

912 **5.4 Limitations and future directions**

913 A limitation of this study is that the stochastic engine adopted from Hamilton et
914 al. (2020) assumes that wholesale power prices are independent from hydrology. In re-
915 ality, fluctuations in hydropower availability can impact wholesale prices across the west-
916 ern United States on multiple timescales (Su, Kern, Reed, & Characklis, 2020; Voisin et
917 al., 2018). This inverse correlation between streamflow and price could alter the utility’s
918 financial risk either for the better (e.g., higher prices received for hydropower sold dur-
919 ing drought) or for the worse (e.g., higher prices paid for replacement power). Future
920 work could integrate these factors into the adaptive hydro-financial risk model using an
921 economic power dispatch model (e.g., Su, Kern, Denaro, et al. (2020)) or a surrogate sta-
922 tistical model (e.g., Madani, Guégan, and Uvo (2014)), but this is beyond the scope of
923 the current investigation.

924 Another limitation of this study is the implicit assumption of stationarity embed-
925 ded in the stochastic engine adopted from Hamilton et al. (2020). Despite this fact, Fig-
926 ure 5 suggests that the EMODPS-derived policies trained on a stationary Monte Carlo
927 ensemble can perform relatively well across a wide range of potential outcomes, many
928 of which are extreme compared to historical data. Additionally, the present study con-
929 cerns purely financial decisions on relatively short time scales, for which interannual cli-
930 mate variability is expected to overwhelm longer-term non-stationarity (Lehner et al.,
931 2020). The reader is referred to Hamilton et al. (2020) for further discussion of these is-
932 sues. Nonetheless, future studies should consider a broader analysis of the impacts of chang-
933 ing climate, markets, etc., on the robustness of adaptive financial risk management strate-
934 gies for hydropower production. This would be especially important if combined with
935 dynamic infrastructure investments (Haasnoot, Kwakkel, Walker, & ter Maat, 2013; Kwakkel,
936 Haasnoot, & Walker, 2015; Zeff, Herman, Reed, & Characklis, 2016), since climate un-
937 certainties become increasingly important for long-term, irreversible decisions (Doss-Gollin,
938 Farnham, Steinschneider, & Lall, 2019; Stakhiv, 2011). Statistical learning approaches
939 can be used to update decision-making based on evolving beliefs about the non-stationary
940 hydro-financial system (Cohen, Zeff, & Herman, 2020; Fletcher, Lickley, & Strzepek, 2019;
941 Fletcher et al., 2017; Herman, Quinn, Steinschneider, Giuliani, & Fletcher, 2020). Ad-
942 ditionally, scenario discovery approaches can be used to search for financial risk man-
943 agement strategies that perform satisfactorily across a wide range of (perhaps deeply)
944 uncertain factors (Bryant & Lempert, 2010; Herman, Reed, Zeff, & Characklis, 2015; Kasprzyk
945 et al., 2013; Lempert, 2002).

946 **6 Conclusions**

947 A substantial body of literature has emerged around optimal control of water reser-
948 voir systems in the face of hydrologic uncertainty (Macian-Sorribes & Pulido-Velazquez,
949 2019). Evolutionary multi-objective direct policy search has emerged as an especially pow-
950 erful tool for overcoming the simultaneous curses of dimensionality, modeling, and mul-
951 tiple objectives that are characteristic of problems in the field (Giuliani et al., 2016, 2018).
952 This paper demonstrates that the same properties of EMODPS that make it ideal for
953 optimal reservoir control problems also make it well suited for the complex, multi-objective
954 financial risk management problems faced by water-reliant organizations as a result of
955 hydrologic variability. The methodology is applied in the context of the hydrologic fi-

956 nancial risk faced by the Power Enterprise of the San Francisco Public Utilities Com-
957 mission, an electricity producer relying primarily on hydropower. EMODPS is used to
958 develop control policies that dynamically balance the use of snowpack-based hedging con-
959 tracts, cash reserves, and debt, based on changing conditions within the model. Perfor-
960 mance is quantified based on four conflicting performance metrics: expected annualized
961 cash flow, 95th percentile maximum debt, expected hedging frequency, and expected max-
962 imum reserve fund balance. The first two metrics represent the classic return vs. risk
963 tradeoff in finance, while the second two metrics represent a decision-maker's preferences
964 for using one risk management instrument over another based on an organization's in-
965 dividual circumstances. By utilizing real-time model state information when making de-
966 cisions, the dynamic policies produced by EMODPS are found to significantly outper-
967 form policies produced under a more static control formulation akin to those commonly
968 used for financial risk management in the water resources literature. *A posteriori* visual
969 analytics and information theoretic sensitivity analysis can be used to help decision-makers
970 better understand how the complex, non-linear operating policies adapt to real-time in-
971 formation when making decisions.

972 The methodology developed in this paper should help decision-makers to better un-
973 derstand the dynamic relationships between hydrology, decision-making, and financial
974 outcomes, and thus facilitate more knowledgeable and effective management of hydro-
975 logic financial risks. Additionally, we note that while the interrelatedness of hydrology
976 and financial risk is conceptually useful for the present study (i.e., water resources re-
977 searchers and practitioners will easily grasp the similarities between reservoir control and
978 financial risk management), it does not represent a necessary condition for the useful-
979 ness of the dynamic financial risk management framework presented herein. In fact, a
980 broad class of financial risk management problems share a similar mathematical struc-
981 ture to reservoir control (i.e., multi-objective Markov Decision Processes). Although the
982 decision-making context and implementation details will vary, the overall framework pre-
983 sented here should thus be applicable to a wide variety of organizations, from water util-
984 ities exposed to hydrologic risk, to renewable energy developers exposed to wind risk,
985 to commodities firms exposed to interest rate risk.

986 Acknowledgments

987 Funding for this work was provided by the National Science Foundation (NSF), Inno-
 988 vations at the Nexus of Food-Energy-Water Systems, Track 2 (Award 1639268). The au-
 989 thors would like to thank Rohini Gupta for helpful discussion and code review, as well
 990 as Alexis Dufour and Darryl Dunn from the San Francisco Public Utilities Commission
 991 (SFPUC) for helpful discussion and data provision. The views expressed in this work
 992 represent those of the authors and do not necessarily reflect the views or policies of the
 993 NSF or SFPUC. All code and data for this project, including figure generation, are avail-
 994 able in a live repository ([https://github.com/ahamilton144/hamilton-2021-EMODPS](https://github.com/ahamilton144/hamilton-2021-EMODPS-financial-risk)
 995 [-financial-risk](https://github.com/ahamilton144/hamilton-2021-EMODPS-financial-risk)) and a permanent archive (<https://doi.org/10.5281/zenodo.5079786>).

996 References

- 997 Aguilar-Rivera, R., Valenzuela-Rendón, M., & Rodríguez-Ortiz, J. J. (2015,
 998 6). Genetic algorithms and Darwinian approaches in financial applica-
 999 tions: A survey. *Expert Systems with Applications*, *42*(21), 7684–7697. doi:
 1000 10.1016/j.eswa.2015.06.001
- 1001 Auder, B., & Iooss, B. (2009). Global sensitivity analysis based on entropy. *Safety,*
 1002 *Reliability and Risk Analysis: Theory, Methods and Applications - Proceedings*
 1003 *of the Joint ESREL and SRA-Europe Conference*, *3*, 2107–2115.
- 1004 Bank, M., & Wiesner, R. (2010). The Use of Weather Derivatives by Small-and
 1005 Medium- Sized Enterprises: Reasons and Obstacles. *Journal of Small Business*
 1006 *and Entrepreneurship*, *23*(4), 581–600. doi: 10.1080/08276331.2010.10593503
- 1007 Basdekas, L. (2014). Is multiobjective optimization ready for water resources prac-
 1008 titioners? Utility’s drought policy investigation. *Journal of Water Resources*
 1009 *Planning and Management*, *140*(3), 275–276. doi: 10.1061/(ASCE)WR.1943
 1010 -5452.0000415
- 1011 Basdekas, L., & Hayslett, R. (2021). *Improving Tradeoff Understanding in Water*
 1012 *Resource Planning Using Multi-Objective Search* (Tech. Rep.).
- 1013 Baumol, W. J. (1952). The Transactions Demand for Cash : An Inventory Theoretic
 1014 Approach. *The Quarterly Journal of Economics*, *66*(4), 545–556.
- 1015 Behzadian, M., Khanmohammadi Otaghsara, S., Yazdani, M., & Ignatius, J. (2012,
 1016 12). A state-of-the-art survey of TOPSIS applications. *Expert Systems with*
 1017 *Applications*, *39*, 13051–13069. doi: 10.1016/j.eswa.2012.05.056

- 1018 Bellman, R. (1957). *Dynamic Programming*. Princeton, NJ: Princeton University
1019 Press.
- 1020 Bertsekas, D. P. (2019). *Reinforcement learning and optimal control*. Nashua, NH:
1021 Athena Scientific.
- 1022 Blomfield, A., & Plummer, J. (2014). The allocation and documentation of hydro-
1023 logical risk. *Hydropower & Dams*(5), 94–108. doi: 10.1093/rfs/15.4.1283
- 1024 Bodnar, G. M., Giambona, E., Graham, J. R., & Harvey, C. R. (2019). A view in-
1025 side corporate risk management. *Management Science*, 65(11), 5001–5026. doi:
1026 10.1287/mnsc.2018.3081
- 1027 Bolton, P., Chen, H., & Wang, N. (2011). A Unified Theory of Tobin’s q, Corpo-
1028 rate Investment, Financing, and Risk Management. *Journal of Finance*, 66(5),
1029 1545–1578. doi: 10.1111/j.1540-6261.2011.01681.x
- 1030 Borgonovo, E. (2007). A new uncertainty importance measure. *Reliability Engineer-
1031 ing and System Safety*, 92(6), 771–784. doi: 10.1016/j.res.2006.04.015
- 1032 Borgonovo, E., Castaings, W., & Tarantola, S. (2011). Moment Independent Impor-
1033 tance Measures: New Results and Analytical Test Cases. *Risk Analysis*, 31(3),
1034 404–428. doi: 10.1111/j.1539-6924.2010.01519.x
- 1035 Brown, C. M., & Carriquiry, M. (2007). Managing hydroclimatological risk to water
1036 supply with option contracts and reservoir index insurance. *Water Resources
1037 Research*, 43, W11423. doi: 10.1029/2007WR006093
- 1038 Brown, C. M., Lund, J. R., Cai, X., Reed, P. M., Zagona, E. A., Ostfeld, A., . . .
1039 Brekke, L. (2015). The future of water resources systems analysis: Toward
1040 a scientific framework for sustainable water management. *Water Resources
1041 Research*, 51(8), 6110–6124. doi: 10.1002/2015WR017114
- 1042 Bryant, B. P., & Lempert, R. J. (2010). Thinking inside the box: A participatory,
1043 computer-assisted approach to scenario discovery. *Technological Forecasting
1044 and Social Change*, 77(1), 34–49. doi: 10.1016/j.techfore.2009.08.002
- 1045 Cai, X., Wallington, K., Shafiee-Jood, M., & Marston, L. (2018). Understand-
1046 ing and managing the food-energy-water nexus - opportunities for water re-
1047 sources research. *Advances in Water Resources*, 111(November), 259–273. doi:
1048 10.1016/j.advwatres.2017.11.014
- 1049 Castelletti, A., Pianosi, F., & Soncini-Sessa, R. (2008). Water reservoir control un-
1050 der economic, social and environmental constraints. *Automatica*, 44(6), 1595–

- 1051 1607. doi: 10.1016/j.automatica.2008.03.003
- 1052 Castelletti, A., & Soncini-Sessa, R. (2006). A procedural approach to strengthening
1053 integration and participation in water resource planning. *Environmental Mod-*
1054 *elling and Software*, 21(10), 1455–1470. doi: 10.1016/j.envsoft.2005.07.013
- 1055 Castelvechi, D. (2016). Can we open the black box of AI? *Nature*, 538(7623), 20–
1056 23. doi: 10.1038/538020a
- 1057 Ceres. (2017). *The Investor Water Toolkit* (Tech. Rep.). Boston, MA: Ceres.
1058 Retrieved from [https://www.ceres.org/sites/default/files/reports/](https://www.ceres.org/sites/default/files/reports/Ceres_InvestWaterToolkit.pdf)
1059 [Ceres_InvestWaterToolkit.pdf](https://www.ceres.org/sites/default/files/reports/Ceres_InvestWaterToolkit.pdf)
- 1060 Chapman, T. A., & Breeding, J. M. (2014). *U.S. Public Finance Waterworks,*
1061 *Sanitary Sewer, and Drainage Utility Systems: Methodology And Assump-*
1062 *tions* (Tech. Rep.). Standard and Poor’s Ratings Services. Retrieved from
1063 https://www.spratings.com/documents/20184/908554/US_PF_Event
1064 [_RFCRndTblsJan2015_Article1/30d125eb-1066-4730-8ab1-f2cd6a6d6f9a](https://www.spratings.com/documents/20184/908554/US_PF_Event)
- 1065 Coello Coello, C. A., Lamont, G. B., & Van Veldhuizen, D. A. (2007). *Evolution-*
1066 *ary Algorithms for Solving Multi-Objective Problems* (2nd ed.). New York, NY:
1067 Springer Science+Business Media, LLC. doi: 10.1046/j.1365-2672.2000.00969
1068 .x
- 1069 Cohen, J. S., Zeff, H. B., & Herman, J. D. (2020). Adaptation of Multiobjective
1070 Reservoir Operations to Snowpack Decline in the Western United States. *Jour-*
1071 *nal of Water Resources Planning and Management*, 146(12), 04020091. doi: 10
1072 .1061/(asce)wr.1943-5452.0001300
- 1073 da Costa Moraes, M. B., & Nagano, M. S. (2013). Cash Management Policies By
1074 Evolutionary Models: A Comparison Using The Miller-Orr Model. *Journal*
1075 *of Information Systems and Technology Management*, 10(3), 561–576. doi:
1076 10.4301/s1807-17752013000300006
- 1077 da Costa Moraes, M. B., Nagano, M. S., & Sobreiro, V. A. (2015). Stochastic
1078 Cash Flow Management Models: A Literature Review Since the 1980s. In
1079 P. Guarnieri (Ed.), *Decision models in engineering and management* (pp.
1080 11–28). Springer International Publishing Switzerland. doi: 10.1007/
1081 978-3-319-11949-6{-}2
- 1082 de Almeida-Filho, A. T., de Lima Silva, D. F., & Ferreira, L. (2020). Finan-
1083 cial modelling with multiple criteria decision making: A systematic liter-

- 1084 ature review. *Journal of the Operational Research Society*, 1–19. doi:
 1085 10.1080/01605682.2020.1772021
- 1086 Denaro, S., Anghileri, D., Giuliani, M., & Castelletti, A. (2017). Informing the
 1087 operations of water reservoirs over multiple temporal scales by direct use of
 1088 hydro-meteorological data. *Advances in Water Resources*, 103, 51–63. doi:
 1089 10.1016/j.advwatres.2017.02.012
- 1090 Denaro, S., Castelletti, A., Giuliani, M., & Characklis, G. (2020). Insurance Port-
 1091 folio Diversification Through Bundling for Competing Agents Exposed to
 1092 Uncorrelated Drought and Flood Risks. *Water Resources Research*, 56(5),
 1093 1–20. doi: 10.1029/2019WR026443
- 1094 Desreumaux, Q., Côté, P., & Leconte, R. (2018). Comparing model-based and
 1095 model-free streamflow simulation approaches to improve hydropower reservoir
 1096 operations. *Journal of Water Resources Planning and Management*, 144(3),
 1097 1–10. doi: 10.1061/(ASCE)WR.1943-5452.0000860
- 1098 di Pierro, F., Khu, S. T., & Savić, D. A. (2007). An investigation on preference or-
 1099 der ranking scheme for multiobjective evolutionary optimization. *IEEE Trans-*
 1100 *actions on Evolutionary Computation*, 11(1), 17–45. doi: 10.1109/TEVC.2006
 1101 .876362
- 1102 Disatnik, D., Duchin, R. A. N., & Schmidt, B. (2014). Cash Flow Hedging and Liq-
 1103 uidity Choices. *Review of Finance*, 18, 715–748. doi: 10.1093/rof/rft006
- 1104 D’Odorico, P., Davis, K. F., Rosa, L., Carr, J. A., Chiarelli, D., Dell’Angelo, J.,
 1105 ... Rulli, M. C. (2018). The Global Food-Energy-Water Nexus. *Reviews of*
 1106 *Geophysics*, 56(3), 456–531. doi: 10.1029/2017RG000591
- 1107 Doss-Gollin, J., Farnham, D. J., Steinschneider, S., & Lall, U. (2019). Robust Adap-
 1108 tation to Multiscale Climate Variability. *Earth’s Future*, 7(7), 734–747. doi: 10
 1109 .1029/2019EF001154
- 1110 Fabozzi, F. J., Focardi, S., & Jonas, C. (2007). Trends in quantitative equity man-
 1111 agement: Survey results. *Quantitative Finance*, 7(2), 115–122. doi: 10.1080/
 1112 14697680701195941
- 1113 Fletcher, S., Lickley, M., & Strzepek, K. (2019). Learning about climate change
 1114 uncertainty enables flexible water infrastructure planning. *Nature Communica-*
 1115 *tions*, 10(1), 1–11. doi: 10.1038/s41467-019-09677-x
- 1116 Fletcher, S., Miotti, M., Swaminathan, J., Klemun, M., Strzepek, K., & Siddiqi,

- 1117 A. (2017). Water supply infrastructure planning decision-making frame-
 1118 work to classify multiple uncertainties and evaluate flexible design. *Journal*
 1119 *of Water Resources Planning and Management*, *143*(10), 04017061. doi:
 1120 10.1061/(ASCE)WR.1943-5452.0000823
- 1121 Flood, M. D., Lemieux, V. L., Varga, M., & William Wong, B. L. (2016). The appli-
 1122 cation of visual analytics to financial stability monitoring. *Journal of Financial*
 1123 *Stability*, *27*, 180–197. doi: 10.1016/j.jfs.2016.01.006
- 1124 Foster, B. T., Kern, J. D., & Characklis, G. W. (2015). Mitigating hydrologic fi-
 1125 nancial risk in hydropower generation using index-based financial instruments.
 1126 *Water Resources and Economics*, *10*, 45–67. doi: 10.1016/j.wre.2015.04.001
- 1127 Froot, K., Scharfstein, D., & Stein, J. (1993). Risk Management: Coordinating
 1128 Corporate Investment and Financing Policies. *The Journal of Finance*, *48*(5),
 1129 1629–1648. doi: 10.1111/j.1540-6261.1993.tb05123.x
- 1130 Giuliani, M., & Castelletti, A. (2016). Is robustness really robust? How different def-
 1131 initions of robustness impact decision-making under climate change. *Climatic*
 1132 *Change*, *135*(3-4), 409–424. doi: 10.1007/s10584-015-1586-9
- 1133 Giuliani, M., Castelletti, A., Pianosi, F., Mason, E., & Reed, P. M. (2016). Curses,
 1134 tradeoffs, and scalable management: Advancing evolutionary multiobjec-
 1135 tive direct policy search to improve water reservoir operations. *Journal*
 1136 *of Water Resources Planning and Management*, *142*(2), 04015050. doi:
 1137 10.1061/(ASCE)WR.1943-5452.0000570
- 1138 Giuliani, M., Herman, J. D., Castelletti, A., & Reed, P. (2014). Many-objective
 1139 reservoir policy identification and refinement to reduce policy inertia and my-
 1140 opia in water management. *Water Resources Research*, *50*(4), 3355–3377. doi:
 1141 10.1002/2013WR014700
- 1142 Giuliani, M., Mason, E., Castelletti, A., Pianosi, F., & Soncini-Sessa, R. (2014).
 1143 Universal approximators for direct policy search in multi-purpose water reser-
 1144 voir management: A comparative analysis. In *Proceedings of the 19th world*
 1145 *congress, international federation of automatic control* (Vol. 19, pp. 6234–
 1146 6239). Cape Town, South Africa. doi: 10.3182/20140824-6-za-1003.01962
- 1147 Giuliani, M., Pianosi, F., & Castelletti, A. (2015). Making the most of data:
 1148 An information selection and assessment framework to improve water sys-
 1149 tems operations. *Water Resources Research*, *51*(11), 9073–9093. doi:

- 1150 10.1002/2015WR017044
- 1151 Giuliani, M., Quinn, J. D., Herman, J. D., Castelletti, A., & Reed, P. M. (2018).
 1152 Scalable Multiobjective Control for Large-Scale Water Resources Systems un-
 1153 der Uncertainty. *IEEE Transactions on Control Systems Technology*, *26*(4),
 1154 1492–1499. doi: 10.1109/TCST.2017.2705162
- 1155 Graham, N. E., & Georgakakos, K. P. (2010). Toward understanding the value of
 1156 climate information for multiobjective reservoir management under present
 1157 and future climate and demand scenarios. *Journal of Applied Meteorology and*
 1158 *Climatology*, *49*(4), 557–573. doi: 10.1175/2009JAMC2135.1
- 1159 Gupta, R. S., Hamilton, A. L., Reed, P. M., & Characklis, G. W. (2020). Can mod-
 1160 ern multi-objective evolutionary algorithms discover high-dimensional financial
 1161 risk portfolio tradeoffs for snow-dominated water-energy systems? *Advances in*
 1162 *Water Resources*, *145*, 103718. doi: 10.1016/j.advwatres.2020.103718
- 1163 Haasnoot, M., Kwakkel, J. H., Walker, W. E., & ter Maat, J. (2013). Dy-
 1164 namic adaptive policy pathways: A method for crafting robust deci-
 1165 sions for a deeply uncertain world. *Global Environmental Change*. doi:
 1166 10.1016/j.gloenvcha.2012.12.006
- 1167 Hadka, D., & Reed, P. (2012). Diagnostic assessment of search controls and fail-
 1168 ure modes in many-objective evolutionary optimization. *Evolutionary Compu-*
 1169 *tation*, *20*(3), 423–452. doi: 10.1162/EVCO{-}a{-}00053
- 1170 Hadka, D., & Reed, P. (2013). Borg: An auto-adaptive many-objective evolution-
 1171 ary computing framework. *Evolutionary Computation*, *21*(2), 231–259. doi: 10
 1172 .1162/EVCO{-}a{-}00075
- 1173 Hamilton, A. L., Characklis, G. W., & Reed, P. M. (2020). Managing financial risk
 1174 tradeoffs for hydropower generation using snowpack-based index contracts.
 1175 *Water Resources Research*, *56*, e2020WR027212. doi: 10.1029/2020wr027212
- 1176 Hejazi, M. I., Cai, X., & Ruddell, B. L. (2008). The role of hydrologic information in
 1177 reservoir operation - Learning from historical releases. *Advances in Water Re-*
 1178 *sources*, *31*(12), 1636–1650. doi: 10.1016/j.advwatres.2008.07.013
- 1179 Herman, J. D., Quinn, J. D., Steinschneider, S., Giuliani, M., & Fletcher, S. (2020).
 1180 Climate adaptation as a control problem: Review and perspectives on dynamic
 1181 water resources planning under uncertainty. *Water Resources Research*, *56*,
 1182 e24389. doi: 10.1029/2019wr025502

- 1183 Herman, J. D., Reed, P. M., Zeff, H. B., & Characklis, G. W. (2015). How should
 1184 robustness be defined for water systems planning under change? *Journal*
 1185 *of Water Resources Planning and Management*, *141*(10), 04015012. doi:
 1186 10.1061/(ASCE)WR.1943-5452.0000509
- 1187 Herman, J. D., Zeff, H. B., Reed, P. M., & Characklis, G. W. (2014). Beyond op-
 1188 timality: Multistakeholder robustness tradeoffs for regional water portfolio
 1189 planning under deep uncertainty. *Water Resources Research*, *50*(10), 7692–
 1190 7713. doi: 10.1002/2014WR015338
- 1191 Hughes, J., & Leurig, S. (2013). *Assessing water system revenue risk: Considera-*
 1192 *tions for market analysts* (Tech. Rep.). Boston, MA: Ceres and Environmental
 1193 Finance Center at the University of North Carolina at Chapel Hill.
- 1194 Hughes, J., Tiger, M., Eskaf, S., Berahzer, S. I., Royster, S., Boyle, C., & Batten, D.
 1195 (2014). *Defining a Resilient Business Model for Water Utilities* (Tech. Rep.).
 1196 Water Research Foundation. Retrieved from [https://efc.sog.unc.edu/](https://efc.sog.unc.edu/sites/default/files/4366_Exec_Summary_0.pdf)
 1197 [sites/default/files/4366_Exec_Summary_0.pdf](https://efc.sog.unc.edu/sites/default/files/4366_Exec_Summary_0.pdf)
- 1198 Hull, J. C. (2009). *Options, Futures, and Other Derivatives* (8th ed.). Boston, MA:
 1199 Prentice Hall.
- 1200 Huskova, I., Matrosov, E. S., Harou, J. J., Kasprzyk, J. R., & Lambert, C. (2016).
 1201 Screening robust water infrastructure investments and their trade-offs under
 1202 global change: A London example. *Global Environmental Change*, *41*, 216–
 1203 227. doi: 10.1016/j.gloenvcha.2016.10.007
- 1204 Iooss, B., & Lemaître, P. (2015). A review on global sensitivity analysis methods.
 1205 *Operations Research/ Computer Science Interfaces Series*, *59*, 101–122. doi:
 1206 10.1007/978-1-4899-7547-8{_}5
- 1207 Kasprzyk, J. R., Nataraj, S., Reed, P. M., & Lempert, R. J. (2013). Many ob-
 1208 jective robust decision making for complex environmental systems under-
 1209 going change. *Environmental Modelling and Software*, *42*, 55–71. doi:
 1210 10.1016/j.envsoft.2012.12.007
- 1211 Kasprzyk, J. R., Reed, P. M., Characklis, G. W., & Kirsch, B. R. (2012). Many-
 1212 objective de Novo water supply portfolio planning under deep uncertainty. *En-*
 1213 *vironmental Modelling and Software*, *34*, 87–104. doi: 10.1016/j.envsoft.2011
 1214 .04.003
- 1215 Kasprzyk, J. R., Reed, P. M., & Hadka, D. M. (2016). Battling Arrow’s Paradox

- 1216 to discover robust water management alternatives. *Journal of Water Resources*
 1217 *Planning and Management*, 142(2), 04015053. doi: 10.1061/(ASCE)WR.1943
 1218 -5452.0000572
- 1219 Kollat, J. B., & Reed, P. (2007). A framework for Visually Interactive Decision-
 1220 making and Design using Evolutionary Multi-objective Optimization
 1221 (VIDEO). *Environmental Modelling and Software*, 22(12), 1691–1704. doi:
 1222 10.1016/j.envsoft.2007.02.001
- 1223 Koutsyiannis, D., & Economou, A. (2003). Evaluation of the parameterization-
 1224 simulation-optimization approach for the control of reservoir systems. *Water*
 1225 *Resources Research*, 39(6). doi: 10.1029/2003WR002148
- 1226 Kouwenberg, R., & Zenios, S. A. (2008). Stochastic Programming Models for Asset
 1227 Liability Management. In *Handbook of asset and liability management* (Vol. 1,
 1228 pp. 253–303). doi: 10.1016/B978-044453248-0.50012-5
- 1229 Krause, T. A., & Tse, Y. (2016). Risk management and firm value: Recent theory
 1230 and evidence. *International Journal of Accounting and Information Manage-*
 1231 *ment*, 24(1), 56–81. doi: 10.1108/IJAIM-05-2015-0027
- 1232 Krzykacz-Hausmann, B. (2001). Epistemic sensitivity analysis based on the concept
 1233 of entropy. In P. Prado & R. Bolado (Eds.), *Proceedings of samo* (pp. 31–35).
 1234 Madrid: CIEMAT.
- 1235 Kwakkel, J. H., Haasnoot, M., & Walker, W. E. (2015). Developing dynamic
 1236 adaptive policy pathways: a computer-assisted approach for developing
 1237 adaptive strategies for a deeply uncertain world. *Climatic Change*. doi:
 1238 10.1007/s10584-014-1210-4
- 1239 Labadie, J. W. (2004). Optimal operation of multireservoir systems: State-of-the-art
 1240 review. *Journal of Water Resources Planning and Management*, 130(2), 93–
 1241 111. doi: 10.1061/(ASCE)0733-9496(2004)130:2(93)
- 1242 Larson, W. M., Freedman, P. L., Passinsky, V., Grubb, E., & Adriaens, P. (2012).
 1243 Mitigating corporate water risk: Financial market tools and supply manage-
 1244 ment strategies. *Water Alternatives*, 5(3), 582–602.
- 1245 Lehner, F., Deser, C., Maher, N., Marotzke, J., Fischer, E., Brunner, L., ...
 1246 Hawkins, E. (2020). Partitioning climate projection uncertainty with mul-
 1247 tiple Large Ensembles and CMIP5/6. *Earth System Dynamics*, 11, 491–508.
 1248 doi: 10.5194/esd-2019-93

- 1249 Lempert, R. J. (2002). A new decision sciences for complex systems. *Proceedings of*
 1250 *the National Academy of Sciences of the United States of America*, 99(SUPPL.
 1251 3), 7309–7313. doi: 10.1073/pnas.082081699
- 1252 Leurig, S. (2010). *The Ripple Effect: Water risk in the municipal bond market*
 1253 (Tech. Rep.). Boston, MA: Ceres. Retrieved from [https://www.ceres.org/](https://www.ceres.org/resources/reports/ripple-effect-water-risk-municipal-bond-market)
 1254 [resources/reports/ripple-effect-water-risk-municipal-bond-market](https://www.ceres.org/resources/reports/ripple-effect-water-risk-municipal-bond-market)
- 1255 Liu, H., Chen, W., & Sudjianto, A. (2006). Relative entropy based method for prob-
 1256 abilistic sensitivity analysis in engineering design. *Journal of Mechanical De-*
 1257 *sign*, 128(2), 326–336.
- 1258 Macian-Sorribes, H., & Pulido-Velazquez, M. (2019). Inferring efficient operat-
 1259 ing rules in multireservoir water resource systems: A review. *Wiley Interdisci-*
 1260 *plinary Reviews: Water*, 7(1), e1400. doi: 10.1002/wat2.1400
- 1261 Madani, K., Guégan, M., & Uvo, C. B. (2014). Climate change impacts on high-
 1262 elevation hydroelectricity in California. *Journal of Hydrology*, 510, 153–163.
 1263 doi: 10.1016/j.jhydrol.2013.12.001
- 1264 Maestro, T., Barnett, B. J., Coble, K. H., Garrido, A., & Bielza, M. (2016). Drought
 1265 index insurance for the Central Valley Project in California. *Applied Economic*
 1266 *Perspectives and Policy*, 38(3), 521–545. doi: 10.1093/aep/ppw013
- 1267 Maier, H. R., Kapelan, Z., Kasprzyk, J., Kollat, J., Matott, L. S., Cunha, M. C.,
 1268 ... Reed, P. M. (2014). Evolutionary algorithms and other metaheuris-
 1269 tics in water resources: Current status, research challenges and future
 1270 directions. *Environmental Modelling and Software*, 62, 271–299. doi:
 1271 10.1016/j.envsoft.2014.09.013
- 1272 Markowitz, H. (1952). Portfolio Selection. *The Journal of Finance*, 7(1), 77–91.
- 1273 Marqués, A. I., García, V., & Sánchez, J. S. (2020). Ranking-based MCDM
 1274 models in financial management applications: analysis and emerging chal-
 1275 lenges. *Progress in Artificial Intelligence*, 9(3), 171–193. doi: 10.1007/
 1276 s13748-020-00207-1
- 1277 Matrosov, E. S., Huskova, I., Kasprzyk, J. R., Harou, J. J., Lambert, C., & Reed,
 1278 P. M. (2015). Many-objective optimization and visual analytics reveal key
 1279 trade-offs for London’s water supply. *Journal of Hydrology*, 531, 1040–1053.
 1280 doi: 10.1016/j.jhydrol.2015.11.003
- 1281 Meyer, E. S., Characklis, G. W., Brown, C., & Moody, P. (2016). Hedging the fi-

- 1282 nancial risk from water scarcity for Great Lakes shipping. *Water Resources Re-*
1283 *search*, 52(1), 227–245. doi: 10.1002/2015WR017855
- 1284 Miller, M. M., & Orr, D. (1966). A Model of the Demand for Money by Firms. *The*
1285 *Quarterly Journal of Economics*, 80(3), 413–435.
- 1286 Moallemi, E. A., Kwakkel, J., de Haan, F. J., & Bryan, B. A. (2020, 11). Ex-
1287 ploratory modeling for analyzing coupled human-natural systems under
1288 uncertainty. *Global Environmental Change*, 65, 102186. doi: 10.1016/
1289 j.gloenvcha.2020.102186
- 1290 Moody’s Investors Service. (2011). *Rating Methodology: U.S. Public Power Electric*
1291 *Utilities with Generation Ownership Exposure* (Tech. Rep.).
- 1292 Moody’s Investors Service. (2019). *Bonneville Power Administration , OR: Credit*
1293 *Update Following Rating Affirmation and Change in Outlook to Negative*
1294 (Tech. Rep.).
- 1295 Mortensen, E., & Block, P. (2018). ENSO Index-Based Insurance for Agricul-
1296 tural Protection in Southern Peru. *Geosciences*, 8(2), 64. doi: 10.3390/
1297 geosciences8020064
- 1298 Mulvey, J. M. (2001). Introduction to financial optimization: Mathematical Pro-
1299 gramming Special Issue. *Mathematical Programming*, 89(2), 205–216. doi: 10
1300 .1007/pl00011395
- 1301 Nicklow, J., Reed, P., Savic, D., Dessalegne, T., Harrell, L., Chan-Hilton, A.,
1302 ... Zechman, E. (2010). State of the Art for Genetic Algorithms and
1303 Beyond in Water Resources Planning and Management. *Journal of Wa-*
1304 *ter Resources Planning and Management*, 136(4), 412–432. doi: 10.1061/
1305 ASCEWR.1943-5452.0000053
- 1306 Pappenberger, F., Beven, K. J., Ratto, M., & Matgen, P. (2008). Multi-method
1307 global sensitivity analysis of flood inundation models. *Advances in Water Re-*
1308 *sources*, 31(1), 1–14. doi: 10.1016/j.advwatres.2007.04.009
- 1309 Pardalos, P. M., Sandström, M., & Zopounidis, C. (1994). On the use of optimiza-
1310 tion models for portfolio selection: A review and some computational results.
1311 *Computational Economics*, 7(4), 227–244. doi: 10.1007/BF01299454
- 1312 Pérez-González, F., & Yun, H. (2013). Risk management and firm value: Evidence
1313 from weather derivatives. *Journal of Finance*, 68(5), 2143–2176. doi: 10.1111/
1314 jofi.12061

- 1315 Pianosi, F., Beven, K., Freer, J., Hall, J. W., Rougier, J., Stephenson, D. B., & Wa-
 1316 gener, T. (2016). Sensitivity analysis of environmental models: A systematic
 1317 review with practical workflow. *Environmental Modelling and Software*, *79*,
 1318 214–232. doi: 10.1016/j.envsoft.2016.02.008
- 1319 Ponsich, A., Jaimes, A. L., & Coello Coello, C. A. (2013). A survey on multiobjec-
 1320 tive evolutionary algorithms for the solution of the portfolio optimization prob-
 1321 lem and other finance and economics applications. *IEEE Transactions on Evo-*
 1322 *lutionary Computation*, *17*(3), 321–344. doi: 10.1109/TEVC.2012.2196800
- 1323 Powell, W. B. (2019). A unified framework for stochastic optimization. *European*
 1324 *Journal of Operational Research*, *275*(3), 795–821. doi: 10.1016/j.ejor.2018.07
 1325 .014
- 1326 Quinn, J. D., Reed, P. M., Giuliani, M., & Castelletti, A. (2017). Rival framings:
 1327 A framework for discovering how problem formulation uncertainties shape risk
 1328 management trade-offs in water resources systems. *Water Resources Research*,
 1329 *53*(8), 7208–7233. doi: 10.1002/2017WR020524
- 1330 Quinn, J. D., Reed, P. M., Giuliani, M., & Castelletti, A. (2019). What Is Control-
 1331 ling Our Control Rules? Opening the Black Box of Multireservoir Operating
 1332 Policies Using Time-Varying Sensitivity Analysis. *Water Resources Research*,
 1333 *55*(7), 5962–5984. doi: 10.1029/2018WR024177
- 1334 Quinn, J. D., Reed, P. M., Giuliani, M., Castelletti, A., Oyler, J. W., & Nicholas,
 1335 R. E. (2018). Exploring How Changing Monsoonal Dynamics and Hu-
 1336 man Pressures Challenge Multireservoir Management for Flood Protection,
 1337 Hydropower Production, and Agricultural Water Supply. *Water Resources*
 1338 *Research*, *54*(7), 4638–4662. doi: 10.1029/2018WR022743
- 1339 Quinn, J. D., Reed, P. M., & Keller, K. (2017). Direct policy search for ro-
 1340 bust multi-objective management of deeply uncertain socio-ecological tip-
 1341 ping points. *Environmental Modelling and Software*, *92*, 125–141. doi:
 1342 10.1016/j.envsoft.2017.02.017
- 1343 Rampini, A. A., Sufi, A., & Viswanathan, S. (2014). Dynamic risk management.
 1344 *Journal of Financial Economics*, *111*(2), 271–296. doi: 10.1016/j.jfineco.2013
 1345 .10.003
- 1346 Reed, P. M., Hadka, D., Herman, J. D., Kasprzyk, J. R., & Kollat, J. B. (2013).
 1347 Evolutionary multiobjective optimization in water resources: The past,

- 1348 present, and future. *Advances in Water Resources*, 51, 438–456. doi:
 1349 10.1016/j.advwatres.2012.01.005
- 1350 Rehan, R., Knight, M. A., Unger, A. J., & Haas, C. T. (2013). Development of
 1351 a system dynamics model for financially sustainable management of mu-
 1352 nicipal watermain networks. *Water Research*, 47(20), 7184–7205. doi:
 1353 10.1016/j.watres.2013.09.061
- 1354 Rehan, R., Unger, A., Knight, M. A., & Haas, C. (2015). Strategic water utility
 1355 management and financial planning using a new system dynamics tool. *Journal*
 1356 *- American Water Works Association*, 107(1), E22-E36. doi: 10.5942/jawwa
 1357 .2015.107.0006
- 1358 Rosenstein, M. T., & Barto, A. G. (2001). Robot weightlifting by direct policy
 1359 search. In *International joint conference on artificial intelligence* (pp. 839–
 1360 844).
- 1361 Roszkowska, E. (2011). Multi-criteria decision making models by applying the TOP-
 1362 SIS method to crisp and interval data. In *Multiple criteria decision making*
 1363 *'10-11* (pp. 200–230). Publisher of The University of Economics in Katowice.
- 1364 Rundo, F., Trenta, F., di Stallo, A. L., & Battiato, S. (2019, 12). Machine learning
 1365 for quantitative finance applications: A survey. *Applied Sciences (Switzerland)*,
 1366 9(24). doi: 10.3390/app9245574
- 1367 Salas-Molina, F., Pla-Santamaria, D., & Rodriguez-Aguilar, J. A. (2018). A multi-
 1368 objective approach to the cash management problem. *Annals of Operations*
 1369 *Research*, 267(1-2), 515–529. doi: 10.1007/s10479-016-2359-1
- 1370 Saltelli, A., Tarantola, S., & Campolongo, F. (2000). Sensitivity Anaysis as an Ingre-
 1371 dient of Modeling. *Statistical Science*, 15(4), 377–395.
- 1372 San Francisco Public Utilities Commission. (2016). *Comprehensive Annual Financial*
 1373 *Report, Fiscal Years Ended June 30, 2016 and 2015* (Tech. Rep.).
- 1374 Savikhin, A., Lam, H. C., Fisher, B., & Ebert, D. S. (2011). An experimental study
 1375 of financial portfolio selection with visual analytics for decision support. *Pro-*
 1376 *ceedings of the Annual Hawaii International Conference on System Sciences*,
 1377 1–10. doi: 10.1109/HICSS.2011.54
- 1378 Scanlon, B. R., Ruddell, B. L., Reed, P. M., Hook, R. I., Zheng, C., Tidwell,
 1379 V. C., & Siebert, S. (2017). The food-energy-water nexus: Transform-
 1380 ing science for society. *Water Resources Research*, 53(5), 3550–3556. doi:

- 1381 10.1002/2017WR020889
- 1382 Shannon, C. (1948). A Mathematical Theory of Communication. *The Bell System*
1383 *Technical Journal*, XXVII(3), 379–423.
- 1384 Singh, L. A., Whittecar, W. R., Diprinzio, M. D., Herman, J. D., Ferringer, M. P.,
1385 & Reed, P. M. (2020). Low cost satellite constellations for nearly con-
1386 tinuous global coverage. *Nature Communications*, 11(200), 1–7. doi:
1387 10.1038/s41467-019-13865-0
- 1388 Smith, R., Kasprzyk, J., & Dilling, L. (2017). Participatory Framework for As-
1389 sessment and Improvement of Tools (ParFAIT): Increasing the impact and
1390 relevance of water management decision support research. *Environmental*
1391 *Modelling and Software*, 95, 432–446. doi: 10.1016/j.envsoft.2017.05.004
- 1392 Smith, R., Kasprzyk, J., & Dilling, L. (2019). Testing the potential of Multiobjec-
1393 tive Evolutionary Algorithms (MOEAs) with Colorado water managers. *Envi-*
1394 *ronmental Modelling and Software*, 117, 149–163. doi: 10.1016/j.envsoft.2019
1395 .03.011
- 1396 Sodhi, M. S. (2005). LP modeling for asset-liability management: A survey of
1397 choices and simplifications. *Operations Research*, 53(2), 181–196. doi:
1398 10.1287/opre.1040.0185
- 1399 Spronk, J., Steuer, R. E., & Zopounidis, C. (2005). Multicriteria Decision
1400 Aid/Analysis in Finance. In *Multiple criteria decision analysis: State of the art*
1401 *surveys. international series in operations research & management science, vol.*
1402 *78* (pp. 799–848). New York, NY: Springer.
- 1403 Stakhiv, E. Z. (2011). Pragmatic approaches for water management under climate
1404 change uncertainty. *Journal of the American Water Resources Association*.
1405 doi: 10.1111/j.1752-1688.2011.00589.x
- 1406 Su, Y., Kern, J. D., Denaro, S., Hill, J., Reed, P., Sun, Y., . . . Characklis, G. W.
1407 (2020). An open source model for quantifying risks in bulk electric power
1408 systems from spatially and temporally correlated hydrometeorological pro-
1409 cesses. *Environmental Modelling & Software*, 126(January), 104667. Re-
1410 trieved from <https://doi.org/10.1016/j.envsoft.2020.104667> doi:
1411 10.1016/j.envsoft.2020.104667
- 1412 Su, Y., Kern, J. D., Reed, P. M., & Characklis, G. W. (2020). Compound hydrom-
1413 eteorological extremes across multiple timescales drive volatility in California

- 1414 electricity market prices and emissions. *Applied Energy*, 276(July), 115541.
 1415 doi: 10.1016/j.apenergy.2020.115541
- 1416 Sun, J., Fang, W., Wu, X., Lai, C. H., & Xu, W. (2011). Solving the multi-stage
 1417 portfolio optimization problem with a novel particle swarm optimization. *Ex-
 1418 pert Systems with Applications*, 38(6), 6727–6735. doi: 10.1016/j.eswa.2010.11
 1419 .061
- 1420 Tapia, M. G. C., & Coello Coello, C. A. (2007). Applications of multi-objective evo-
 1421 lutionary algorithms in economics and finance: A survey. *2007 IEEE Congress
 1422 on Evolutionary Computation*, 532–539. doi: 10.1109/CEC.2007.4424516
- 1423 Tejada-Guibert, J. A., Johnson, S. A., & Stedinger, J. R. (1995). The Value of
 1424 Hydrologic Information in Stochastic Dynamic Programming Models of a
 1425 Multireservoir System. *Water Resources Research*, 31(10), 2571–2579. doi:
 1426 10.1029/95WR02172
- 1427 Turvey, C. G. (2001). Weather Derivatives for Specific Event Risks in Agricul-
 1428 ture. *Review of Agricultural Economics*, 23(2), 333–351. doi: 10.1111/1467
 1429 -9353.00065
- 1430 Turvey, C. G., & Nayak, G. (2003). The semivariance-minimizing hedge ra-
 1431 tio. *Journal of Agricultural and Resource Economics*, 28(1), 100–115. doi:
 1432 10.2307/40987175
- 1433 Voisin, N., Kintner-Meyer, M., Wu, D., Skaggs, R., Fu, T., Zhou, T., . . . Kraucu-
 1434 nas, I. (2018). Opportunities for joint water-energy management: Sensi-
 1435 tivity of the 2010 western U.S. electricity grid operations to climate oscilla-
 1436 tions. *Bulletin of the American Meteorological Society*, 99(2), 299–312. doi:
 1437 10.1175/BAMS-D-16-0253.1
- 1438 Wild, T. B., Reed, P. M., Loucks, D. P., Mallen-Cooper, M., & Jensen, E. D. (2019).
 1439 Balancing Hydropower Development and Ecological Impacts in the Mekong:
 1440 Tradeoffs for Sambor Mega Dam. *Journal of Water Resources Planning and
 1441 Management*, 145(2), 05018019. doi: 10.1061/(asce)wr.1943-5452.0001036
- 1442 Woodruff, M. J., Reed, P. M., & Simpson, T. W. (2013). Many objective visual an-
 1443 alytics: Rethinking the design of complex engineered systems. *Structural and
 1444 Multidisciplinary Optimization*, 48, 201–219. doi: 10.1007/s00158-013-0891-z
- 1445 Wu, W., Maier, H. R., Dandy, G. C., Leonard, R., Bellette, K., Cuddy, S., & Ma-
 1446 heepala, S. (2016). Including stakeholder input in formulating and solving real-

- 1447 world optimisation problems: Generic framework and case study. *Environmental Modelling and Software*, 79, 197–213. doi: 10.1016/j.envsoft.2016.02.012
- 1448
- 1449 Yeh, W. G. (1985). Reservoir Management and Operations Models: A State-of-
1450 the-Art Review. *Water Resources Research*, 21(12), 1797–1818. doi: 10.1029/
1451 WR021i012p01797
- 1452 Zatarain Salazar, J., Reed, P. M., Herman, J. D., Giuliani, M., & Castelletti, A.
1453 (2016). A diagnostic assessment of evolutionary algorithms for multi-objective
1454 surface water reservoir control. *Advances in Water Resources*, 92, 172–185.
1455 doi: 10.1016/j.advwatres.2016.04.006
- 1456 Zatarain Salazar, J., Reed, P. M., Quinn, J. D., Giuliani, M., & Castelletti, A.
1457 (2017). Balancing exploration, uncertainty and computational demands in
1458 many objective reservoir optimization. *Advances in Water Resources*, 109,
1459 196–210. doi: 10.1016/j.advwatres.2017.09.014
- 1460 Zeff, H. B., & Characklis, G. W. (2013). Managing water utility financial risks
1461 through third-party index insurance contracts. *Water Resources Research*,
1462 49(8), 4939–4951. doi: 10.1002/wrcr.20364
- 1463 Zeff, H. B., Herman, J. D., Reed, P. M., & Characklis, G. W. (2016). Cooperative
1464 drought adaptation: Integrating infrastructure development, conservation, and
1465 water transfers into adaptive policy pathways. *Water Resources Research*,
1466 52(9), 7327–7346. doi: 10.1002/2016WR018771
- 1467 Zeff, H. B., Kasprzyk, J. R., Herman, J. D., Reed, P. M., & Characklis, G. W.
1468 (2014). Navigating financial and supply reliability tradeoffs in regional drought
1469 management portfolios. *Water Resources Research*, 50(6), 4906–4923. doi:
1470 10.1002/2013WR015126
- 1471 Zopounidis, C., Doumpos, M., & Niklis, D. (2018). Financial decision support: an
1472 overview of developments and recent trends. *EURO Journal on Decision Pro-
1473 cesses*, 6(1-2), 63–76. doi: 10.1007/s40070-018-0078-3
- 1474 Zopounidis, C., Galariotis, E., Doumpos, M., Sarri, S., & Andriosopoulos, K. (2015).
1475 Multiple criteria decision aiding for finance: An updated bibliographic sur-
1476 vey. *European Journal of Operational Research*, 247(2), 339–348. doi:
1477 10.1016/j.ejor.2015.05.032