

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<sup>1,2</sup>, Gregory W. Characklis<sup>1,2</sup>, and Patrick M. Reed<sup>3</sup>**

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

7                   <sup>2</sup>Center 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                   <sup>3</sup>Department 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](mailto: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 financial risk management policies in the context of hydropower production in a snow-dominated region. 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.

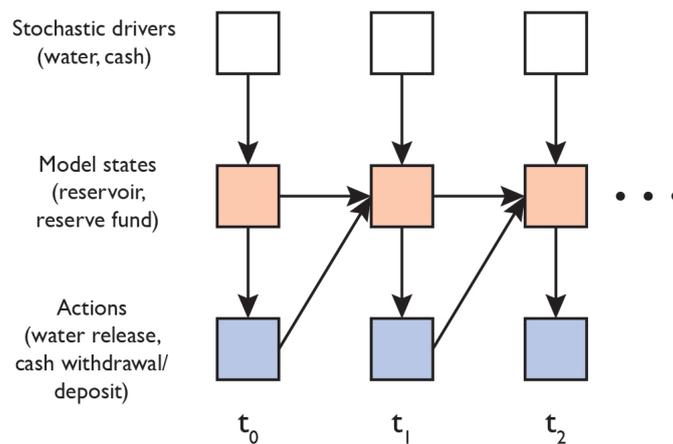
**Keywords**

Hydropower, water resources, snow, financial risk, decision support, uncertainty

**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 a reserve fund) that is filled in times of abundance and drawn down in times of scarcity

(Figure 1). Other risk management tools may also be used to limit the impact of low-flow periods, but at a cost (e.g., water desalination or demand management for stream-flow deficits, and borrowing or financial hedging for cash flow deficits). In both cases, the manager must make decisions under an array of uncertainties, and may need to navigate tradeoffs between conflicting objectives (e.g., flood control vs. water supply for reservoir control, risk vs. cost for financial risk management). And in both cases, as systems dynamically evolve, managers will have to adapt to new information as it becomes available. In other words, reservoir control and financial risk management can be formulated as very similar Markov Decision Processes (MDPs) (Bertsekas, 2019; Powell, 2019), whether managers attempt to solve this problem explicitly, using programmatic approaches such as stochastic dynamic programming, or implicitly, relying on expert specified rules. Additionally, reservoir control and financial risk management are strongly interdependent activities for water-reliant organizations in the Food-Energy-Water Nexus, such as hydropower producers, municipal water utilities, and irrigation districts (Cai, Wallington, Shafiee-Jood, & Marston, 2018; D’Odorico et al., 2018; Scanlon et al., 2017). Such organizations rely on water for the provision of services, and as a result, their revenues and/or costs can be highly dependent on hydrologic inflows (Blomfield & Plummer, 2014; Larson, Freedman, Passinsky, Grubb, & Adriaens, 2012). This suggests that an understanding of complex water resource system dynamics can be used to better characterize and 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.

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

81 A variety of approximation methods have been developed to overcome these chal-  
82 lenges, such as approximate dynamic programming, reinforcement learning, and model  
83 predictive control (Bertsekas, 2019). Direct Policy Search (DPS) (Rosenstein & Barto,  
84 2001), or parameterization-simulation-optimization (Koutsoyiannis & Economou, 2003),  
85 has become increasingly popular in the field of water resources systems analysis (Macian-  
86 Sorribes & Pulido-Velazquez, 2019). DPS is an approximation in policy space (Powell,  
87 2019), wherein the optimal operating policy is assumed to lie in the space of a certain  
88 parametric family of functions, and the policy parameters are optimized rather than the  
89 decisions themselves (i.e., optimizing state-aware adaptive rule systems instead of spe-  
90 cific actions). This drastically reduces the “curse of dimensionality” that limits the tractabil-  
91 ity of large SDP problems. Additionally, DPS allows for “model-free” representation of  
92 stochastic inputs, meaning that observational data, synthetically generated data, and  
93 process-based simulation model output can all be used in lieu of explicit probability dis-  
94 tributions (Desreumaux, Côté, & Leconte, 2018; Giuliani, Quinn, Herman, Castelletti,  
95 & Reed, 2018). A simulation-based approach to optimization also allows for flexible con-  
96 struction of mixed multi-objective formulations (Giuliani et al., 2016; Kasprzyk, Reed,  
97 & Hadka, 2016; Quinn, Reed, & Keller, 2017). In Evolutionary Multi-Objective Direct  
98 Policy Search (EMODPS) (Giuliani, Herman, Castelletti, & Reed, 2014), the policies are  
99 parameterized with a non-linear approximating network and optimized using a multi-  
100 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 DMs 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 DMs to better understand the complex tradeoffs in their system and choose the solution that best suits their needs (Herman, Zeff, Reed, & Characklis, 2014; Huskova, Matrosova, Harou, Kasprzyk, & Lambert, 2016; Kollat & Reed, 2007). These tools can also help DMs to refine their conceptualization of the problem at hand through iterative problem 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 DMs 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 (Basdekas, 2014; Brown et al., 2015), a crucial element in the production of actionable knowledge (Caniglia et al., 2020; Cash et al., 2003).

Water-reliant organizations such as water utilities and hydropower producers rely on water for the provision of services, so that water or power sales may be significantly diminished during a drought (Hughes et al., 2014; Larson et al., 2012). This can result in severe cash flow deficits that leave an organization at risk of defaulting on its obligations (e.g., debt service, operations and maintenance) (Ceres, 2017; Hughes & Leurig, 2013; Leurig, 2010). Water utilities and hydropower-reliant power utilities are therefore vulnerable to significant financial disruption during drought, and hydrologic financial risk

134 can have an outsized impact on the long-term viability of the utility; indeed, credit rat-  
135 ing agencies have noted that the ability to manage the financial impacts of drought is  
136 an important factor in determining a utility's creditworthiness (Chapman & Breeding,  
137 2014; Moody's Investors Service, 2011, 2019). Financial risk management tools, such as  
138 reserve funds, financial hedging contracts, and lines of credit, can reduce the variabil-  
139 ity in net cash flows. This, in turn, can reduce an organization's likelihood of bankruptcy,  
140 improve its credit rating, and reduce its future borrowing costs (Bank & Wiesner, 2010;  
141 Pérez-González & Yun, 2013), in addition to helping risk-averse staff feel more comfort-  
142 able (Bodnar, Giambona, Graham, & Harvey, 2019; Krause & Tse, 2016).

143 Despite the critical role of financial risk management in water resources, decision  
144 support for practitioners in this area has remained limited. There is a long history of con-  
145 sidering financial objectives such as expected revenues and costs in water resources sys-  
146 tems analysis (e.g., see references in Labadie (2004); Macian-Sorribes and Pulido-Velazquez  
147 (2019); Yeh (1985)). However, fewer studies have explicitly accounted for variability in  
148 costs and revenues, or the financial risk management actions that an organization can  
149 take to combat this variability. Those that do have tended to propose static, non-adaptive  
150 management strategies. For example, the literature on using hydrology-based index con-  
151 tracts to hedge exposure to drought has generally assumed that the same hedging con-  
152 tract is purchased each year, not allowing for risk management to be adjusted over time  
153 as conditions change; applications include hydropower (Foster, Kern, & Characklis, 2015;  
154 Hamilton, Characklis, & Reed, 2020; Meyer, Characklis, Brown, & Moody, 2016), wa-  
155 ter supply (Brown & Carriquiry, 2007; Maestro, Barnett, Coble, Garrido, & Bielza, 2016;  
156 Zeff & Characklis, 2013), and agriculture (Denaro, Castelletti, Giuliani, & Characklis,  
157 2020; Mortensen & Block, 2018; Turvey, 2001). Modeling of financial reserves is not com-  
158 mon in the water resources literature, and the limited examples tend to assume that the  
159 utility will contribute either a fixed amount or a fixed fraction of revenues to the reserve  
160 fund each year (Rehan, Knight, Unger, & Haas, 2013; Rehan, Unger, Knight, & Haas,  
161 2015; Zeff, Kasprzyk, Herman, Reed, & Characklis, 2014).

162 However, financial researchers have demonstrated that adaptive, state-aware ac-  
163 tion is crucial to financial risk management (Bolton, Chen, & Wang, 2011; Disatnik, Duchin,  
164 & Schmidt, 2014; Froot, Scharfstein, & Stein, 1993; Rampini, Sufi, & Viswanathan, 2014).  
165 Just as a reservoir operator should consider current reservoir levels and expected future  
166 inflows when making release decisions, so should a financial risk manager consider the

167 utility's current bank account balance and projected future revenues and costs when de-  
168 ciding whether to withdraw money from the bank, or whether to hedge its drought ex-  
169 posure using index contracts. A variety of optimization methods have been applied to  
170 financial problems such as investment portfolio selection (Markowitz, 1952; Mulvey, 2001;  
171 Pardalos, Sandström, & Zopounidis, 1994), asset-liability management (Kouwenberg &  
172 Zenios, 2008; Sodhi, 2005), and cash flow management (Baumol, 1952; da Costa Moraes,  
173 Nagano, & Sobreiro, 2015; Miller & Orr, 1966). As in water resources systems analysis,  
174 some researchers in finance have moved towards multi-objective formulations (Salas-Molina,  
175 Pla-Santamaria, & Rodriguez-Aguilar, 2018; Spronk, Steuer, & Zopounidis, 2005; Zo-  
176 pounidis, Galariotis, Doumpos, Sarri, & Andriosopoulos, 2015), model-free information  
177 (Sun, Fang, Wu, Lai, & Xu, 2011), heuristic solution methods (da Costa Moraes & Nagano,  
178 2013; Ponsich, Jaimes, & Coello Coello, 2013; Tapia & Coello Coello, 2007), and visual  
179 analytics (Flood, Lemieux, Varga, & William Wong, 2016; Savikhin, Lam, Fisher, & Ebert,  
180 2011), in order to provide more meaningful decision support for practitioners in the fi-  
181 nancial sector (Steuer, Qi, & Hirschberger, 2007; Zopounidis, Doumpos, & Niklis, 2018).  
182 However, decision support remains limited for actors outside of the financial sector, such  
183 as water and power utilities, who nevertheless face significant financial risks.

184 This paper bridges the gap between reservoir control and financial risk manage-  
185 ment to show how computational tools developed for the former can be adapted to the  
186 latter in coupled water resource systems. This research builds on prior work by the au-  
187 thors dealing with drought-related financial risk management by a hydropower producer  
188 in a snow-dominated region. First, Hamilton et al. (2020) developed a hydro-financial  
189 simulation model that abstracts the hydroclimatology, hydropower generation, cash flows,  
190 and financial risk management of the Power Enterprise of the San Francisco Public Util-  
191 ities Commission (SFPUC). The authors used this model to evaluate different static fi-  
192 nancial risk management portfolios within a Monte Carlo framework and search for op-  
193 timal portfolios using an MOEA. In related work, Gupta, Hamilton, Reed, and Charack-  
194 lis (2020) introduced an adaptive EMODPS formulation of a simplified financial risk man-  
195 agement problem, which was used to diagnostically benchmark if modern MOEAs are  
196 capable of addressing this new class of problem. The present study builds on these prior  
197 works by contributing the most detailed and actionable representation to date of how  
198 EMODPS can be used to craft operating policies that adapt to changing conditions over  
199 time when managing drought-related financial risk. The advantages of dynamic decision-

200 making are demonstrated relative to a simplified static operating policy akin to those  
201 commonly applied to financial risk management in the water resources literature. This  
202 paper also demonstrates the value of higher-dimensional problem framings that explic-  
203 itly account for DM preferences with respect to the use of different management tools.  
204 Lastly, a framework is contributed for combining *a posteriori* visual analytics with in-  
205 formation theoretic sensitivity analysis (ITSA) in order to help DMs better understand  
206 how complex, non-linear operating policies achieve their goals by adapting to real-time  
207 information when making decisions.

## 208 **2 Study context**

### 209 **2.1 Study area**

210 San Francisco Public Utilities Commission (SFPUC) owns and operates three reser-  
211 voirs (Hetch Hetchy Reservoir, Cherry Lake, and Lake Eleanor) in the upper Tuolumne  
212 River basin in the Sierra Nevada mountains (Figure S1 in Supporting Information (SI)).  
213 These reservoirs deliver drinking water to much of the San Francisco Bay area, and en  
214 route, the water also provides hydroelectric power. SFPUC uses this hydropower to sell  
215 retail electricity at fixed rates to San Francisco International Airport, municipal build-  
216 ings in San Francisco, and a number of other retail customer classes within the Bay area.  
217 Irrigation districts along the Tuolumne River also have the right to buy surplus hydropower,  
218 when available, at a fixed rate. When hydropower production is in excess of retail and  
219 irrigation district demands, it is sold at floating market rates into the Western Systems  
220 Power Pool (hereafter “wholesale market”). On the other hand, when hydropower is in-  
221 sufficient to meet the demand from retail customers, SFPUC is obligated to purchase  
222 the remainder on the wholesale market (San Francisco Public Utilities Commission, 2016).

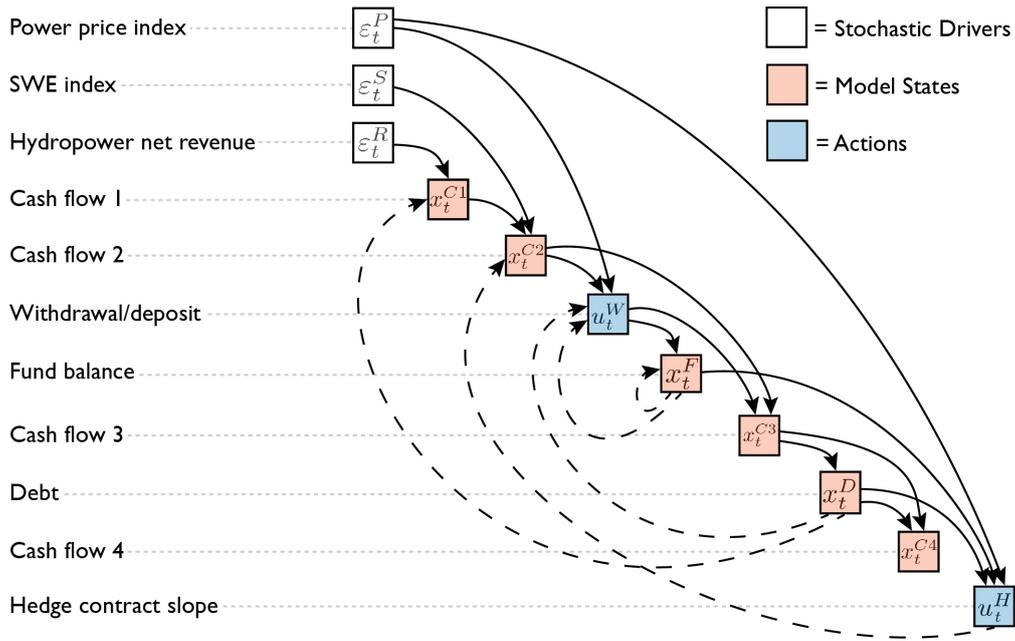
### 223 **2.2 Hydro-financial simulation model**

224 The hydro-financial simulation model from Hamilton et al. (2020) consists of three  
225 types of variables which are updated at an annual time step: stochastic drivers, actions,  
226 and model states (Figure 2). The stochastic drivers are derived from a million-year syn-  
227 thetic dataset that links monthly snow water equivalent depth (SWE), hydropower gen-  
228 eration, wholesale power price, and hydropower net revenue, and is found to closely match

229 the historical record, while providing a wider sampling of possible outcomes than can  
 230 be found in the limited historical data (Hamilton et al., 2020).

231 Three annual quantities are derived from this monthly synthetic dataset and used  
 232 as stochastic drivers for the present study. Firstly, the SWE index ( $\varepsilon^S$ , in inches) is a  
 233 weighted average of February and April SWE observations. The inflows to SFPUC’s reser-  
 234 voirs are dominated by the seasonal dynamics of snow accumulation and melt, so SWE  
 235 measurements taken upstream of the reservoirs in the late winter/early spring can be  
 236 used to predict the magnitude of streamflows when the snow melts in the late spring/early  
 237 summer. Subsequently, the weighted average SWE index is highly correlated with an-  
 238 nual hydropower production and can be used to design index contracts used for finan-  
 239 cial hedging (see below). The second stochastic driver is hydropower net revenue ( $\varepsilon^R$ ,  
 240 in \$M), defined as the total annual cash flow resulting from retail and wholesale hydropower  
 241 sales, minus wholesale power purchases, minus the annual “fixed costs” (debt service pay-  
 242 ments, operations and maintenance, salaries, etc.) that must be paid each year. Lastly,  
 243 the power price index ( $\varepsilon^P$ , in \$/MWh) takes advantage of autocorrelation in the whole-  
 244 sale power market to predict (using linear regression) whether wholesale power prices  
 245 over the coming water year will be favorable or unfavorable for the utility’s net hydropower  
 246 revenues. This index thus provides valuable information for making decisions regarding  
 247 financial risk, and is used as one of the inputs to the dynamic control policies (Section  
 248 3.1.2). For more details on the monthly synthetic records and the calculation of  $\varepsilon^S$  and  
 249  $\varepsilon^R$ , see Hamilton et al. (2020). For more details on the calculation of  $\varepsilon^P$ , see SI Section  
 250 S1.

251 Absent any financial risk management, the utility will experience years in which  
 252 costs outweigh revenues (i.e., net revenue is negative). This situation can be extremely  
 253 disruptive because the utility risks defaulting on its obligations (e.g., debt service or op-  
 254 erations and maintenance). The utility has three financial risk management tools which  
 255 can be used to avoid such negative outcomes. Firstly, They can purchase a snowpack-  
 256 based hedging contract called a capped Contract for Differences (CFD). The CFD (SI  
 257 Figure S2) provides payouts to the utility in low-SWE years (below 24.7 inches), when  
 258 it expects to have low hydropower and thus low revenue, in return for the utility mak-  
 259 ing payments in high-SWE years (above 24.7 inches), when the utility expects to have  
 260 abundant hydropower and surplus revenue. The negative correlation between hydropower  
 261 revenue and CFD payout has been found to significantly reduce the volatility of the com-



**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.

262 bined cash flow, suggesting its value as a financial risk management tool (Hamilton et  
 263 al., 2020). The second risk management tool is a reserve fund, into which the utility can  
 264 deposit surplus cash flows. This allows them to withdraw from the fund when hydropower  
 265 revenues are insufficient to pay their bills. Lastly, the utility has a letter of credit with  
 266 a bank, under which they can borrow money (i.e., issue short-term debt). The debt is  
 267 paid back each year (with interest), and is assumed to take up the slack in situations where  
 268 the other two tools fail to generate sufficient cash flows to avoid defaulting on the util-  
 269 ity’s obligations. Note that the short-term debt considered in this model is distinct from  
 270 longer-term debt service obligations related to past bond offerings, typically associated  
 271 with infrastructure investments, and which are assumed to be part of the “fixed costs”  
 272 above.

273 Figure 2 shows how these financial operations are abstracted in the hydro-financial  
 274 simulation model (See Table 1 for a list of variable names, symbols, units, and constants).  
 275 The sequence of operations occurs at the end of each water year, September 30, based  
 276 on the stochastic outcomes that occur over the course of that water year,  $\epsilon_t$ . Two state-  
 277 aware “actions” each year are governed by the control policy (to be described in Section

278 3.1): the amount of cash withdrawn from/deposited to the reserve fund ( $u_t^W$ , in \$M, where  
 279  $u_t^W > 0$  represents a withdrawal and  $u_t^W < 0$  represents a deposit), and the hedging  
 280 contract slope ( $u_t^H$ , in \$M/inch of SWE). All other variables (“model states”) are au-  
 281 tomatically updated according to the following rules:

$$282 \quad x_t^{C1} = \varepsilon_t^R - r^D x_{t-1}^D \quad (1)$$

$$283 \quad x_t^{C2} = x_t^{C1} + u_{t-1}^H h(\varepsilon_t^S) \quad (2)$$

$$284 \quad x_t^F = r^F x_{t-1}^F - u_t^W \quad (3)$$

$$285 \quad x_t^{C3} = x_t^{C2} + u_t^W \quad (4)$$

$$286 \quad x_t^D = \max(-x_t^{C3}, 0) \quad (5)$$

$$287 \quad x_t^{C4} = x_t^{C3} + x_t^D \quad (6)$$

288 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  
 289 year  $t$ ;  $x_t^D$  and  $x_t^F$  are the short-term debt and reserve fund balance at the end of time  
 290 step  $t$ ;  $r^D$  and  $r^F$  are the annual real interest rates on debt and reserves; and  $h(\varepsilon_t^S)$  is  
 291 the CFD payout function (SI Figure S2). This function converts the stochastic SWE in-  
 292 dex value from the current year into a number of inches of SWE for which the utility will  
 293 receive compensation (if  $h(\varepsilon_t^S) > 0$ ) or owe payment (if  $h(\varepsilon_t^S) < 0$ ). To get the util-  
 294 ity’s total payout received (or payment due), this output is multiplied by the CFD slope,  
 295  $u_{t-1}^H$ , as chosen by the control policy at the end of the previous year (Section 3.1). The  
 296 reader is referred to Hamilton et al. (2020) for more details on construction of the CFD.

297 A full realization of the hydro-financial simulation model requires iterating this se-  
 298 quence for  $T = 20$  years, subject to a randomly sampled  $(T+1)$ -year sequence of stochas-  
 299 tic drivers. The multi-year simulation accounts for the path-dependent dynamics of the  
 300 reserve fund and debt, as well as the autocorrelation within the stochastic power prices.  
 301 The reserve fund and debt are assumed to be zero at  $t = 0$  (although in practice these  
 302 values could be reset based on circumstance). The hedging contract policy in year 0 (the  
 303 slope to be used for the payout in year 1) is calculated using  $x_0^F$ ,  $x_0^D$ , and  $\varepsilon_0^P$ .

### 304 3 Methods

#### 305 3.1 Control formulations

306 Within the hydro-financial simulation model, there are two important decisions that  
 307 must be made each year: the hedging contract slope and the withdrawal from/deposit

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

Variable	Symbol	Value	Units
Power price index	$\varepsilon_t^P$	-	\$/MWh
SWE index	$\varepsilon_t^S$	-	inches
Annual net revenue	$\varepsilon_t^R$	-	\$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	$\epsilon$	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

308 to the reserve fund. A control policy refers to a structured set of rules for making these  
309 two decisions each year. This study introduces two types of control: static (or open-loop)  
310 policies, which perform the same actions with each time step (Section 3.1.1), and dynamic  
311 (or closed-loop) policies, which adapt to changing conditions over time (Section 3.1.2).  
312 Dynamic policies are considered state-aware because the decisions at each time step are  
313 conditioned on the current state of the model. Under both static and dynamic formu-  
314 lations, a policy is defined by a parameter vector which governs its operations. Multi-

315 objective evolutionary optimization (Section 3.3) will be used to search for parameter  
 316 vectors that perform well across four important objectives (Section 3.2).

### 317 **3.1.1 Static policies**

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

$$319 \quad \boldsymbol{\theta}_{stat} = [u^H, x_{max}^F] \quad (7)$$

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

### 333 **3.1.2 Dynamic policies using Direct Policy Search (DPS)**

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

$$337 \quad u_t^{\mathcal{D}} = \mathcal{P}^{\mathcal{D}}(\mathcal{I}_{t'}^{\mathcal{D}} | \boldsymbol{\theta}_{dyn}^{\mathcal{D}}) \quad (8)$$

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

344 In DPS,  $\mathcal{P}$  is assumed to be a family of parametric functions (Rosenstein & Barto,  
 345 2001). This approximation drastically reduces the number of decision variables in the  
 346 search relative to SDP (Bertsekas, 2019; Powell, 2019). Many parametric function fam-  
 347 ilies are available (e.g., piecewise linear, polynomial, artificial neural network), but ra-  
 348 dial basis functions (RBFs) have been shown to be efficient universal approximators for  
 349 DPS (Giuliani, Mason, Castelletti, Pianosi, & Soncini-Sessa, 2014). In this work, a sum  
 350 of RBFs is paired with a constant shift parameter, along with an outer function that per-  
 351 forms operations such as normalization and constraints. Equation 8 can be rewritten as:

$$352 \quad u_t^{\mathcal{D}} = \phi^{\mathcal{D}} \left( a^{\mathcal{D}} + \sum_{m=1}^M w_m^{\mathcal{D}} \varphi_m \left( \mathcal{I}_{t'}^{\mathcal{D}} \right) \right) \quad (9)$$

353 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  
 354 given to the  $m$ th out of  $M$  total RBFs,  $\varphi_m$ . The weights must be chosen such that  $\sum_{m=1}^M w_m^{\mathcal{D}} =$   
 355 1, and  $w_m^{\mathcal{D}} \geq 0$  for all  $m$ . The RBF is defined

$$356 \quad \varphi_m(\mathcal{I}_{t'}^{\mathcal{D}}) = \exp \left( - \sum_{l=1}^L \frac{\left( [\mathcal{I}_{t'}^{\mathcal{D}}]_l - c_{l,m} \right)^2}{(b_{l,m})^2} \right) \quad (10)$$

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

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

$$362 \quad \mathcal{I}_{t'}^W = [r^F \tilde{x}_{t-1}^F, \quad r^D \tilde{x}_{t-1}^D, \quad \tilde{\varepsilon}_t^P, \quad \tilde{x}_t^{C2}] \quad (11)$$

$$363 \quad \mathcal{I}_{t'}^H = [\tilde{x}_t^F, \quad \tilde{x}_t^D, \quad \tilde{\varepsilon}_t^P] \quad (12)$$

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

371 The outer functions  $\phi^W$  and  $\phi^H$  (Equation 9) each consist of multiple nested func-  
 372 tions performing specific operations. The more straightforward  $\phi^H$  consists of a normal-  
 373 ization function,  $\phi^{HN}$ , and a constraint function,  $\phi^{HC}$ . Let  $z_t$  be the argument to  $\phi^H$ ,

374 the action prescribed by the constant shift and sum of radial basis functions in Equa-  
 375 tion 9 when  $H$  is substituted for  $\mathcal{D}$ . This equation can be decomposed as

$$376 \quad u_t^H = \phi^H(z_t) = \phi^{HC}(\phi^{HN}(z_t)) \quad (13)$$

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

$$379 \quad z'_t = \phi^{HN}(z_t) = k^H \max(\min(z_t, 1), 0) \quad (14)$$

380  $\phi^{HC}$  constrains the contract slope to be greater than or equal to a constant thresh-  
 381 old,  $k^H d^H$ , where the threshold parameter  $d^H \in [0, 1]$  is included in the policy param-  
 382 eter vector to be optimized, along with  $a^H$ ,  $\mathbf{w}^H$ ,  $\mathbf{c}$ , and  $\mathbf{b}$ .

$$383 \quad u_t^H = \phi^{HC}(z'_t) = \begin{cases} z'_t, & \text{if } z'_t \geq k^H d^H \\ 0, & \text{otherwise} \end{cases} \quad (15)$$

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

$$387 \quad u_t^W = \phi^W(z_t) = \phi^{WCO}(\phi^{WCI}(\phi^{WW}(\phi^{WN}(z_t)))) \quad (16)$$

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

$$394 \quad z'_t = \phi^{WN}(z_t) = k^R \max(\min(2z_t - 1, 1), -1) \quad (17)$$

395 The withdrawal transformation function,  $\phi^{WW}$ , transforms  $z'_t$  from a cash flow into  
 396 a withdrawal/deposit using the relationship between incoming and outgoing cash flow:

$$397 \quad z''_t = \phi^{WW}(z'_t) = z'_t - x_t^{C2} \quad (18)$$

398 The inner constraint function,  $\phi^{WCI}$ , ensures that the withdrawal/deposit is con-  
 399 sistent with cash-balance equations:

$$400 \quad z_t''' = \phi^{WCI}(z_t'') = \begin{cases} \min(z_t'', r^F x_{t-1}^F), & \text{if } z_t'' \geq 0 \\ \max(z_t'', -\max(x_t^{C2}, 0)), & \text{otherwise} \end{cases} \quad (19)$$

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

405 Lastly, the outer constraint,  $\phi^{WCO}$ , ensures that the reserve fund balance (after  
 406 withdrawal/deposit) cannot be larger than a constant threshold,  $k^F d^W$ , where  $k^F$  (\$M)  
 407 is the normalization constant used for the reserve fund and debt in Table 1, and  $d^W \in$   
 408  $[0, 1]$  is another decision variable to be optimized.

$$409 \quad u_t^W = \phi^{WCO}(z_t''') = \begin{cases} r^F x_{t-1}^F - k^F d^W, & \text{if } (r^F x_{t-1}^F - z_t''') > k^F d^W \\ z_t''', & \text{otherwise} \end{cases} \quad (20)$$

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

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

$$414 \quad \boldsymbol{\theta}_{dyn}^{\mathcal{D}} = [a^{\mathcal{D}}, \quad d^{\mathcal{D}}, \quad \mathbf{w}^{\mathcal{D}}, \quad \mathbf{c}, \quad \mathbf{b}] \quad (21)$$

415 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-  
 416 rameter vector to be optimized,  $\boldsymbol{\theta}_{dyn}$ , is the set of unique parameters,

$$417 \quad \boldsymbol{\theta}_{dyn} = [a^W, \quad a^H, \quad d^W, \quad d^H, \quad \mathbf{w}^W, \quad \mathbf{w}^H, \quad \mathbf{c}, \quad \mathbf{b}] \quad (22)$$

### 418 3.2 Objective formulations

419 This study uses “noisy” objective formulations to account for the uncertainty of  
 420 outcomes under the stochastic drivers. Each candidate policy is assessed using a Monte  
 421 Carlo (MC) ensemble of  $N$  realizations, each representing one possible trajectory of the  
 422 hydro-financial system under a  $T$ -year sample of the stochastic drivers. To convert an  
 423 ensemble of time series into a scalar objective thus requires both a time aggregation step

424 (e.g., the maximum of debt over a  $T$ -year realization) and a noise filtering step (e.g., the  
 425 95th percentile over  $N$  realizations in the ensemble). Four objectives are considered.

426 First is the expected annualized cash flow,  $J^{cash}$ , a measure of “average” cash flows.  
 427 A high value represents a low-cost risk management policy, so  $J^{cash}$  should be maximized.

$$428 \quad J^{cash} \left( x_{t \in (1, \dots, T)}^{C4}, x_T^F, x_T^D \right) = E_{\epsilon} \left[ ANN_t \left( x_{t \in (1, \dots, T)}^{C4}, x_T^F, x_T^D \right) \right] \quad (23)$$

429 where  $x_t^{C4}$  is the final cash flow for year  $t$ ;  $x_T^F$  and  $x_T^D$  are the reserve fund balance and  
 430 debt at the end of the simulation;  $E_{\epsilon}$  is the expectation over the stochastic drivers (ap-  
 431 proximated by the mean of  $N$  MC samples); and  $ANN_t$  is the annualization operator:

$$432 \quad ANN_t \left( x_{t \in (1, \dots, T)}^{C4}, x_T^F, x_T^D \right) = \frac{1}{\sum_{t=1}^T (r^A)^t} \left( \sum_{t=1}^T ((r^A)^t x_t^{C4}) + (r^A)^{T+1} (r^F x_T^F - r^D x_T^D) \right) \quad (24)$$

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

440 The second objective,  $J^{debt}$ , is the 95th percentile of maximum debt. This is a mea-  
 441 sure of the short-term debt load that would be needed to meet fixed costs in an extremely  
 442 bad year (or sequence of years). This is used as a proxy for “risk”, and a DM would want  
 443 to minimize this quantity in order to avoid compromising the utility’s credit rating, in-  
 444 creasing future borrowing costs, and/or risking bankruptcy.

$$445 \quad J^{debt} \left( x_{t \in (1, \dots, T)}^D \right) = Q95_{\epsilon} \left[ \max_{t \in (1, \dots, T)} [x_t^D] \right] \quad (25)$$

446 where the *max* operator takes the maximum debt over a  $T$ -year realization, and the  $Q95$   
 447 operator takes the 95th percentile over the MC ensemble.

448 The third objective,  $J^{hedge}$ , is the expected hedging frequency, to be minimized.

$$449 \quad J^{hedge} \left( u_{t \in (0, \dots, T-1)}^H \right) = E_{\epsilon} \left[ \max_{t \in (0, \dots, T-1)} [\mathbf{1}_{u_t^H > 0}] \right] \quad (26)$$

450 where the indicator function  $\mathbf{1}_{u_t^H > 0}$  returns a 1 if the hedging contract slope is non-zero,  
 451 and a 0 otherwise.  $J^{hedge}$  represents the likelihood that the utility will enter into at least

452 one hedging contract over the course of 20 years. Note that each hedging contract does  
 453 have an annual cost, a “loading” applied by the contract seller that makes the expected  
 454 payout of  $h$  (SI Figure S2) negative (Hamilton et al., 2020). However, this cost is already  
 455 accounted for by  $J^{cash}$ , and does not need to be double-counted.  $J^{hedge}$ , rather, relates  
 456 to the significant extra costs (in time, personnel, and/or money) of having to set up the  
 457 first hedging contract within a realization, assuming that this start-up cost will be sig-  
 458 nificantly diminished in subsequent contract purchases.

459 The last objective is the expected maximum reserve fund balance,  $J^{fund}$ :

$$460 \quad J^{fund} \left( x_{t \in (1, \dots, T)}^F \right) = E_{\epsilon} \left[ \max_{t \in (1, \dots, T)} [x_t^F] \right] \quad (27)$$

461 This objective represents the expected value of the largest reserve fund used in a  $T$ -year  
 462 realization, and should be minimized. This objective would be important if a utility is  
 463 worried that rate-payers or regulators would be critical of large liquid reserves.

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

$$467 \quad E_{\epsilon} [x_T^D - x_{T-1}^D] < \epsilon \quad (28)$$

468 where  $\epsilon$  is a small constant (Table 1). This “noisy” constraint is calculated from the en-  
 469 tire MC ensemble; there is no constraint on debt use in individual extreme realizations.

### 470 **3.3 Multi-objective evolutionary optimization of control policies**

471 As described in Sections 1 and 3.1.2, DPS has a number of advantages relative to  
 472 traditional methods such as SDP, especially when combined with non-linear approximat-  
 473 ing networks such as RBFs. However, RBF parameterization can result in a highly non-  
 474 linear and non-convex search space that is difficult to traverse with gradient-based meth-  
 475 ods, especially when combined with noisy multi-objective formulations (Giuliani & Castel-  
 476 letti, 2016; Giuliani, Mason, et al., 2014; Giuliani et al., 2018). These problems are bet-  
 477 ter handled by MOEAs, which use evolution-inspired strategies (e.g., selection, mating,  
 478 mutation) to iteratively improve a population of solutions competing on multiple objec-  
 479 tives (Coello Coello, Lamont, & Van Veldhuizen, 2007). Population-based methods can  
 480 approximate the entire Pareto set in a single run, rather than rerunning many single-  
 481 objective optimizations, making them quite efficient on many-objective problems. Ad-

ditionally, these heuristic approaches require no information on the topology of a problem and are well-adapted to the types of nonlinear, non-convex, high-dimensional, and stochastic problems that are common in both water resources (Maier et al., 2014; Nicklow et al., 2010; Reed, Hadka, Herman, Kasprzyk, & Kollat, 2013) and finance (Ponsich et al., 2013; Tapia & Coello Coello, 2007).

This study employs the Borg Multiobjective Evolutionary Algorithm (MOEA) (Hadka & Reed, 2013), which has been particularly successful across a range of difficult problems in water resources (Gupta et al., 2020; Hadka & Reed, 2012; Reed et al., 2013; Zatarain Salazar, Reed, Herman, Giuliani, & Castelletti, 2016) and engineering design (Singh et al., 2020; Woodruff, Reed, & Simpson, 2013). The Borg MOEA includes novel components such as adaptive search operator selection, adaptive population sizing, stagnation detection via epsilon-progress, and epsilon-dominance archiving. Its self-adaptive nature makes the Borg MOEA highly controllable (Hadka & Reed, 2013; Reed et al., 2013), and the master-worker parallel variant used in this study is scalable on high-performance computing infrastructure (Giuliani et al., 2018; Zatarain Salazar et al., 2017).

### 3.4 Information theoretic sensitivity analysis

A sensitivity analysis (SA) is an evaluation of the effects of a model’s input factors on its output factors, and a wide range of methods are available to suit different purposes. According to the taxonomy of SA introduced by Pianosi et al. (2016), the method that follows would be considered a quantitative, global, “all-at-a-time” SA, based on simulation model output. This SA is used to explore how different policies adapt their actions to changing conditions; more specifically, it will probe the sensitivity of the prescribed hedging and withdrawal decisions (Equation 8) to changing informational inputs (Equations 11-12). This type of analysis can help to “open the black box” of control policies, helping DMs better understand how different policies respond to changing information (Quinn et al., 2019).

However, commonly-used variance-based methods, which decompose the variance of an output variable into contributions from covariance with different input variables, are inappropriate in the proposed context. First, the policies described by Equations 9-20 are highly non-linear and discontinuous, so that variance and covariance are inappropriate measures of variability and relationship. Secondly, most variance decomposition

513 methods assume independence between the input variables, and can lead to misleading  
 514 results when this independence is violated (Borgonovo, 2007; Borgonovo, Castaings, &  
 515 Tarantola, 2011). This is especially problematic in the current context because most Pareto-  
 516 optimal solutions will impose the following relationship between the reserve fund and debt:  
 517 if one is large, the other is usually zero. For these reasons, moment-independent global  
 518 SA methods, such as entropy-based SA (Auder & Iooss, 2009; Krzykacz-Hausmann, 2001),  
 519 are preferred. Hejazi, Cai, and Ruddell (2008) use ITSA to study the impact of hydro-  
 520 logic information on historical release decisions made by reservoir operators under dif-  
 521 ferent conditions. A similar approach is adopted here to study how different policies along  
 522 the Pareto front use model state information to make decisions.

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

$$528 \quad H(U^{\mathcal{D}}) = - \sum_{u^{\mathcal{D}} \in v^{\mathcal{D}}} p(u^{\mathcal{D}}) \log_2 p(u^{\mathcal{D}}) \quad (29)$$

529 where  $p(u^{\mathcal{D}})$  is the probability mass function (PMF) after discretizing the outcome to  
 530 a discrete domain,  $v^{\mathcal{D}}$ . The entropy (in bits when written with a base-2 logarithm) can  
 531 be thought of as a moment-free measure of uncertainty, or dispersion, in the probabil-  
 532 ity distribution of a random variable. A variable whose outcome is known determinis-  
 533 tically has zero entropy, while a uniformly distributed variable is the most uncertain and  
 534 has the largest possible entropy. Although a continuous variant of entropy based on Kullback-  
 535 Leibler divergence can also be used for SA (Auder & Iooss, 2009; Liu, Chen, & Sudjianto,  
 536 2006; Pappenberger, Beven, Ratto, & Matgen, 2008), the discrete version is more straight-  
 537 forward when the random variable's distribution is unknown.

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

$$540 \quad MI(\mathbf{I}_i^{\mathcal{D}}, U^{\mathcal{D}}) = H(U^{\mathcal{D}}) - H(U^{\mathcal{D}} | \mathbf{I}_i^{\mathcal{D}}) \quad (30)$$

$$541 \quad = - \sum_{\mathcal{I}_i^{\mathcal{D}} \in \iota_i^{\mathcal{D}}} \sum_{u^{\mathcal{D}} \in v^{\mathcal{D}}} p(\mathcal{I}_i^{\mathcal{D}}, u^{\mathcal{D}}) \log_2 \frac{p(\mathcal{I}_i^{\mathcal{D}}, u^{\mathcal{D}})}{p(\mathcal{I}_i^{\mathcal{D}})p(u^{\mathcal{D}})} \quad (31)$$

542 where  $\mathbf{I}_i^{\mathcal{D}}$  is the random variable for the  $i$ th informational input (e.g., reserve fund bal-  
 543 ance or power price index),  $H(U^{\mathcal{D}} | \mathbf{I}_i^{\mathcal{D}})$  is the entropy of the action conditional on the in-

544 put,  $p(\mathcal{I}_i^{\mathcal{D}})$  is the PMF for the input on the discrete domain  $\mathcal{I}_i^{\mathcal{D}}$ , and  $p(\mathcal{I}_i^{\mathcal{D}}, u^{\mathcal{D}})$  is the  
 545 joint PMF on the discrete domain  $\mathcal{I}_i^{\mathcal{D}} \times \mathcal{U}^{\mathcal{D}}$ . This mutual information is a measure how  
 546 much information the outcome of one random variable contains about the outcome of  
 547 the other: how much does knowledge of a particular informational input reduce the un-  
 548 certainty in the prescribed action?

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

$$551 \quad \eta_i^{\mathcal{D}} = \frac{MI(\mathcal{I}_i^{\mathcal{D}}, U^{\mathcal{D}})}{H(U^{\mathcal{D}})} \quad (32)$$

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

## 555 4 Computational experiments

### 556 4.1 Problem formulations

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

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

$$565 \quad \boldsymbol{\theta}^* = \arg \min_{\boldsymbol{\theta}} [-J^{cash}(\boldsymbol{\theta}), \quad J^{debt}(\boldsymbol{\theta})] \quad (33)$$

566 while the four-objective problem can be written:

$$567 \quad \boldsymbol{\theta}^* = \arg \min_{\boldsymbol{\theta}} [-J^{cash}(\boldsymbol{\theta}), \quad J^{debt}(\boldsymbol{\theta}), \quad J^{hedge}(\boldsymbol{\theta}), \quad J^{fund}(\boldsymbol{\theta})] \quad (34)$$

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

## 4.2 MOEA parameters

An ensemble of  $N = 50,000$  realizations is run for each function evaluation, balancing computational demand against the need to minimize sampling error in the noisy objective/constraint evaluations (see discussions in Kasprzyk et al. (2012); Quinn, Reed, Giuliani, and Castelletti (2017); Zatarain Salazar et al. (2017)). In order to select the appropriate number of RBFs, the dynamic 4-objective formulation is repeated with 1, 2, 3, 4, 8, and 12 RBFs. Due to the inherent stochasticity of evolutionary algorithms, each optimization is repeated with 10 different random seeds. Each seed is run for 150,000 function evaluations (candidate policy trials). Final populations are assessed in terms of hypervolume, additive epsilon indicator, and generational distance (SI Figure S3), three common metrics for assessing convergence, consistency, and diversity of multi-objective solution sets (Coello Coello et al., 2007; Hadka & Reed, 2012; Reed et al., 2013). Results are found to be relatively insensitive to the number of RBFs used in the dynamic control policies, but  $M = 2$  RBFs is chosen due to their robust performance across seeds. Next, 20 additional seeds are run for the dynamic 4-objective formulation with  $M = 2$ , and 30 seeds each are also run for the dynamic 2-objective, static 2-objective, and static 4-objective formulations. The best known Pareto approximate set for each formulation is the set of non-dominated solutions from across the 30 seeds. After using the same 50,000-member ensemble of 20-year simulations for all formulations/seeds in the initial optimization, each solution in the final Pareto approximate set for each formulation is rerun on a separate 50,000-member ensemble, for which results are reported. Important parameter values for the optimization can be found in SI Table S1; all other Borg MOEA parameters besides those listed are set to the default values (Hadka & Reed, 2013; Reed et al., 2013).

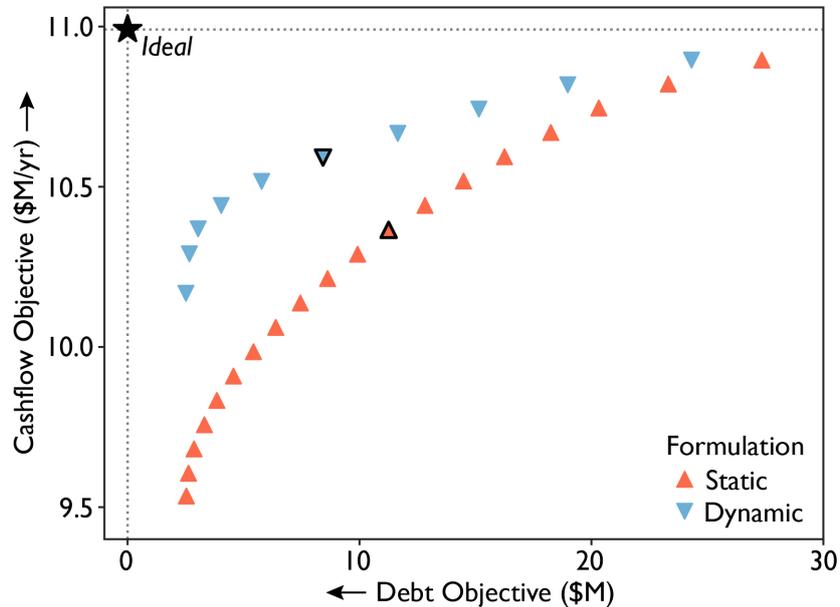
## 4.3 Information theoretic sensitivity analysis parameters

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

## 5 Results and discussion

### 5.1 Static vs. dynamic financial risk management

Figure 3 shows the resulting Pareto approximate sets from the 2-objective optimization problem (Equation 33), under both static and dynamic control formulations. Each point represents a different financial risk management policy. The ideal performance, denoted by a black star, would be achieved with a cash flow objective ( $J^{cash}$ ) of \$10.99M (the average net revenue in the absence of any financial risk management) and a debt objective ( $J^{debt}$ ) of zero. However, this is not possible due to the strong tradeoff between “risk” and “return” that is standard in financial risk applications: in order to achieve higher expected cash flows, the utility must forego costly risk management actions and therefore risk more extreme debt burdens in less favorable realizations. As discussed in Section 3.2, large short-term debt in our model can be viewed as a proxy for larger financial disruptions such as credit rating downgrades or bankruptcy in practice. DMs will have to balance this tradeoff when selecting a particular policy for the utility to use, based on risk aversion, access to credit, and other organizational factors.



**Figure 3.** 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.

617 However, DMs can drastically reduce the risk management tradeoff by using adap-  
618 tive operating rules that respond to changing conditions. The Pareto approximate set  
619 from the dynamic EMODPS control formulation is found to dominate the Pareto approx-  
620 imate set from the static formulation, suggesting that one can improve on both the cash  
621 flow and debt objectives simultaneously. For example, consider the two example poli-  
622 cies outlined in black in Figure 3 and listed in rows 1-2 in Table 2. These are chosen as  
623 the “best compromise” policies near the centers of their respective Pareto approximate  
624 sets (as selected using the TOPSIS method with equal weights on each objective (Be-  
625 hzadian, Khanmohammadi Otaghsara, Yazdani, & Ignatius, 2012; Roszkowska, 2011)).  
626 The dynamic policy is found to reduce  $J^{debt}$  by \$2.83M, or 25.1%, relative to the static  
627 policy. At the same time, it increases  $J^{cash}$  by \$0.23M, representing a 36.1% reduction  
628 in risk management cost relative to the “ideal”  $J^{cash}$  value of \$10.99M. This dual im-  
629 provement highlights the value of dynamic financial risk management: the utility can  
630 improve on both objectives simultaneously without requiring any investment in its in-  
631 frastructure or changes to its physical operations. All that is required is to switch to a  
632 more flexible and efficient financial risk management policy.

**Table 2.** Six example policies referenced in the results, along with their four-objective perfor-  
mance and their information theoretic sensitivity indices related to the hedging action.

Row	Figure	$J^{cash}$ (\$M/yr)	$J^{debt}$ (\$M)	$J^{hedge}$ (unitless)	$J^{fund}$ (\$M)	Fund Sensitivity	Debt Sensitivity	Power Sensitivity
1	3 red	10.37	11.25	1.00	16.11	–	–	–
2	3 blue	10.59	8.42	1.00	19.31	0.74	0.11	0.12
3	7	10.75	15.90	0.77	12.01	0.36	0.72	0.01
4	8a	10.20	3.22	1.00	24.55	0.93	0.12	0.00
5	8b	10.71	15.72	0.40	16.83	0.44	0.96	0.01
6	8c	9.84	8.96	1.00	1.53	0.02	0.03	0.72

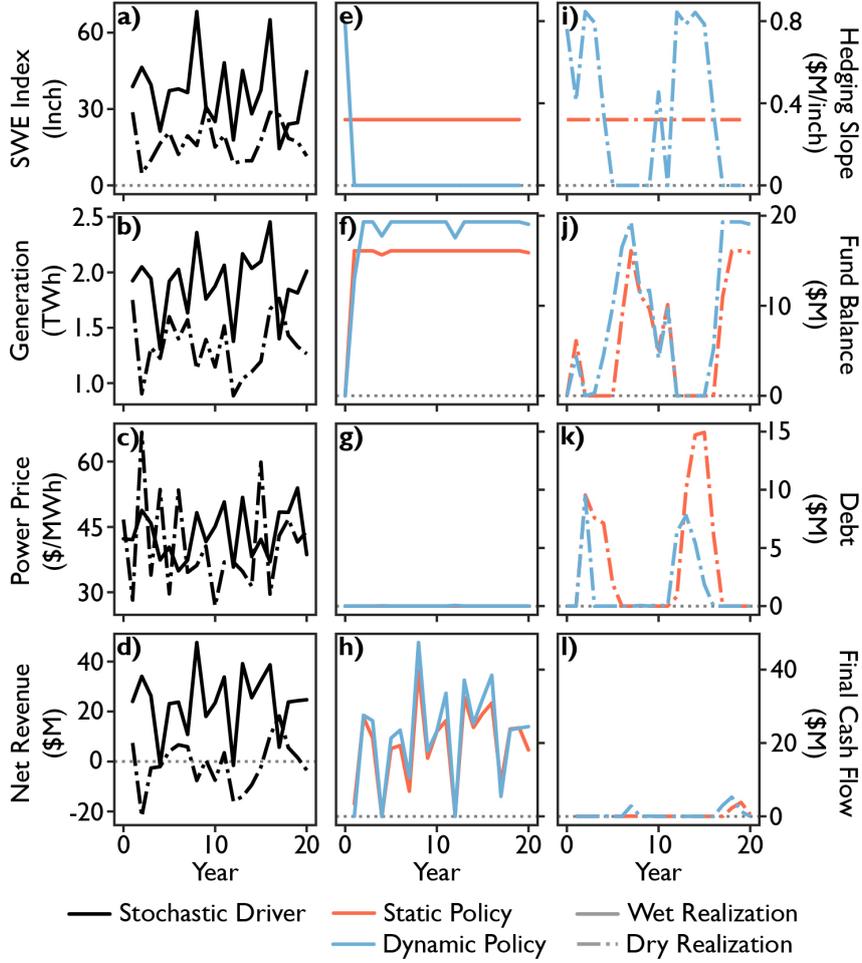
633 The dynamic formulation allows the utility to take different sequences of actions  
634 under different stochastic realizations, using parameterized control rules that allow for  
635 the actions taken at any particular time to be better tailored to the current state of the  
636 system. To elucidate the differences between static and dynamic financial risk manage-

637 ment, the two best compromise policies are simulated under two different 20-year real-  
 638 izations from the synthetic record: an unusually wet period and an unusually dry pe-  
 639 riod (Figure 4). Differences in SWE (4a) lead to drastic differences in hydropower gen-  
 640 eration (4b) and net revenues (4d) under the two realizations, and the dry scenario ex-  
 641periences lengthy periods of drought-related cash flow deficits. The two scenarios also  
 642 yield very different responses in terms of the hedging policy (4e & 4i), reserve fund bal-  
 643 ance (4f & 4j), debt (4g & 4k), and final cash flow (4h & 4l). In the wet scenario, the  
 644 reserve funds fill up quickly and stay nearly full. Neither policy requires any significant  
 645 debt, and final cash flows are generally positive and rather large. In the dry scenario,  
 646 the reserve funds fluctuate up and down, including two periods in which they reach zero.  
 647 During these periods, significant debt is required to overcome further cash flow deficits.  
 648 The final cash flows are close to zero throughout the dry simulation, as both policies strug-  
 649 gle to fill their reserve funds.

650 With respect to the hedging contract, the static policy uses the same contract each  
 651 year in both the wet and dry scenarios, with a payout slope of \$0.32M/inch. The dynamic  
 652 policy, on the other hand, adjusts its contract slope from year to year. In the wet sce-  
 653 nario, it opts not to hedge at all after year 0, while in the dry scenario, it fluctuates be-  
 654 tween \$0 and \$0.85M/inch. Comparing the hedging slope dynamics to the other model  
 655 state variables suggests that this policy opts to hedge only when the reserve fund bal-  
 656 ance is low and/or when debt is non-zero. This strategy allows the dynamic policy to  
 657 achieve higher cash flows than the static policy in wet scenarios (Sub-Figure 4h), by fore-  
 658 going the cost of hedging contracts when the utility already has sufficient protection from  
 659 a large reserve fund. On the other hand, when the reserve is empty and/or there is out-  
 660 standing debt (presumably after a very dry year or sequence of years), the utility pur-  
 661 chases large hedging contracts in order to increase its financial risk coverage and thus  
 662 reduce the risk of extreme debt levels (Sub-Figure 4k). This adaptivity allows the dy-  
 663 namic policy to improve on both the cash flow objective and the debt objective simul-  
 664 taneously, compared to the static policy. As will be seen in Section 5.3, there are a mul-  
 665 tiplicity of ways that utilities can adapt to changing conditions to meet their goals.

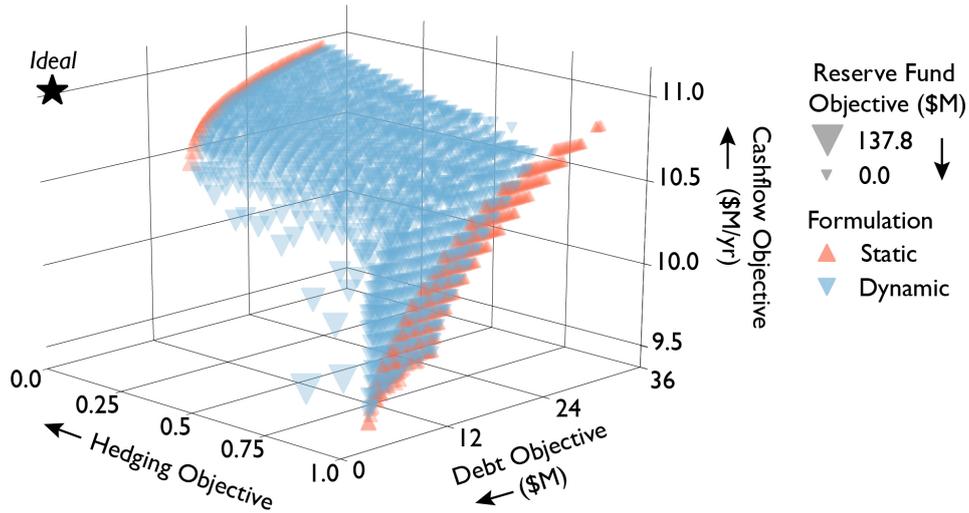
## 666 5.2 Many-objective decision-making

667 As discussed in Section 3.2, a DM choosing a financial risk management policy may  
 668 actually consider other factors beyond risk ( $J^{debt}$ ) and return ( $J^{cash}$ ). For example, the



**Figure 4.** Trajectories for hydro-financial simulation model, over both wet and dry 20-year realizations, for the example static and dynamic policies shown in Figure 3 and rows 1-2 of Table 2. Sub-Figures show the SWE index (a), hydropower generation (b), power price (c), net hydropower revenue (d), hedging slope action (e & i), fund balance (f & j), debt (g & k), and final annual cash flow (h & l). Middle column (e-h) shares its y-axis with the right-hand column (i-l).

669 utility might also worry about the size of the reserve fund needed to enact a particular  
 670 policy ( $J^{fund}$ ), or the likelihood of needing to enact a complicated hedging program ( $J^{hedge}$ ).  
 671 Such DMs are likely to find that none of the solutions found under the 2-objective prob-  
 672 lem (Figure 3) can meet their needs. The 2-objective problem cannot adequately rep-  
 673 resent the tradeoffs that a utility manager must weigh when making these decisions be-  
 674 cause it does not account for DM preferences with respect to the use of different risk man-  
 675 agement tools. For this reason,  $J^{hedge}$  and  $J^{fund}$  can be explicitly included in the op-  
 676 timization using the 4-objective problem (Equation 34).

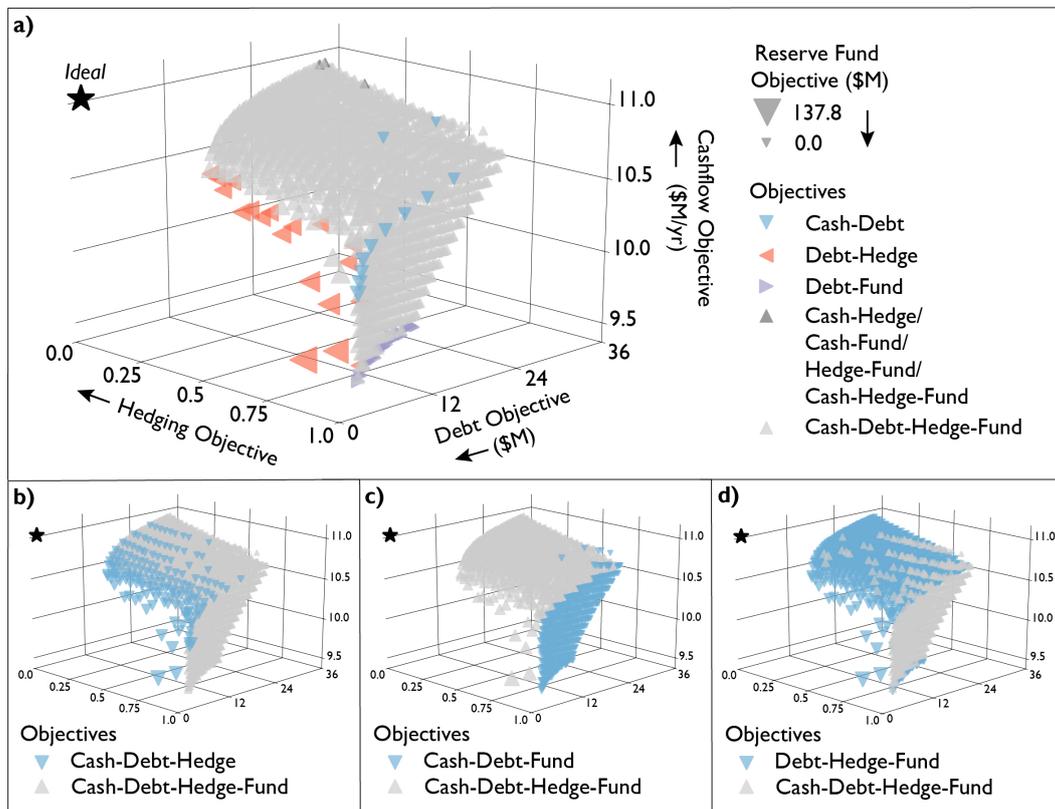


**Figure 5.** Comparison of 4-objective Pareto approximate sets under static and dynamic control formulations.

677 Both the static and dynamic formulations produce much larger Pareto approximate  
 678 sets in this higher-dimensional problem (Figure 5), representing the more complex set  
 679 of tradeoffs across the four objectives. The dynamic Pareto approximate set is found to  
 680 generally outperform the static Pareto approximate set, especially in terms of the over-  
 681 all diversity of solutions. For the static formulation, where the hedging contract slope  
 682 is fixed,  $J^{hedge}$  must be equal to 1 or 0. The dynamic formulation, on the other hand,  
 683 is able to find policies with  $J^{hedge}$  spanning the entire range from 0 to 1. Note that  $J^{hedge}$   
 684 is defined as the fraction of 20-year realizations that contain any hedging, not the frac-  
 685 tion of years which hedge (see Equation 26). Thus, intermediate values between 0 and  
 686 1 represent solutions that are unlikely to hedge in any given year, but maintain the op-  
 687 tion to do so under particularly problematic circumstances. This valuable optionality  
 688 is only possible with a dynamic control strategy. Additionally, the dynamic solution set  
 689 occupies a much larger region within the ridge where  $J^{hedge} = 1$ . These policies out-  
 690 perform the nearest static policies with respect to  $J^{cash}$  and  $J^{debt}$ , but may require the  
 691 use of larger reserve funds. Because the dynamic control method produces a much more  
 692 complete and continuous Pareto approximate set, it allows DMs to find control policies  
 693 that more precisely match their preferences.

694 A major benefit of solving the larger-dimensional problem is that the solution set  
 695 will already contain all of the tradeoffs for all possible lower-dimensional problems (di

696 Pierro, Khu, & Savić, 2007). In the present context, the 4-objective Pareto front will include within it the Pareto fronts for the four 3-objective problems, six 2-objective problems, and four 1-objective problems that are embedded within the 4-objective problem (Figure 6). In Sub-Figure 6a, the blue triangles show the subset of the 4-objective Pareto approximate set that is non-dominated with respect to the original two objectives,  $J^{cash}$  and  $J^{debt}$ . When compared to the original 2-objective solutions (Figure 3), the 4-objective policies are very similar with respect to the first two objectives. However, they can achieve improvements with respect to the two new objectives (see SI Figure 4). In other words, it is possible to improve  $J^{fund}$  and/or  $J^{hedge}$  with no penalty in  $J^{cash}$  or  $J^{debt}$ , but they must be included in the optimization explicitly to realize this benefit.

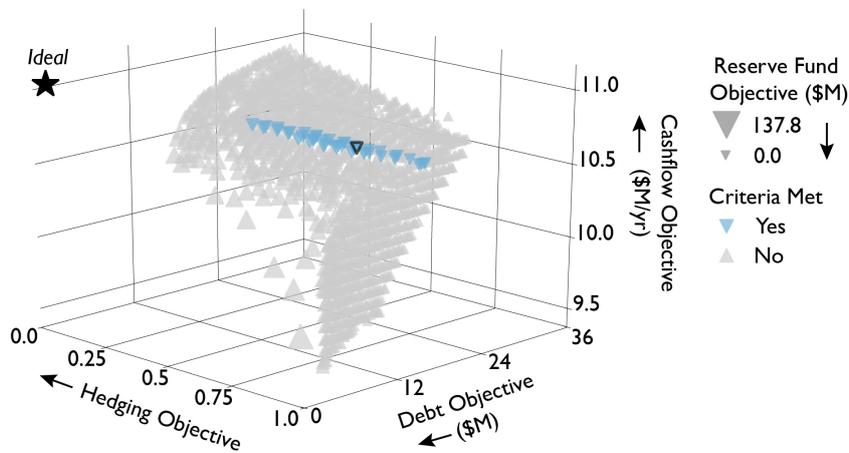


**Figure 6.** Visualization of Pareto approximate sets for different sub-problems. Colored points represent solutions that are non-dominated with respect to a particular sub-problem; for example, orange points in sub-figure (a) represent solutions that are non-dominated with respect to  $J^{debt}$  and  $J^{hedge}$ . Light grey points in all sub-figures represent solutions from the 4-objective problem that are not captured in the lower-dimensional problems.

706 More broadly, the lower-dimensional sub-problems tend to produce Pareto approx-  
 707 imate sets that are near the extreme boundaries of the larger-dimensional problem. Sub-  
 708 Figure 6a includes four sub-problems for which the Pareto approximate set consists of  
 709 a single solution ( $J^{cash\_Jhedge}$ ,  $J^{cash\_Jfund}$ ,  $J^{hedge\_Jfund}$ ,  $J^{cash\_Jhedge\_Jfund}$ ). Each  
 710 of these sub-problems excludes debt, leading to a single optimal policy that performs es-  
 711 sentially no risk management. This is consistent with prior work finding that conflicts  
 712 in higher-dimensional problems can remain hidden in lower-dimensional sub-problems  
 713 (Woodruff et al., 2013). Sub-Figure 6a also shows results for the  $J^{cash\_Jdebt}$ ,  $J^{debt\_Jhedge}$ ,  
 714 and  $J^{debt\_Jfund}$  sub-problems. Each subset of solutions is concentrated along an outer  
 715 border of the larger Pareto front, where performance of the two explicitly-considered ob-  
 716 jectives is optimized at the expense of the other two objectives. The same pattern is ev-  
 717 ident in the 3-objective sub-problems of Sub-Figures 6b ( $J^{cash\_Jdebt\_Jhedge}$ ), 6c ( $J^{cash\_}$   
 718  $J^{debt\_Jfund}$ ), and 6d ( $J^{debt\_Jhedge\_Jfund}$ ). These solution sets are larger, but still oc-  
 719 cupy extremal regions of the overall Pareto front. Thus, by choosing to optimize a 2- or  
 720 3-objective sub-problem, DMs may unwittingly produce an incomplete and biased Pareto  
 721 approximate set.

722 The larger-dimensional problem leads to a fuller set of alternatives that better rep-  
 723 represents the tradeoffs associated with DM preferences for different financial risk manage-  
 724 ment tools. However, it is a non-trivial task to select a single operating policy from among  
 725 the large Pareto approximate set. Interactive visualization approaches can help with this  
 726 task. One example is to allow DMs to apply *a posteriori* performance criteria and “brush  
 727 away” solutions that fail to meet these constraints (Kollat & Reed, 2007; Zeff et al., 2014).  
 728 The strictness of the constraints can be iteratively increased until DMs are relatively ag-  
 729 nostic about the tradeoffs across the feasible solution set. For example, consider a util-  
 730 ity whose financial team (perhaps in consultation with its regulatory commission) de-  
 731 velops the following criteria: if  $\bar{R}$  = \$10.99 million is the mean annual net hydropower  
 732 revenue in the absence of any risk management, then (1) the risk management policy should  
 733 not reduce expected annualized cash flows by more than 2.5% ( $J^{cash} \geq 0.975\bar{R}$ ); (2)  
 734 the utility should rarely be forced to borrow more than 150% of mean net revenue to cover  
 735 cash flow deficits ( $J^{debt} \leq 1.5\bar{R}$ ); and (3) the utility should not maintain reserves larger  
 736 150% of mean net revenue ( $J^{fund} \leq 1.5\bar{R}$ ). These constraints drastically reduce the  
 737 set of feasible solutions (Figure 7). At this point, a quantitative method such as TOP-  
 738 SIS (Behzadian et al., 2012; Roszkowska, 2011) can be used to select one of the remain-

739 ing policies for the utility to use (e.g., the policy outlined in Figure 7 and listed in row  
 740 3 of Table 2).



**Figure 7.** 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

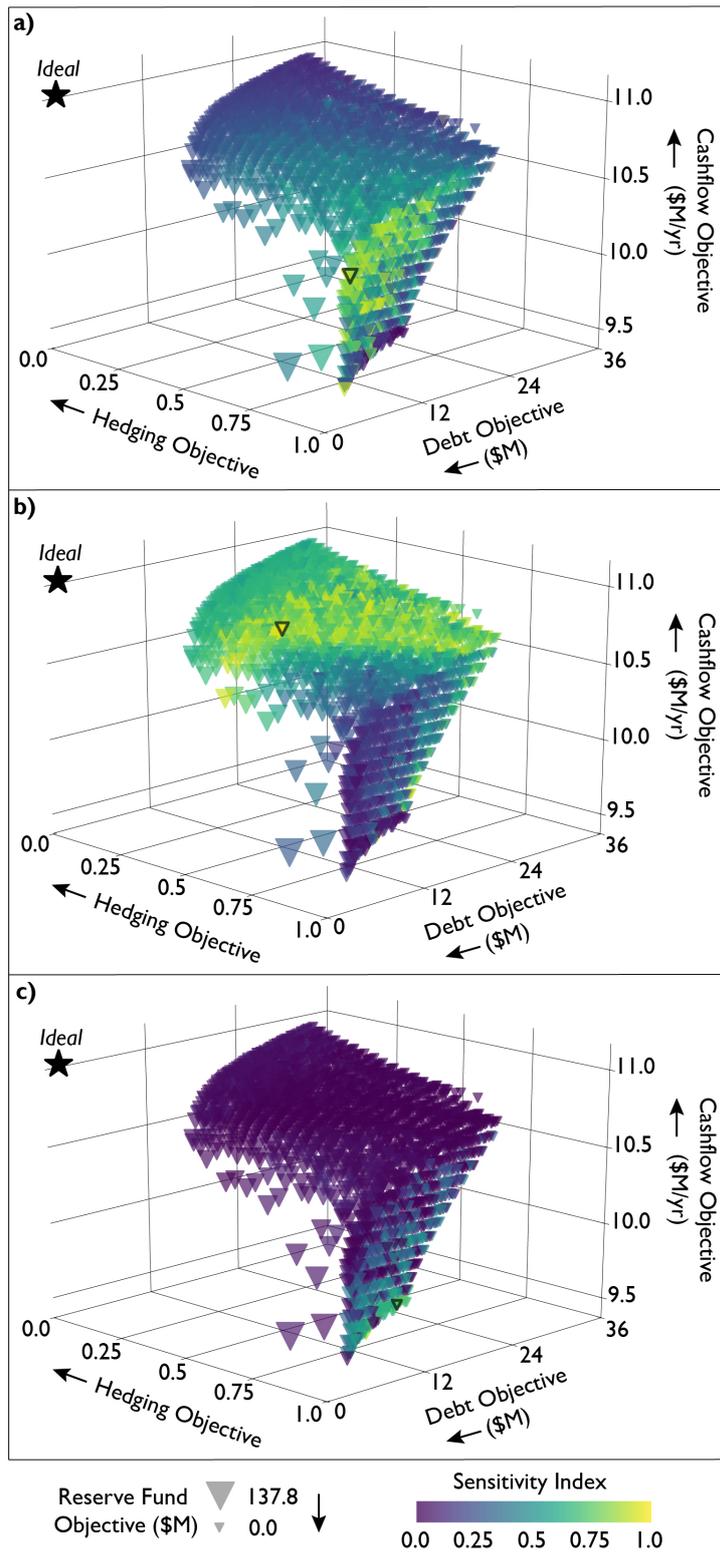
741 While these constraints could, in theory, be applied *a priori* and used to reduce  
 742 the number of objectives in the optimization, it is very difficult in practice for DMs to  
 743 effectively set the constraint values without first understanding the topology of the trade-  
 744 off surface (Kasprzyk et al., 2016; Spronk et al., 2005). This confirms the value of the  
 745 EMODPS approach, which is scalable to extremely large problems on modern high-performance  
 746 computing infrastructure (Giuliani et al., 2018; Zatarain Salazar et al., 2016), suggest-  
 747 ing that the formulation used here could be expanded to include additional objectives  
 748 such as customer rates, social equity, and environmental quality.

### 749 5.3 Value of state information for control

750 As demonstrated above, the EMODPS method can be used to develop control poli-  
 751 cies that perform well across a range of stakeholder preferences. However, DMs may be  
 752 unwilling to adopt a complex, non-linear control policy if its operating rules remain opaque;  
 753 it may be necessary to “open the black box” for users if they are to apply such tools in  
 754 practice (Castelvecchi, 2016; Quinn et al., 2019). Each policy represents a map from a  
 755 vector of inputs (e.g., reserve fund balance) to its outputs (e.g., the hedging contract slope).  
 756 ITSA (Section 3.4) can help DMs to better understand how different policies respond

757 to changing model state information. Figure 8 shows the hedging policy sensitivity indices  
 758 for each solution in the Pareto approximate set, representing the degree to which  
 759 each policy adjusts its annual hedging decision based on each of the three inputs: the  
 760 reserve fund balance ( $\eta_F^H$ , Sub-Figure 8a), the debt ( $\eta_D^H$ , 8b), and the power price index  
 761 ( $\eta_P^H$ , 8c). Each index is a measure of the importance of a particular input variable for  
 762 controlling a state-aware policy;  $\eta = 1$  implies that the policy is entirely controlled by  
 763 the input, while  $\eta = 0$  means that the input has no impact on the policy. Interestingly,  
 764 Figure 8 shows that each input has a different “region of specialization” in objective space.  
 765 The reserve fund balance is the most important input for policies along the top of the  
 766 ridge where  $J^{hedge} = 1$ . These are policies that achieve a relatively low levels of debt  
 767 and high levels of cash flow, in return for frequent hedging and a relatively large reserve  
 768 fund. The debt information, on the other hand, is critical for policies occupying the swath  
 769 of objective space with  $J^{hedge}$  between 0 and 1. The power price index is less informa-  
 770 tive, but it does provide value for policies requiring minimal reserve funds and debt, along  
 771 the bottom edge of the Pareto front.

772 In order to better understand how these policies utilize the information that is avail-  
 773 able to them, it is helpful to visualize the policies themselves. To that end, one high-sensitivity  
 774 policy is chosen for each input (as outlined in Figure 8, and listed in rows 4-6 of Table  
 775 2). Each policy is used to simulate 20 random trajectories of length 20 years, for a to-  
 776 tal of 400 annual decisions. These decisions are visualized in state-action space using parallel-  
 777 coordinate plots (Figure 9). The first three vertical axes represent the three hedging pol-  
 778 icy inputs: reserve fund balance, debt, and power price index. The policy output, the  
 779 hedging contract slope, is represented by the fourth vertical axis as well as the colorbar  
 780 in order to aid interpretation. Each of the colored lines connecting the four axes repre-  
 781 sents one of the 400 simulated decisions. These visualizations, in combination with the  
 782 sensitivity indices, can be useful in understanding how each policy operates. For exam-  
 783 ple, the policy in Sub-Figure 9a appears to hedge selectively, when the reserve fund bal-  
 784 ance has fallen below a certain threshold. Above the threshold, no hedging contract is  
 785 purchased, and below the threshold, the hedging slope increases as the fund balance falls.  
 786 The policy in Sub-Figure 9b has a similar strategy, but structured around debt; hedg-  
 787 ing is zero below some threshold, and increases with debt above the threshold. Lastly,  
 788 the bottom policy always utilizes hedging contracts, the magnitude of which tend to be

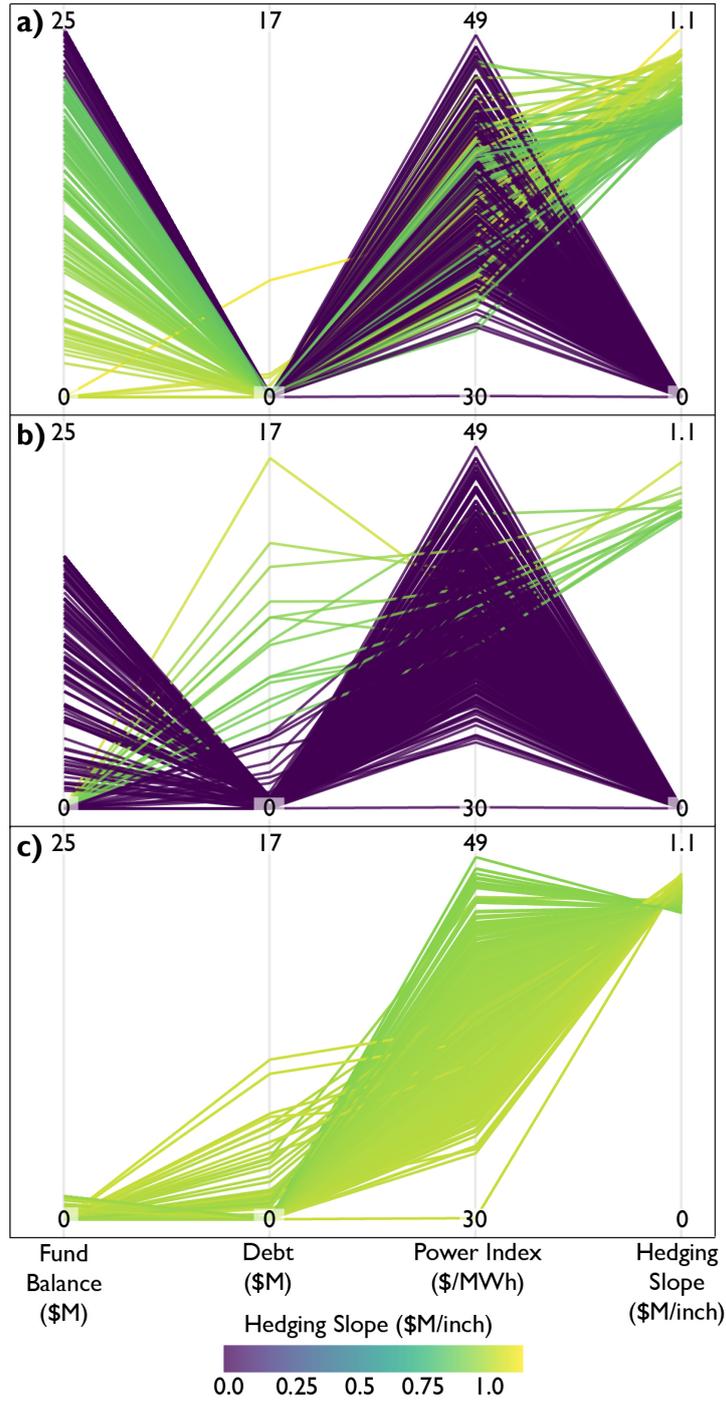


**Figure 8.** Information theoretic sensitivity indices, relative to hedging contract slope decision, for the reserve fund balance (a), debt (b), and power price index (c). One high-sensitivity solution for each input is outlined in black and described in rows 4-6 of Table 2

789 inversely proportional to the power price index. Each of these patterns is consistent with  
790 the sensitivity indices in Figure 8 and Table 2.

791         These plots can be used to build intuition about how the different risk management  
792 policies achieve their competitive advantages. For example, compare the fund-sensitive  
793 policy (a) to the debt-sensitive policy (b). The former maintains a relatively large re-  
794 serve fund for its risk management needs, and uses hedging contracts as a substitute to  
795 maintain its risk protection when the reserve fund is inadequate. This is qualitatively  
796 similar behavior to the example policy simulated in Section 5.1 (Figure 4). The debt-  
797 sensitive policy, on the other hand, keeps a much smaller reserve fund, which results in  
798 more frequent cash flow shortfalls and debt during dry years. In order to reduce the like-  
799 lihood of extreme debt spirals during longer droughts, this policy begins to use hedging  
800 contracts when it has significant debt, and ceases hedging once it has paid off this debt.  
801 The result is that the debt-sensitive policy is significantly more risky than the fund-sensitive  
802 policy, but in return, it is less costly and requires less frequent hedging and a smaller re-  
803 serve fund. The power-sensitive policy (c) takes a more consistent approach, purchas-  
804 ing similar hedging contracts each year. This makes it the most expensive contract of  
805 the three due to the cost of these contracts. However, the risk coverage from hedging al-  
806 lows it to maintain a very small reserve fund and still avoid substantial debt. This pol-  
807 icy does adjust its hedging contract in response to projected wholesale power prices us-  
808 ing the power price index. If the index is high, then the utility expects that its net rev-  
809 enue per unit of hydropower will be higher than average, and vice versa when the index  
810 is low. By purchasing hedging contracts in inverse proportion to this index, the utility  
811 can dampen the overall variability of its combined cash flow (hydropower net revenue  
812 plus the net payout from the hedging contract), and thus reduce its financial risk.

813         ITSA and policy visualization plots for the withdrawal/deposit decision can be found  
814 in SI Figures S5-S6. However, withdrawals and deposits are found to be much less sen-  
815 sitive to model state information than hedging, suggesting that the gains from dynamic  
816 financial risk management in this study largely accrue from dynamic hedging rather than  
817 dynamic reserve fund management. In future problems with a larger number of candi-  
818 date actions, an iterative scheme for selecting the most sensitive decisions to control dy-  
819 namically would be beneficial. One final takeaway from Figures 8 and 9 is that the most  
820 important model states to include in a state-aware control policy can vary widely across  
821 the Pareto approximate set. This implies that the most important input(s) cannot be



**Figure 9.** Hedging control policy visualization for three chosen policies in Figure 8 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.

822 known *a priori* without accounting for DM preferences. This is consistent with both an-  
823 analytical (Graham & Georgakakos, 2010; Tejada-Guibert, Johnson, & Stedinger, 1995)  
824 and empirical (Hejazi et al., 2008) studies in the reservoir control literature, which have  
825 found that the objective(s) of the operator can affect which hydrologic factors are deemed  
826 most informative. However, computational constraints often require that the total set  
827 of potentially informative data be culled to a small subset of the most important vari-  
828 ables. The results of this study confirm the importance of accounting for the multi-dimensional  
829 nature of information value during this process (Denaro et al., 2017; Giuliani et al., 2015).

830 A limitation of this study is the implicit assumption of stationarity embedded in  
831 the hydro-financial simulation model adopted from Hamilton et al. (2020). Despite this  
832 fact, Figure 4 suggests that the EMODPS-derived policies trained on a stationary MC  
833 ensemble can perform relatively well across a wide range of potential outcomes, many  
834 of which are extreme compared to historical data. Additionally, the present study con-  
835 cerns purely financial decisions on relatively short time scales, for which interannual cli-  
836 mate variability is expected to overwhelm longer-term non-stationarity (Lehner et al.,  
837 2020). The reader is referred to Hamilton et al. (2020) for further discussion of these is-  
838 sues. Nonetheless, future studies should consider a broader analysis of the impacts of chang-  
839 ing climate, markets, etc., on the robustness of adaptive financial risk management strate-  
840 gies for hydropower production. This would be especially important if the framework  
841 proposed here were to be combined with dynamic infrastructure investments (Haasnoot,  
842 Kwakkel, Walker, & ter Maat, 2013; Kwakkel, Haasnoot, & Walker, 2015; Zeff, Herman,  
843 Reed, & Characklis, 2016). For large, irreversible decisions such as infrastructure devel-  
844 opment, medium- to long-term uncertainties become increasingly important (Doss-Gollin,  
845 Farnham, Steinschneider, & Lall, 2019). Future work should focus on integrating addi-  
846 tional sources of information regarding climate, power markets, etc. Statistical learning  
847 approaches can be used to update decision-making based on evolving beliefs about the  
848 non-stationary hydro-financial system (Fletcher, Lickley, & Strzepek, 2019; Fletcher et  
849 al., 2017; Herman, Quinn, Steinschneider, Giuliani, & Fletcher, 2020). Additionally, sce-  
850 nario discovery approaches can be used to search for financial risk management strate-  
851 gies that perform satisfactorily across a wide range of (perhaps deeply) uncertain fac-  
852 tors (Bryant & Lempert, 2010; Herman, Reed, Zeff, & Characklis, 2015; Kasprzyk, Nataraj,  
853 Reed, & Lempert, 2013; Lempert, 2002; Quinn, Hadjimichael, Reed, & Steinschneider,  
854 2020).

## 6 Conclusions

A substantial body of literature has emerged around optimal control of water reservoir systems in the face of hydrologic uncertainty (Macian-Sorribes & Pulido-Velazquez, 2019). Evolutionary multi-objective direct policy search (EMODPS) has emerged as an especially powerful tool for overcoming the simultaneous curses of dimensionality, modeling, and multiple objectives that are characteristic of problems in the field (Giuliani et al., 2016, 2018). This paper demonstrates that the same properties of EMODPS that make it ideal for optimal reservoir control problems also make it well suited for the complex, multi-objective financial risk management problems faced by water-reliant organizations as a result of hydrologic variability. The methodology is applied in the context of the hydrologic financial risk faced by the Power Enterprise of the San Francisco Public Utilities Commission (SFPUC), an electricity producer relying primarily on hydropower from a snow-dominated watershed. EMODPS is used to develop control policies that dynamically balance financial hedging, cash reserves, and debt, based on changing conditions within the model. Performance is quantified based on four conflicting objectives: expected annualized cash flow, 95th percentile maximum debt, expected hedging frequency, and expected maximum reserve fund balance. The first two objectives represent the classic return vs. risk tradeoff in finance, while the second two objectives represent a decision-maker's preferences for using one risk management instrument over another based on an organization's individual circumstances. By utilizing real-time model state information when making decisions, the dynamic policies produced by EMODPS are found to significantly outperform policies produced under a more static control formulation akin to those commonly used for financial risk management in the water resources literature. *A posteriori* visual analytics and information theoretic sensitivity analysis can be used to help decision-makers better understand how the complex, non-linear operating policies adapt to real-time information when making decisions. The methodology developed in this paper should help stakeholders to better understand the dynamic relationships between hydrology, decision-making, and financial outcomes, and facilitate more knowledgeable and effective management of hydrologic financial risks. This work is applicable to other electric utilities that rely on hydropower, as well as other stakeholders for whom environmental variability poses a significant financial risk, such as water utilities, agricultural producers, and renewable energy developers.

## 887 **Acknowledgments**

888 Funding for this work was provided by the National Science Foundation (NSF), Inno-  
 889 vations at the Nexus of Food-Energy-Water Systems, Track 2 (Award 1639268). The au-  
 890 thors would like to thank Rohini Gupta for helpful discussion and code review, as well  
 891 as Alexis Dufour and Darryl Dunn from the San Francisco Public Utilities Commission  
 892 (SFPUC) for helpful discussion and data provision. The views expressed in this work  
 893 represent those of the authors and do not necessarily reflect the views or policies of the  
 894 NSF or SFPUC. All code and data for this project, including figure generation, are avail-  
 895 able in a live repository ([https://github.com/ahamilton144/hamilton-2021-EMODPS](https://github.com/ahamilton144/hamilton-2021-EMODPS-financial-risk)  
 896 [-financial-risk](https://github.com/ahamilton144/hamilton-2021-EMODPS-financial-risk)) and a permanent archive (<https://doi.org/10.5281/zenodo.4499088>).

## 897 **References**

- 898 Auder, B., & Iooss, B. (2009). Global sensitivity analysis based on entropy. *Safety,*  
 899 *Reliability and Risk Analysis: Theory, Methods and Applications - Proceedings*  
 900 *of the Joint ESREL and SRA-Europe Conference, 3*, 2107–2115.
- 901 Bank, M., & Wiesner, R. (2010). The Use of Weather Derivatives by Small-and  
 902 Medium- Sized Enterprises: Reasons and Obstacles. *Journal of Small Business*  
 903 *and Entrepreneurship, 23*(4), 581–600. doi: 10.1080/08276331.2010.10593503
- 904 Basdekas, L. (2014). Is multiobjective optimization ready for water resources prac-  
 905 titioners? Utility’s drought policy investigation. *Journal of Water Resources*  
 906 *Planning and Management, 140*(3), 275–276. doi: 10.1061/(ASCE)WR.1943  
 907 -5452.0000415
- 908 Baumol, W. J. (1952). The Transactions Demand for Cash : An Inventory Theoretic  
 909 Approach. *The Quarterly Journal of Economics, 66*(4), 545–556.
- 910 Behzadian, M., Khanmohammadi Otaghsara, S., Yazdani, M., & Ignatius, J. (2012,  
 911 12). *A state-of the-art survey of TOPSIS applications* (Vol. 39) (No. 17). Else-  
 912 vier Ltd. doi: 10.1016/j.eswa.2012.05.056
- 913 Bellman, R. (1957). *Dynamic Programming*. Princeton, NJ: Princeton University  
 914 Press.
- 915 Bertsekas, D. P. (2019). *Reinforcement learning and optimal control*. Nashua, NH:  
 916 Athena Scientific.
- 917 Blomfield, A., & Plummer, J. (2014). The allocation and documentation of hydro-  
 918 logical risk. *Hydropower & Dams*(5), 94–108. doi: 10.1093/rfs/15.4.1283

- 919 Bodnar, G. M., Giambona, E., Graham, J. R., & Harvey, C. R. (2019). A view in-  
 920 side corporate risk management. *Management Science*, *65*(11), 5001–5026. doi:  
 921 10.1287/mnsc.2018.3081
- 922 Bolton, P., Chen, H., & Wang, N. (2011). A Unified Theory of Tobin’s q, Corpo-  
 923 rate Investment, Financing, and Risk Management. *Journal of Finance*, *66*(5),  
 924 1545–1578. doi: 10.1111/j.1540-6261.2011.01681.x
- 925 Borgonovo, E. (2007). A new uncertainty importance measure. *Reliability Engineer-*  
 926 *ing and System Safety*, *92*(6), 771–784. doi: 10.1016/j.res.2006.04.015
- 927 Borgonovo, E., Castaings, W., & Tarantola, S. (2011). Moment Independent Impor-  
 928 tance Measures: New Results and Analytical Test Cases. *Risk Analysis*, *31*(3),  
 929 404–428. doi: 10.1111/j.1539-6924.2010.01519.x
- 930 Brown, C. M., & Carriquiry, M. (2007). Managing hydroclimatological risk to water  
 931 supply with option contracts and reservoir index insurance. *Water Resources*  
 932 *Research*, *43*, W11423. doi: 10.1029/2007WR006093
- 933 Brown, C. M., Lund, J. R., Cai, X., Reed, P. M., Zagona, E. A., Ostfeld, A., ...  
 934 Brekke, L. (2015). The future of water resources systems analysis: Toward  
 935 a scientific framework for sustainable water management. *Water Resources*  
 936 *Research*, *51*(8), 6110–6124. doi: 10.1002/2015WR017114
- 937 Bryant, B. P., & Lempert, R. J. (2010). Thinking inside the box: A participatory,  
 938 computer-assisted approach to scenario discovery. *Technological Forecast-*  
 939 *ing and Social Change*, *77*(1), 34–49. Retrieved from [http://dx.doi.org/](http://dx.doi.org/10.1016/j.techfore.2009.08.002)  
 940 [10.1016/j.techfore.2009.08.002](http://dx.doi.org/10.1016/j.techfore.2009.08.002) doi: 10.1016/j.techfore.2009.08.002
- 941 Cai, X., Wallington, K., Shafiee-Jood, M., & Marston, L. (2018). Understand-  
 942 ing and managing the food-energy-water nexus - opportunities for water re-  
 943 sources research. *Advances in Water Resources*, *111*(November), 259–273.  
 944 Retrieved from <https://doi.org/10.1016/j.advwatres.2017.11.014> doi:  
 945 10.1016/j.advwatres.2017.11.014
- 946 Caniglia, G., Luederitz, C., von Wirth, T., Fazey, I., Martín-López, B., Hon-  
 947 drila, K., ... Lang, D. J. (2020). A pluralistic and integrated approach to  
 948 action-oriented knowledge for sustainability. *Nature Sustainability*. doi:  
 949 10.1038/s41893-020-00616-z
- 950 Cash, D. W., Clark, W. C., Alcock, F., Dickson, N. M., Eckley, N., Guston, D. H.,  
 951 ... Mitchell, R. B. (2003). Knowledge systems for sustainable development.

- 952 *Proceedings of the National Academy of Sciences of the United States of Amer-*  
 953 *ica*, 100(14), 8086–8091. doi: 10.1073/pnas.1231332100
- 954 Castelletti, A., Pianosi, F., & Soncini-Sessa, R. (2008). Water reservoir control un-  
 955 der economic, social and environmental constraints. *Automatica*, 44(6), 1595–  
 956 1607. doi: 10.1016/j.automatica.2008.03.003
- 957 Castelletti, A., & Soncini-Sessa, R. (2006). A procedural approach to strengthening  
 958 integration and participation in water resource planning. *Environmental Mod-*  
 959 *elling and Software*, 21(10), 1455–1470. doi: 10.1016/j.envsoft.2005.07.013
- 960 Castelvechi, D. (2016). Can we open the black box of AI? *Nature*, 538(7623), 20–  
 961 23. doi: 10.1038/538020a
- 962 Ceres. (2017). *The Investor Water Toolkit* (Tech. Rep.). Boston, MA: Ceres.  
 963 Retrieved from [https://www.ceres.org/sites/default/files/reports/](https://www.ceres.org/sites/default/files/reports/Ceres_InvestWaterToolkit.pdf)  
 964 [Ceres\\_InvestWaterToolkit.pdf](https://www.ceres.org/sites/default/files/reports/Ceres_InvestWaterToolkit.pdf)
- 965 Chapman, T. A., & Breeding, J. M. (2014). *U.S. Public Finance Waterworks,*  
 966 *Sanitary Sewer, and Drainage Utility Systems: Methodology And Assump-*  
 967 *tions* (Tech. Rep.). Standard and Poor’s Ratings Services. Retrieved from  
 968 [https://www.spratings.com/documents/20184/908554/US\\_PF\\_Event](https://www.spratings.com/documents/20184/908554/US_PF_Event_RFCRndTblsJan2015_Article1/30d125eb-1066-4730-8ab1-f2cd6a6d6f9a)  
 969 [\\_RFCRndTblsJan2015\\_Article1/30d125eb-1066-4730-8ab1-f2cd6a6d6f9a](https://www.spratings.com/documents/20184/908554/US_PF_Event_RFCRndTblsJan2015_Article1/30d125eb-1066-4730-8ab1-f2cd6a6d6f9a)
- 970 Coello Coello, C. A., Lamont, G. B., & Van Veldhuizen, D. A. (2007). *Evolution-*  
 971 *ary Algorithms for Solving Multi-Objective Problems* (2nd ed.). New York, NY:  
 972 Springer Science+Business Media, LLC. doi: 10.1046/j.1365-2672.2000.00969  
 973 .x
- 974 da Costa Moraes, M. B., & Nagano, M. S. (2013). Cash Management Policies By  
 975 Evolutionary Models: A Comparison Using The Miller-Orr Model. *Journal*  
 976 *of Information Systems and Technology Management*, 10(3), 561–576. doi:  
 977 10.4301/s1807-17752013000300006
- 978 da Costa Moraes, M. B., Nagano, M. S., & Sobreiro, V. A. (2015). Stochastic cash  
 979 flow management models: A literature review since the 1980s. *Decision Engi-*  
 980 *neering*(January), 322. doi: 10.1007/978-3-319-11949-6{-}2
- 981 Denaro, S., Anghileri, D., Giuliani, M., & Castelletti, A. (2017). Informing the  
 982 operations of water reservoirs over multiple temporal scales by direct use of  
 983 hydro-meteorological data. *Advances in Water Resources*, 103, 51–63. doi:  
 984 10.1016/j.advwatres.2017.02.012

- 985 Denaro, S., Castelletti, A., Giuliani, M., & Characklis, G. (2020). Insurance Port-  
986 folio Diversification Through Bundling for Competing Agents Exposed to  
987 Uncorrelated Drought and Flood Risks. *Water Resources Research*, *56*(5),  
988 1–20. doi: 10.1029/2019WR026443
- 989 Desreumaux, Q., Côté, P., & Leconte, R. (2018). Comparing model-based and  
990 model-free streamflow simulation approaches to improve hydropower reservoir  
991 operations. *Journal of Water Resources Planning and Management*, *144*(3),  
992 1–10. doi: 10.1061/(ASCE)WR.1943-5452.0000860
- 993 di Pierro, F., Khu, S. T., & Savić, D. A. (2007). An investigation on preference or-  
994 der ranking scheme for multiobjective evolutionary optimization. *IEEE Trans-*  
995 *actions on Evolutionary Computation*, *11*(1), 17–45. doi: 10.1109/TEVC.2006  
996 .876362
- 997 Disatnik, D., Duchin, R. A. N., & Schmidt, B. (2014). Cash Flow Hedging and Liq-  
998 uidity Choices. *Review of Finance*, *18*, 715–748. doi: 10.1093/rof/rft006
- 999 D’Odorico, P., Davis, K. F., Rosa, L., Carr, J. A., Chiarelli, D., Dell’Angelo, J.,  
1000 ... Rulli, M. C. (2018). The Global Food-Energy-Water Nexus. *Reviews of*  
1001 *Geophysics*, *56*(3), 456–531. doi: 10.1029/2017RG000591
- 1002 Doss-Gollin, J., Farnham, D. J., Steinschneider, S., & Lall, U. (2019). Robust  
1003 Adaptation to Multiscale Climate Variability. *Earth’s Future*, *7*(7), 734–747.  
1004 Retrieved from [https://onlinelibrary.wiley.com/doi/abs/10.1029/](https://onlinelibrary.wiley.com/doi/abs/10.1029/2019EF001154)  
1005 [2019EF001154](https://onlinelibrary.wiley.com/doi/abs/10.1029/2019EF001154) doi: 10.1029/2019EF001154
- 1006 Fletcher, S., Lickley, M., & Strzepek, K. (2019). Learning about climate change  
1007 uncertainty enables flexible water infrastructure planning. *Nature Communica-*  
1008 *tions*, *10*(1), 1–11. Retrieved from [http://dx.doi.org/10.1038/s41467-019-](http://dx.doi.org/10.1038/s41467-019-09677-x)  
1009 [-09677-x](http://dx.doi.org/10.1038/s41467-019-09677-x) doi: 10.1038/s41467-019-09677-x
- 1010 Fletcher, S., Miotti, M., Swaminathan, J., Klemun, M., Strzepek, K., & Siddiqi,  
1011 A. (2017). Water supply infrastructure planning decision-making frame-  
1012 work to classify multiple uncertainties and evaluate flexible design. *Journal*  
1013 *of Water Resources Planning and Management*, *143*(10), 04017061. doi:  
1014 10.1061/(ASCE)WR.1943-5452.0000823
- 1015 Flood, M. D., Lemieux, V. L., Varga, M., & William Wong, B. L. (2016). The appli-  
1016 cation of visual analytics to financial stability monitoring. *Journal of Financial*  
1017 *Stability*, *27*, 180–197. doi: 10.1016/j.jfs.2016.01.006

- 1018 Foster, B. T., Kern, J. D., & Characklis, G. W. (2015). Mitigating hydrologic fi-  
1019 nancial risk in hydropower generation using index-based financial instruments.  
1020 *Water Resources and Economics*, *10*, 45–67. doi: 10.1016/j.wre.2015.04.001
- 1021 Froot, K., Scharfstein, D., & Stein, J. (1993). Risk Management: Coordinating  
1022 Corporate Investment and Financing Policies. *The Journal of Finance*, *48*(5),  
1023 1629–1648. doi: 10.1111/j.1540-6261.1993.tb05123.x
- 1024 Giuliani, M., & Castelletti, A. (2016). Is robustness really robust? How different def-  
1025 initions of robustness impact decision-making under climate change. *Climatic*  
1026 *Change*, *135*(3-4), 409–424. doi: 10.1007/s10584-015-1586-9
- 1027 Giuliani, M., Castelletti, A., Pianosi, F., Mason, E., & Reed, P. M. (2016). Curses,  
1028 tradeoffs, and scalable management: Advancing evolutionary multiobjec-  
1029 tive direct policy search to improve water reservoir operations. *Journal*  
1030 *of Water Resources Planning and Management*, *142*(2), 04015050. doi:  
1031 10.1061/(ASCE)WR.1943-5452.0000570
- 1032 Giuliani, M., Herman, J. D., Castelletti, A., & Reed, P. (2014). Many-objective  
1033 reservoir policy identification and refinement to reduce policy inertia and my-  
1034 opia in water management. *Water Resources Research*, *50*(4), 3355–3377. doi:  
1035 10.1002/2013WR014700
- 1036 Giuliani, M., Mason, E., Castelletti, A., Pianosi, F., & Soncini-Sessa, R. (2014).  
1037 Universal approximators for direct policy search in multi-purpose water reser-  
1038 voir management: A comparative analysis. In *Proceedings of the 19th world*  
1039 *congress, international federation of automatic control* (Vol. 19, pp. 6234–  
1040 6239). Cape Town, South Africa. doi: 10.3182/20140824-6-za-1003.01962
- 1041 Giuliani, M., Pianosi, F., & Castelletti, A. (2015). Making the most of data:  
1042 An information selection and assessment framework to improve water sys-  
1043 tems operations. *Water Resources Research*, *51*(11), 9073–9093. doi:  
1044 10.1002/2015WR017044
- 1045 Giuliani, M., Quinn, J. D., Herman, J. D., Castelletti, A., & Reed, P. M. (2018).  
1046 Scalable Multiobjective Control for Large-Scale Water Resources Systems un-  
1047 der Uncertainty. *IEEE Transactions on Control Systems Technology*, *26*(4),  
1048 1492–1499. doi: 10.1109/TCST.2017.2705162
- 1049 Graham, N. E., & Georgakakos, K. P. (2010). Toward understanding the value of  
1050 climate information for multiobjective reservoir management under present

- 1051 and future climate and demand scenarios. *Journal of Applied Meteorology and*  
1052 *Climatology*, 49(4), 557–573. doi: 10.1175/2009JAMC2135.1
- 1053 Gupta, R. S., Hamilton, A. L., Reed, P. M., & Characklis, G. W. (2020). Can mod-  
1054 ern multi-objective evolutionary algorithms discover high-dimensional financial  
1055 risk portfolio tradeoffs for snow-dominated water-energy systems? *Advances in*  
1056 *Water Resources*, 145, 103718. Retrieved from [https://doi.org/10.1016/](https://doi.org/10.1016/j.advwatres.2020.103718)  
1057 [j.advwatres.2020.103718](https://doi.org/10.1016/j.advwatres.2020.103718) doi: 10.1016/j.advwatres.2020.103718
- 1058 Haasnoot, M., Kwakkel, J. H., Walker, W. E., & ter Maat, J. (2013). Dy-  
1059 namic adaptive policy pathways: A method for crafting robust deci-  
1060 sions for a deeply uncertain world. *Global Environmental Change*. doi:  
1061 10.1016/j.gloenvcha.2012.12.006
- 1062 Hadka, D., & Reed, P. (2012). Diagnostic assessment of search controls and fail-  
1063 ure modes in many-objective evolutionary optimization. *Evolutionary Compu-*  
1064 *tation*, 20(3), 423–452. doi: 10.1162/EVCO{\\_}a{\\_}00053
- 1065 Hadka, D., & Reed, P. (2013). Borg: An auto-adaptive many-objective evolution-  
1066 ary computing framework. *Evolutionary Computation*, 21(2), 231–259. doi: 10  
1067 .1162/EVCO{\\_}a{\\_}00075
- 1068 Hamilton, A. L., Characklis, G. W., & Reed, P. M. (2020). Managing financial risk  
1069 tradeoffs for hydropower generation using snowpack-based index contracts.  
1070 *Water Resources Research*, 56, e2020WR027212. doi: 10.1029/2020wr027212
- 1071 Hejazi, M. I., Cai, X., & Ruddell, B. L. (2008). The role of hydrologic information in  
1072 reservoir operation - Learning from historical releases. *Advances in Water Re-*  
1073 *sources*, 31(12), 1636–1650. doi: 10.1016/j.advwatres.2008.07.013
- 1074 Herman, J. D., Quinn, J. D., Steinschneider, S., Giuliani, M., & Fletcher, S. (2020).  
1075 Climate adaptation as a control problem: Review and perspectives on dynamic  
1076 water resources planning under uncertainty. *Water Resources Research*, 56,  
1077 e24389. doi: 10.1029/2019wr025502
- 1078 Herman, J. D., Reed, P. M., Zeff, H. B., & Characklis, G. W. (2015). How should  
1079 robustness be defined for water systems planning under change? *Journal*  
1080 *of Water Resources Planning and Management*, 141(10), 04015012. doi:  
1081 10.1061/(ASCE)WR.1943-5452.0000509
- 1082 Herman, J. D., Zeff, H. B., Reed, P. M., & Characklis, G. W. (2014). Beyond op-  
1083 timality: Multistakeholder robustness tradeoffs for regional water portfolio

- 1084 planning under deep uncertainty. *Water Resources Research*, 50(10), 7692–  
 1085 7713. doi: 10.1002/2014WR015338
- 1086 Hughes, J., & Leurig, S. (2013). *Assessing water system revenue risk: Considera-*  
 1087 *tions for market analysts* (Tech. Rep.). Boston, MA: Ceres and Environmental  
 1088 Finance Center at the University of North Carolina at Chapel Hill.
- 1089 Hughes, J., Tiger, M., Eskaf, S., Berahzer, S. I., Royster, S., Boyle, C., & Batten, D.  
 1090 (2014). *Defining a Resilient Business Model for Water Utilities* (Tech. Rep.).  
 1091 Water Research Foundation. Retrieved from [https://efc.sog.unc.edu/](https://efc.sog.unc.edu/sites/default/files/4366_Exec_Summary_0.pdf)  
 1092 [sites/default/files/4366\\_Exec\\_Summary\\_0.pdf](https://efc.sog.unc.edu/sites/default/files/4366_Exec_Summary_0.pdf)
- 1093 Huskova, I., Matrosov, E. S., Harou, J. J., Kasprzyk, J. R., & Lambert, C. (2016).  
 1094 Screening robust water infrastructure investments and their trade-offs under  
 1095 global change: A London example. *Global Environmental Change*, 41, 216–  
 1096 227. Retrieved from <http://dx.doi.org/10.1016/j.gloenvcha.2016.10.007>  
 1097 doi: 10.1016/j.gloenvcha.2016.10.007
- 1098 Iooss, B., & Lemaître, P. (2015). A review on global sensitivity analysis methods.  
 1099 *Operations Research/ Computer Science Interfaces Series*, 59, 101–122. doi:  
 1100 10.1007/978-1-4899-7547-8{\\_}5
- 1101 Kasprzyk, J. R., Nataraj, S., Reed, P. M., & Lempert, R. J. (2013). Many ob-  
 1102 jective robust decision making for complex environmental systems under-  
 1103 going change. *Environmental Modelling and Software*, 42, 55–71. doi:  
 1104 10.1016/j.envsoft.2012.12.007
- 1105 Kasprzyk, J. R., Reed, P. M., Characklis, G. W., & Kirsch, B. R. (2012). Many-  
 1106 objective de Novo water supply portfolio planning under deep uncertainty. *En-*  
 1107 *vironmental Modelling and Software*, 34, 87–104. doi: 10.1016/j.envsoft.2011  
 1108 .04.003
- 1109 Kasprzyk, J. R., Reed, P. M., & Hadka, D. M. (2016). Battling Arrow’s Paradox  
 1110 to discover robust water management alternatives. *Journal of Water Resources*  
 1111 *Planning and Management*, 142(2), 04015053. doi: 10.1061/(ASCE)WR.1943  
 1112 -5452.0000572
- 1113 Kollat, J. B., & Reed, P. (2007). A framework for Visually Interactive Decision-  
 1114 making and Design using Evolutionary Multi-objective Optimization  
 1115 (VIDEO). *Environmental Modelling and Software*, 22(12), 1691–1704. doi:  
 1116 10.1016/j.envsoft.2007.02.001

- 1117 Koutsoyiannis, D., & Economou, A. (2003). Evaluation of the parameterization-  
 1118 simulation-optimization approach for the control of reservoir systems. *Water*  
 1119 *Resources Research*, 39(6). doi: 10.1029/2003WR002148
- 1120 Kouwenberg, R., & Zenios, S. A. (2008). Stochastic Programming Models for Asset  
 1121 Liability Management. In *Handbook of asset and liability management* (Vol. 1,  
 1122 pp. 253–303). doi: 10.1016/B978-044453248-0.50012-5
- 1123 Krause, T. A., & Tse, Y. (2016). Risk management and firm value: Recent theory  
 1124 and evidence. *International Journal of Accounting and Information Manage-*  
 1125 *ment*, 24(1), 56–81. doi: 10.1108/IJAIM-05-2015-0027
- 1126 Krzykacz-Hausmann, B. (2001). Epistemic sensitivity analysis based on the concept  
 1127 of entropy. In P. Prado & R. Bolado (Eds.), *Proceedings of samo* (pp. 31–35).  
 1128 Madrid: CIEMAT.
- 1129 Kwakkel, J. H., Haasnoot, M., & Walker, W. E. (2015). Developing dynamic  
 1130 adaptive policy pathways: a computer-assisted approach for developing  
 1131 adaptive strategies for a deeply uncertain world. *Climatic Change*. doi:  
 1132 10.1007/s10584-014-1210-4
- 1133 Labadie, J. W. (2004). Optimal operation of multireservoir systems: State-of-the-art  
 1134 review. *Journal of Water Resources Planning and Management*, 130(2), 93–  
 1135 111. doi: 10.1061/(ASCE)0733-9496(2004)130:2(93)
- 1136 Larson, W. M., Freedman, P. L., Passinsky, V., Grubb, E., & Adriaens, P. (2012).  
 1137 Mitigating corporate water risk: Financial market tools and supply manage-  
 1138 ment strategies. *Water Alternatives*, 5(3), 582–602.
- 1139 Lehner, F., Deser, C., Maher, N., Marotzke, J., Fischer, E., Brunner, L., ...  
 1140 Hawkins, E. (2020). Partitioning climate projection uncertainty with mul-  
 1141 tiple Large Ensembles and CMIP5/6. *Earth System Dynamics*, 11, 491–508.  
 1142 doi: 10.5194/esd-2019-93
- 1143 Lempert, R. J. (2002). A new decision sciences for complex systems. *Proceedings of*  
 1144 *the National Academy of Sciences of the United States of America*, 99(SUPPL.  
 1145 3), 7309–7313. doi: 10.1073/pnas.082081699
- 1146 Leurig, S. (2010). *The Ripple Effect: Water risk in the municipal bond market*  
 1147 (Tech. Rep.). Boston, MA: Ceres. Retrieved from [https://www.ceres.org/  
 1148 resources/reports/ripple-effect-water-risk-municipal-bond-market](https://www.ceres.org/resources/reports/ripple-effect-water-risk-municipal-bond-market)
- 1149 Liu, H., Chen, W., & Sudjianto, A. (2006). Relative entropy based method for prob-

- 1150           abilistic sensitivity analysis in engineering design. *Journal of Mechanical De-*  
1151           *sign*, 128(2), 326–336.
- 1152   Macian-Sorribes, H., & Pulido-Velazquez, M.   (2019).    Inferring efficient operat-  
1153           ing rules in multireservoir water resource systems: A review. *Wiley Interdisci-*  
1154           *plinary Reviews: Water*, 7(1), e1400. doi: 10.1002/wat2.1400
- 1155   Maestro, T., Barnett, B. J., Coble, K. H., Garrido, A., & Bielza, M. (2016). Drought  
1156           index insurance for the Central Valley Project in California. *Applied Economic*  
1157           *Perspectives and Policy*, 38(3), 521–545. doi: 10.1093/aep/38(3)/ppw013
- 1158   Maier, H. R., Kapelan, Z., Kasprzyk, J., Kollat, J., Matott, L. S., Cunha, M. C.,  
1159           ... Reed, P. M.   (2014).    Evolutionary algorithms and other metaheuris-  
1160           tics in water resources: Current status, research challenges and future  
1161           directions. *Environmental Modelling and Software*, 62, 271–299.    doi:  
1162           10.1016/j.envsoft.2014.09.013
- 1163   Markowitz, H. (1952). Portfolio Selection. *The Journal of Finance*, 7(1), 77–91.
- 1164   Meyer, E. S., Characklis, G. W., Brown, C., & Moody, P.   (2016).    Hedging the fi-  
1165           nancial risk from water scarcity for Great Lakes shipping. *Water Resources Re-*  
1166           *search*, 52(1), 227–245. doi: 10.1002/2015WR017855
- 1167   Miller, M. M., & Orr, D. (1966). A Model of the Demand for Money by Firms. *The*  
1168           *Quarterly Journal of Economics*, 80(3), 413–435.
- 1169   Moody’s Investors Service. (2011). *Rating Methodology: U.S. Public Power Electric*  
1170           *Utilities with Generation Ownership Exposure* (Tech. Rep.).
- 1171   Moody’s Investors Service. (2019). *Bonneville Power Administration , OR: Credit*  
1172           *Update Following Rating Affirmation and Change in Outlook to Negative*  
1173           (Tech. Rep.).
- 1174   Mortensen, E., & Block, P.   (2018).    ENSO Index-Based Insurance for Agricul-  
1175           tural Protection in Southern Peru. *Geosciences*, 8(2), 64.    doi: 10.3390/  
1176           geosciences8020064
- 1177   Mulvey, J. M.   (2001).    Introduction to financial optimization: Mathematical Pro-  
1178           gramming Special Issue. *Mathematical Programming*, 89(2), 205–216. doi: 10  
1179           .1007/pl00011395
- 1180   Nicklow, J., Reed, P., Savic, D., Dessalegne, T., Harrell, L., Chan-Hilton, A.,  
1181           ... Zechman, E.   (2010).    State of the Art for Genetic Algorithms and  
1182           Beyond in Water Resources Planning and Management. *Journal of Wa-*

- 1183 *ter Resources Planning and Management*, 136(4), 412–432. doi: 10.1061/  
 1184 ASCEWR.1943-5452.0000053
- 1185 Pappenberger, F., Beven, K. J., Ratto, M., & Matgen, P. (2008). Multi-method  
 1186 global sensitivity analysis of flood inundation models. *Advances in Water Re-*  
 1187 *sources*, 31(1), 1–14. doi: 10.1016/j.advwatres.2007.04.009
- 1188 Pardalos, P. M., Sandström, M., & Zopounidis, C. (1994). On the use of optimiza-  
 1189 tion models for portfolio selection: A review and some computational results.  
 1190 *Computational Economics*, 7(4), 227–244. doi: 10.1007/BF01299454
- 1191 Pérez-González, F., & Yun, H. (2013). Risk management and firm value: Evidence  
 1192 from weather derivatives. *Journal of Finance*, 68(5), 2143–2176. doi: 10.1111/  
 1193 jofi.12061
- 1194 Pianosi, F., Beven, K., Freer, J., Hall, J. W., Rougier, J., Stephenson, D. B., & Wa-  
 1195 gener, T. (2016). Sensitivity analysis of environmental models: A systematic  
 1196 review with practical workflow. *Environmental Modelling and Software*, 79,  
 1197 214–232. doi: 10.1016/j.envsoft.2016.02.008
- 1198 Ponsich, A., Jaimes, A. L., & Coello Coello, C. A. (2013). A survey on multiobjec-  
 1199 tive evolutionary algorithms for the solution of the portfolio optimization prob-  
 1200 lem and other finance and economics applications. *IEEE Transactions on Evo-*  
 1201 *lutionary Computation*, 17(3), 321–344. doi: 10.1109/TEVC.2012.2196800
- 1202 Powell, W. B. (2019). A unified framework for stochastic optimization. *European*  
 1203 *Journal of Operational Research*, 275(3), 795–821. doi: 10.1016/j.ejor.2018.07  
 1204 .014
- 1205 Quinn, J. D., Hadjimichael, A., Reed, P. M., & Steinschneider, S. (2020, 11). Can  
 1206 Exploratory Modeling of Water Scarcity Vulnerabilities and Robustness Be  
 1207 Scenario Neutral? *Earth's Future*, 8(11). doi: 10.1029/2020EF001650
- 1208 Quinn, J. D., Reed, P. M., Giuliani, M., & Castelletti, A. (2017). Rival framings:  
 1209 A framework for discovering how problem formulation uncertainties shape risk  
 1210 management trade-offs in water resources systems. *Water Resources Research*,  
 1211 53(8), 7208–7233. doi: 10.1002/2017WR020524
- 1212 Quinn, J. D., Reed, P. M., Giuliani, M., & Castelletti, A. (2019). What Is Control-  
 1213 ling Our Control Rules? Opening the Black Box of Multireservoir Operating  
 1214 Policies Using Time-Varying Sensitivity Analysis. *Water Resources Research*,  
 1215 55(7), 5962–5984. doi: 10.1029/2018WR024177

- 1216 Quinn, J. D., Reed, P. M., Giuliani, M., Castelletti, A., Oyler, J. W., & Nicholas,  
 1217 R. E. (2018). Exploring How Changing Monsoonal Dynamics and Hu-  
 1218 man Pressures Challenge Multireservoir Management for Flood Protection,  
 1219 Hydropower Production, and Agricultural Water Supply. *Water Resources*  
 1220 *Research*, *54*(7), 4638–4662. doi: 10.1029/2018WR022743
- 1221 Quinn, J. D., Reed, P. M., & Keller, K. (2017). Direct policy search for ro-  
 1222 bust multi-objective management of deeply uncertain socio-ecological tip-  
 1223 ping points. *Environmental Modelling and Software*, *92*, 125–141. doi:  
 1224 10.1016/j.envsoft.2017.02.017
- 1225 Rampini, A. A., Sufi, A., & Viswanathan, S. (2014). Dynamic risk management.  
 1226 *Journal of Financial Economics*, *111*(2), 271–296. doi: 10.1016/j.jfineco.2013  
 1227 .10.003
- 1228 Reed, P. M., Hadka, D., Herman, J. D., Kasprzyk, J. R., & Kollat, J. B. (2013).  
 1229 Evolutionary multiobjective optimization in water resources: The past,  
 1230 present, and future. *Advances in Water Resources*, *51*, 438–456. doi:  
 1231 10.1016/j.advwatres.2012.01.005
- 1232 Rehan, R., Knight, M. A., Unger, A. J., & Haas, C. T. (2013). Development of  
 1233 a system dynamics model for financially sustainable management of mu-  
 1234 nicipal watermain networks. *Water Research*, *47*(20), 7184–7205. doi:  
 1235 10.1016/j.watres.2013.09.061
- 1236 Rehan, R., Unger, A., Knight, M. A., & Haas, C. (2015). Strategic water utility  
 1237 management and financial planning using a new system dynamics tool. *Journal*  
 1238 *- American Water Works Association*, *107*(1), E22-E36. doi: 10.5942/jawwa  
 1239 .2015.107.0006
- 1240 Rosenstein, M. T., & Barto, A. G. (2001). Robot weightlifting by direct policy  
 1241 search. In *International joint conference on artificial intelligence* (pp. 839–  
 1242 844).
- 1243 Roszkowska, E. (2011). *Multi-criteria decision making models by applying the TOP-*  
 1244 *SIS method to crisp and interval data* (Tech. Rep.). Retrieved from [http://](http://www.mcdm.ue.katowice.pl/files/papers/mcdm11(6)_11.pdf)  
 1245 [www.mcdm.ue.katowice.pl/files/papers/mcdm11\(6\)\\_11.pdf](http://www.mcdm.ue.katowice.pl/files/papers/mcdm11(6)_11.pdf)
- 1246 Salas-Molina, F., Pla-Santamaria, D., & Rodriguez-Aguilar, J. A. (2018). A multi-  
 1247 objective approach to the cash management problem. *Annals of Operations*  
 1248 *Research*, *267*(1-2), 515–529. doi: 10.1007/s10479-016-2359-1

- 1249 Saltelli, A., Tarantola, S., & Campolongo, F. (2000). Sensitivity Analysis as an Ingre-  
 1250 dient of Modeling. *Statistical Science*, *15*(4), 377–395.
- 1251 San Francisco Public Utilities Commission. (2016). *Comprehensive Annual Financial*  
 1252 *Report, Fiscal Years Ended June 30, 2016 and 2015* (Tech. Rep.).
- 1253 Savikhin, A., Lam, H. C., Fisher, B., & Ebert, D. S. (2011). An experimental study  
 1254 of financial portfolio selection with visual analytics for decision support. *Pro-*  
 1255 *ceedings of the Annual Hawaii International Conference on System Sciences*,  
 1256 1–10. doi: 10.1109/HICSS.2011.54
- 1257 Scanlon, B. R., Ruddell, B. L., Reed, P. M., Hook, R. I., Zheng, C., Tidwell,  
 1258 V. C., & Siebert, S. (2017). The food-energy-water nexus: Transform-  
 1259 ing science for society. *Water Resources Research*, *53*(5), 3550–3556. doi:  
 1260 10.1002/2017WR020889
- 1261 Shannon, C. (1948). A Mathematical Theory of Communication. *The Bell System*  
 1262 *Technical Journal*, *XXVII*(3), 379–423.
- 1263 Singh, L. A., Whittecar, W. R., Diprinzio, M. D., Herman, J. D., Ferringer, M. P.,  
 1264 & Reed, P. M. (2020). Low cost satellite constellations for nearly con-  
 1265 tinuous global coverage. *Nature Communications*, *11*(200), 1–7. doi:  
 1266 10.1038/s41467-019-13865-0
- 1267 Sodhi, M. S. (2005). LP modeling for asset-liability management: A survey of  
 1268 choices and simplifications. *Operations Research*, *53*(2), 181–196. doi:  
 1269 10.1287/opre.1040.0185
- 1270 Spronk, J., Steuer, R. E., & Zopounidis, C. (2005). Multicriteria Decision  
 1271 Aid/Analysis in Finance. In *Multiple criteria decision analysis: State of the art*  
 1272 *surveys. international series in operations research & management science, vol.*  
 1273 *78* (pp. 799–848). New York, NY: Springer.
- 1274 Steuer, R. E., Qi, Y., & Hirschberger, M. (2007). Suitable-portfolio investors, non-  
 1275 dominated frontier sensitivity, and the effect of multiple objectives on standard  
 1276 portfolio selection. *Annals of Operations Research*, *152*(1), 297–317. doi:  
 1277 10.1007/s10479-006-0137-1
- 1278 Sun, J., Fang, W., Wu, X., Lai, C. H., & Xu, W. (2011). Solving the multi-stage  
 1279 portfolio optimization problem with a novel particle swarm optimization. *Ex-*  
 1280 *pert Systems with Applications*, *38*(6), 6727–6735. doi: 10.1016/j.eswa.2010.11  
 1281 .061

- 1282 Tapia, M. G. C., & Coello Coello, C. A. (2007). Applications of multi-objective evolutionary algorithms in economics and finance: A survey. *2007 IEEE Congress on Evolutionary Computation*, 532–539. doi: 10.1109/CEC.2007.4424516
- 1283  
1284
- 1285 Tejada-Guibert, J. A., Johnson, S. A., & Stedinger, J. R. (1995). The Value of Hydrologic Information in Stochastic Dynamic Programming Models of a
- 1286  
1287  
1288  
1289  
1290  
1291  
1292  
1293  
1294  
1295  
1296  
1297  
1298  
1299  
1300  
1301  
1302  
1303  
1304  
1305  
1306  
1307  
1308  
1309  
1310  
1311  
1312  
1313  
1314
- Multireservoir System. *Water Resources Research*, 31(10), 2571–2579. doi: 10.1029/95WR02172
- Turvey, C. G. (2001). Weather Derivatives for Specific Event Risks in Agriculture. *Review of Agricultural Economics*, 23(2), 333–351. doi: 10.1111/1467-9353.00065
- Woodruff, M. J., Reed, P. M., & Simpson, T. W. (2013). Many objective visual analytics: Rethinking the design of complex engineered systems. *Structural and Multidisciplinary Optimization*, 48, 201–219. doi: 10.1007/s00158-013-0891-z
- Yeh, W. G. (1985). Reservoir Management and Operations Models: A State-of-the-Art Review. *Water Resources Research*, 21(12), 1797–1818. doi: 10.1029/WR021i012p01797
- Zatarain Salazar, J., Reed, P. M., Herman, J. D., Giuliani, M., & Castelletti, A. (2016). A diagnostic assessment of evolutionary algorithms for multi-objective surface water reservoir control. *Advances in Water Resources*, 92, 172–185. doi: 10.1016/j.advwatres.2016.04.006
- Zatarain Salazar, J., Reed, P. M., Quinn, J. D., Giuliani, M., & Castelletti, A. (2017). Balancing exploration, uncertainty and computational demands in many objective reservoir optimization. *Advances in Water Resources*, 109, 196–210. doi: 10.1016/j.advwatres.2017.09.014
- Zeff, H. B., & Characklis, G. W. (2013). Managing water utility financial risks through third-party index insurance contracts. *Water Resources Research*, 49(8), 4939–4951. doi: 10.1002/wrcr.20364
- Zeff, H. B., Herman, J. D., Reed, P. M., & Characklis, G. W. (2016). Cooperative drought adaptation: Integrating infrastructure development, conservation, and water transfers into adaptive policy pathways. *Water Resources Research*, 52(9), 7327–7346. doi: 10.1002/2016WR018771
- Zeff, H. B., Kasprzyk, J. R., Herman, J. D., Reed, P. M., & Characklis, G. W. (2014). Navigating financial and supply reliability tradeoffs in regional drought

- 1315 management portfolios. *Water Resources Research*, 50(6), 4906–4923. doi:  
1316 10.1002/2013WR015126
- 1317 Zopounidis, C., Doumpos, M., & Niklis, D. (2018). Financial decision support: an  
1318 overview of developments and recent trends. *EURO Journal on Decision Pro-*  
1319 *cesses*, 6(1-2), 63–76. doi: 10.1007/s40070-018-0078-3
- 1320 Zopounidis, C., Galariotis, E., Doumpos, M., Sarri, S., & Andriosopoulos, K. (2015).  
1321 Multiple criteria decision aiding for finance: An updated bibliographic sur-  
1322 vey. *European Journal of Operational Research*, 247(2), 339–348. doi:  
1323 10.1016/j.ejor.2015.05.032